

1. INTRODUCTION

1.1. Problem statement

There are several optimisation techniques and heuristic approaches that provide theoretically optimal, or near-optimal results for production planning problems. However, these techniques are not widely used by planning operators and supervisors [1-6]. Among the reasons put forward to explain the low utilisation of these methods, lack of credibility is most frequently cited [1,2,3,6].

In the search for more effective decision support techniques, the use of knowledge-based systems provides a feasible alternative. As discussed by O'Keefe *et al* [7] a number of knowledge-based planning systems are attacking problems that have previously been investigated in operational research, replacing or augmenting existing quantitative methods.

However, the current generation of knowledge-based planning systems has failed to grow rapidly in the manufacturing community [8]. In many cases knowledge-based planning systems are not accepted by the end-users or are not used for their intended purpose. Unsuccessful implementations typically result from a failure to meet the expected final performance specifications. Two related factors contribute to this problem:

- real-world planning problems are generally ill-structured
- the planning decision process is often pervaded with uncertainty

1.2. Proposed approach

This report describes a successfully implemented knowledge-based planning system. The system will henceforth be referred to as *the prototype*. The prototype is effectively a co-operative planning decision support system for a steelmaking plant. It generates detailed production sequences over a flexible planning horizon, based on operator input. The prototype represents a less abstract model of production planning than conventional quantitative methods.

LIST OF TABLES

<i>Table</i>	<i>Page</i>
4.1. Data Attributes for Each Steelmaking Order in the Problem Space	21
4.2. Top Level Breakdown of the External Environment for the Steelmaking Domain	22
4.3. Example of Hot-Rolling Standards	24
4.4. Casting Speed vs Slab Width at the Casters	28
5.1 The Implications of $\tau_i \in TS(1 \circ 1 \circ 1 \circ 1)$	37
9.1. Planner-Prototype Coherence in the Prioritisation of Instantiations	68
10.1. Comparative Analysis (Prototype A \leftrightarrow Prototype B)	82

LIST OF FIGURES

<i>Figure</i>	<i>Page</i>
1.1. Customer Oriented Production Planning	2
3.1. The Basic Structure of a Knowledge-Based System	14
4.1. Production Planning Activities at an Integrated Steel Works	19
4.2. The Steelmaking Process	20
4.3. Quality Precedence Constraint Specifications	27
5.1. Task Space Normalisation Schematic	35
5.2. State-Space Schematic	38
6.1. Degree of Human Intervention	44
6.2. The Semantics of the Temporal Delivery Date Constraint	47
6.3. Rule-Base Evolution (Prototype A)	49
6.4. Average Sequence Generation Time (Prototype A)	50
6.5. Conformance to Selected Planning Objectives (Prototype A)	51
7.1. The Decision Variable <i>IMMINENCE</i> and the Fuzzy Set <i>LOW</i>	55
7.2. Min-Max Method of Implication With an Output Fuzzy Region	59
7.3. Min-Max Method of Implication With an Output Singleton Geometry	61
9.1. Application of Bias (Operator Preference) on the Δt and Δw Soft Decision Variables	70
9.2. Subjective Variations on the Semantics of Δt	72
9.3. Degree of Human Intervention	73
9.4. Average Sequence generation Time	77
9.5. Conformance to Planning Objectives (Prototype B)	75
10.1. Hypothetical Schematic of a Multi-Peaked Problem Space	84

APPENDIX 7	Evaluation Data (Prototype Conformance to Selected Planning Objectives)	102
APPENDIX 8	Evaluation Data (Planner-Prototype Coherence in the Prioritisation of Instantiations)	103
APPENDIX 9	Evaluation Data (Sensitivity Analysis With Respect to Bias on the Soft Decision Variables Δt and Δs)	104

8.4.	The Search Method	67
9.	EVALUATION (PROTOTYPE B)	68
9.1.	Evaluations Prior to Implementation	68
	9.1.1. Response to Operator Preferences	69
	9.1.2. Sensitivity Analysis With Respect to Fuzzy Set Morphology	69
9.2.	Evaluations After Implementation	73
	9.2.1. Degree of Human Intervention	73
	9.2.2. Sequence Generation Time	73
	9.2.3. Conformance to Planning Objectives	74
	9.2.4. Interpretation of Results	75
10.	SUMMARY AND CONCLUSIONS	79
10.1.	Acceptance of the Prototype	79
	10.1.1. Prototype A	79
	10.1.2. Prototype B	80
10.2.	Comparative Analysis (Prototype A ↔ Prototype B)	81
10.3.	Additional Benefits (Fuzzy Decision Modeling)	82
	10.3.1. Reduced Model Complexity	82
	10.3.2. Improved System Maintainability	82
	10.3.3. The Ability to Model Conflicting Knowledge	82
	10.3.4. Improved Handling of Uncertainty	82
10.4.	Pertinent Areas for Further Research and Development	84
	10.4.1. Improvement of the Search Method	84
	10.4.2. Expansion of the Fuzzy Knowledge Base	85
	REFERENCES	86
	APPENDICES	93
APPENDIX 1	Response From the Planning Supervisor	93
APPENDIX 2	The Visualisation Interface	95
APPENDIX 3	The Meta-Plan Generation Interface	96
APPENDIX 4	The Fuzzy Constraint Spaces (Interactive Screen)	97
APPENDIX 5	Fuzzy Propositional Table (Interactive Screen)	98
APPENDIX 6	Singleton Geometry Output Space	101

6.	EVALUATION (PROTOTYPE A)	43
6.1.	Degree of Human Intervention	43
	6.1.1. Ambiguity	45
	6.1.2. Vagueness	46
6.2.	Rule-Base Evolution	48
6.3.	Sequence Generation Time	50
6.4.	Conformance to Planning Objectives	51
6.5.	Interpretation of Results	52
7.	A FUZZY SET METHODOLOGY FOR SOFT CONSTRAINT INFERENCE	54
7.1.	Approximate Reasoning	54
7.2.	Representing Soft Constraints as Fuzzy Sets	54
	7.2.1. Fuzzy Sets	54
	7.2.2. Decision Variables	56
7.3.	Fuzzy Inference	56
	7.3.1. Fuzzy propositions	57
	7.3.2. Method of Implication	57
	7.3.3. Method of Defuzzification	58
	7.3.4. Singleton Geometry Output Spaces	60
7.4.	Probability and Possibility	62
8.	A FUZZY INFERENCE SCHEME (PROTOTYPE B)	63
8.1.	The Planning Procedure	63
8.2.	Knowledge Representation	63
	8.2.1. Soft Constraints and Operator Preference	63
	8.2.2. The Fuzzy Planning Propositions	64
	8.2.3. The Output Fuzzy region	64
8.3.	The Decision Model	65
	8.3.1. Inference Mechanism	65
	8.3.2. Responding to the Preferences of the Operator	67

2.6.	Knowledge-based Systems	12
3.	KNOWLEDGE-BASED PLANNING	14
3.1.	Definition of Knowledge-based Systems	14
	3.1.1. Knowledge-Based Systems and Expert Systems	14
	3.1.2. Knowledge-Based Planning Systems	14
3.2.	Knowledge-based Planning Applications	15
	3.2.1. Knowledge-based Planning Applications at Integrated Steel Works	16
3.3.	Current Research Topics in Knowledge-based Planning	17
4.	THE STEELMAKING DOMAIN - A REAL-WORLD CASE STUDY	18
4.1.	The Steelmaking Process	18
4.2.	The Planning Decision Process	18
4.3.	The Set of Planning Constraints	22
	4.3.1. Organisational Goals	22
	4.3.2. Precedence Constraints	25
	4.3.3. Physical Constraints	27
	4.3.4. Resource Availability	28
	4.3.5. Preference Constraints	29
5.	A BINARY HEURISTIC INFERENCE SCHEME (PROTOTYPE A)	31
5.1.	Software Platform	31
5.2.	The Planning Procedure	31
	5.2.1. The Visualisation Interface	31
	5.2.2. The Meta-Plan Generation Interface	32
5.3.	Knowledge Representation	33
	5.3.1. Hard Planning Constraints	34
	5.3.2. Soft Constraints and Operator Preference	34
5.4.	The Decision Model	36
	5.4.1. State-Space Representation	37
	5.4.2. The Search Method	39

TABLE OF CONTENTS

<i>Title</i>	<i>Page</i>
Declaration	i
Abstract	ii
Acknowledgments	iii
List of Figures	ix
List of Tables	x
1. INTRODUCTION	1
1.1. Problem Statement	1
1.2. Proposed Approach	1
1.2.1. A Structured Decision Model	2
1.2.2. Constraint-Directed Search	3
1.2.3. Uncertainty in Soft Constraint Inferencing	4
1.2.4. The Application of Fuzzy Set Theory	4
1.3. Report Structure	5
2. THE PRODUCTION PLANNING PROBLEM	7
2.1. The Planning Decision Process	7
2.2. Production Planning Constraints	7
2.2.1. Constraint Categories	8
2.3. Potential Problems in Formulating a Solution	9
2.3.1. Domain Specificity	9
2.3.2. Conflicting Objectives	10
2.3.3. The Combinatorial Complexity of the Problem Space	11
2.4. The Limitations of Quantitative Planning Techniques	11
2.5. Heuristic Techniques	12

ACKNOWLEDGMENTS

I wish to acknowledge Mr. Keith Sandrock for his supervision and guidance. I am grateful to Prof. I.S. Shaw of the Rand Afrikaans University (RAU) for initiating me on the theory of fuzzy sets and systems.

ABSTRACT

Developers of decision support systems for production planning domains often encounter ill-structured problems that are characterised by uncertainty. This report describes a knowledge-based decision modelling approach which successfully addressed these issues at a steelmaking plant. All planning knowledge was reduced to a set of hard and soft constraints explicitly derived from the relevant domain experts. A prototype system was designed which allowed the operators to vary the relative priorities of the soft constraints according to the often unstable requirements of the various *stakeholders* in the planning decision process.

Soft constraint inferencing was modelled in two ways - a binary heuristic approach and a fuzzy constraint analysis. The two approaches are evaluated and compared. The binary heuristic approach enforced a *rigid* semantic partitioning of the constraint spaces. The consequent *fragmentation* of the problem resulted in a prohibitively complex system, which provided neither consistent nor reliable decision support.

The fuzzy constraint analysis provided significantly improved results. It allowed an *intuitive* representation of soft constraint semantics. Uncertainty was represented as an *intrinsic* part of the decision model. The rule base was significantly reduced and the inference rules were closer to the manner in which the domain experts reason about the problem. The operators were able to *effectively* impose their variable decision criteria on the model.

D.ECLARATION

I declare that this dissertation is my own unaided work. It is being submitted for the degree of Master of Science in Engineering at the University of the Witwatersrand, Johannesburg. It has not been submitted before for any degree or examination at any other University



LML Besteiro

4th day of April 1997

**CUSTOMER ORIENTED PRODUCTION PLANNING
AT AN INTEGRATED STEEL WORKS**

L M L BESTEIRO

**A dissertation submitted to the faculty of Engineering, University of the
Witwatersrand, Johannesburg, in fulfillment of the requirements for the degree of
Master of Science in Engineering**

Johannesburg, 1997

However, there continue to be problems with the quantitative approaches [35]:

- the applied algorithms are generally too complex or impractical for real-world applications
- the models often do not consider all the constraints
- only linear relations among constraint parameters are generally considered
- the models demand exact information about constraints
- it is often not possible to model antagonistic or conflicting information

2.5. Heuristic techniques

Heuristics are criteria, methods or principles for deciding which among several alternative courses of action promises to be the most effective. Unlike the analytical techniques mentioned above (which are theoretically guaranteed to reach a solution) heuristics are not necessarily guaranteed to identify the most effective actions to take, but do so sufficiently often to be useful. Heuristics are most useful for solving ill-structured problems where there are an immense number of possibilities to consider or compare - they act to reduce the number of possibilities considered and thereby reach a satisfactory solution within a reasonable amount of time.

Various heuristic techniques have been devised for the production planning problem (e.g. [36-39]). The basic premise of these techniques is to use priority rules (or a combination of priority rules) to rank the operations and then sequence the ranked operations such that the desired objectives are obtained [40,41]. Most priority rules rely on the critical path of the network; the critical path is generally obtained by the standard PERT/CPM method - studies [36] have shown that such approaches are unreliable for complex, large-sized problems.

2.6. Knowledge-based systems

Recently, some newly developed approaches have been applied to the production planning problem, such as simulated annealing [42], parallel tabu search [43], neural networks [44] and genetic algorithms [45]. In particular, knowledge-based systems [46-49] that incorporate *domain-specific knowledge* have generally proved effective in real-world planning domains.

should be to arrive at a *satisfactory* solution within a reasonable time-frame - note that *satisfactory* is a linguistic expression, generally fuzzy and related to the current preference.

2.3.3. The combinatorial complexity of the problem space

The production planning problem has been classified as NP-complete [12-14]. That is, the set of possible solutions $[S = \{s_1, \dots, s_n\}]$ is combinatorially explosive as the number of operations increases. Even simple planning problems can be shown to be exponential. It is thus important to *prune* ineffectual paths of investigation during the search for a feasible solution.

2.4. The limitations of quantitative planning techniques

Researchers have devised a wide spectrum of quantitative techniques for various simplifications of the production planning problem. Several review papers and books have been written on this subject (e.g. [15-19]). Extensive bibliographies are also available in books by Muth and Thompson [20], Conway et al [21], Eilon and King [22], Elmaghraby [23], Baker [24], Lenstra [25], Kan [26], Coffman [27] and French [28].

Quantitative planning techniques include integer programming [29], branch and bound methods [30-32], dynamic programming [33] and PERT/CPM techniques [34]. These techniques have not been used on a broad basis in real-world domains. Vollmann *et al* [3] indicate that many planning operators and supervisors do not understand the analytical foundations of quantitative models. This viewpoint is supported by the results of various empirical studies (e.g.[1,2]).

The theoretical and mathematical nature of quantitative techniques requires very restrictive assumptions which may not be tenable in practice. In particular, attempts to provide mathematically optimal solutions to real-world planning problems have largely forced implementors to highly simplify the decision model and force fit it to linear continuous functions. Such implementations rely on a considerable amount of human intervention to ensure the generation of feasible solutions. In short, theoretical work on the planning problem has resulted in important improvements.

The explicit nature of real-world planning constraints is often specific to the particular domain. Consequently, domain-specific knowledge is required to completely define the set of constraints and their effect on the planning decision process. The planner essentially acquires all domain-specific planning knowledge from the external environment and *applies* that knowledge to formulate feasible solutions.

2.3.2. Conflicting objectives

The set of objectives [$O = \{o_1, \dots, o_k\}$] are considered to be *conflicting* or at least *competitive* in the sense that each objective is exclusively concerned with obtaining its maximal satisfaction. Moreover, the *relative importance* of planning objectives in real-world domains is often:

- *Variable* - the external environment is essentially a variable and unstable environment. Progress'ive plans are often required to emphasise alternate constraints as real-world scenarios change.
- *Imprecisely defined* - the planning constraints are usually defined and expressed linguistically by the domain experts. Such linguistic expressions are often vague or imprecise. Consequently, the planner generally cannot define a clear *preference policy*.

The *preference policy* of the planner defines the extent to which the various objectives must be satisfied. When all the soft constraints are taken into account, there is generally a fuzzy non-linear relationship between the *state* of a particular constraint - the value of a particular decision variable - and the extent to which the associated objective(s) must be satisfied. Moreover, it is not possible to treat each input to the decision process as an independent decision variable that can be solved in isolation. There are interactions or inter-dependencies such that the combined effects of the decision variable *values* must be considered.

In effect, the production planning problem incorporates multi peaked problem spaces. Any attempt to converge to a uniquely optimal solution is ill-advised. Rather, the fundamental goal

steelmaking process. For example, certain steel grades may only be followed in sequence by certain other steel grades during continuous casting.

Resource availability

This is essentially a temporal constraint. During a particular period, the availability of a resource may be constrained by planned or unplanned maintenance. During the same period the *capacity* of the resource is determined the time available for production runs, as well as set-up times and product mix (different products may imply different processing times).

Preference constraints

A preference constraint may be viewed as an abstraction of other types of constraints. Examples of this constraint include order sequencing preferences and preferential order routing. The reason for the preference may be due to cost or quality factors. Once again, sufficient information generally does not exist to derive actual costs.

2.3. Potential problems in formulating a solution

Advances in planning theory and techniques have generally not produced substantial results in real-world domains due to *domain-specific* features that prevent the application of generic approaches.

2.3.1. Domain specificity

Researchers have traditionally classified the planning problem into two categories:

- *domain-independent* planning, which concentrates on the control mechanisms of planning in general
- *domain-dependent* planning which concentrates on using domain-specific heuristics to encourage efficient search within the particular problem space.

2.2. Production planning constraints

The formulation of a feasible solution (a z^*) to the planning problem must clearly incorporate the set of production planning constraints to be considered and their effect on the decision process.

2.2.1. Constraint categories

Analysis of the planning process at a steelmaking site yielded five broad categories of planning constraints:

Organisational goals

One can view all organisational goal constraints to be approximations of a profit constraint. The goal of an organisation is to maximise profits. Production planning decisions are thus made on the basis of current and future costs incurred. For example, the longer the work-in-process time is, the greater the carrying cost will be for raw materials and value-added operations. In practice, most of these costs cannot be accurately determined, and the planner must therefore make decisions based on expertise or intuitive approximations.

Physical constraints

These constraints specify characteristics which limit functionality. For example, the maximal casting thickness on a continuous casting machine may limit the orders that can be processed through it.

Precedence constraints

This type of constraint defines what conditions must be satisfied before initiating an operation. A precedence rule on an operation states that another operation must take place before it (or after it). Precedence plays a major role in the development of production sequences at the

2. THE PRODUCTION PLANNING PROBLEM

2.1. The planning decision process

The production planning problem is essentially a multi-criteria decision problem comprising a finite set $[S = \{s_1, \dots, s_i\}]$ of possible solutions and a finite set $[C = \{c_1, \dots, c_j\}]$ of planning constraints. A finite set of optimisation goals or objectives $[O = \{o_1, \dots, o_k\}]$ to be satisfied is *implicitly derived* from the set of planning constraints.

For example, temporal delivery-date constraints may imply that orders must be processed through a particular resource before a certain time, in order to be delivered to the customer on-time. Moreover, organisational (policy) constraints may imply that work-in-process (WIP) turnover on the shop floor must remain within certain limits. The semantic intersection of these constraints implicitly defines the following objectives:

- sequence every order before its latest planned starting time (LPST) at the particular resource (maximise the on-time delivery of orders)
- sequence every order as close as possible to its LPST (minimise WIP turnover)

The planning decision process may be viewed as a core I/O process surrounded by an external environment to which the decision maker is accountable. The external environment typically consists of *multiple* domain experts in various sectors. The domain experts impose specific conditions on the decision process, thus, in effect, defining the set of planning constraints.

The *quality* of the decision process is defined as the extent to which the outputs *conform* to the set of planning objectives. Alternatively, the quality of the decision process may be defined as the extent to which the outputs *do not violate* the planning constraints imposed by the external environment.

inferred through a combination of two techniques - an implicit normalisation of the problem space with respect to the semantics of the soft constraint spaces, and (conventional) binary heuristic prioritisation procedures.

- *Chapter 6* provides an evaluation of the results obtained with the binary heuristic inference scheme - the author postulates that most unsuccessful implementations of real-world applications result, at least partially, from the inability of the applied inference scheme to manage the uncertainty inherent in soft constraint inferencing.
- *Chapter 7* provides a theoretical background to the fuzzy set methodology applied in Prototype B.
- *Chapter 8* describes the knowledge representation method and general architecture of the fuzzy inference scheme (Prototype B). Soft constraints are modelled as fuzzy constraints and instantiations are prioritised according to a fuzzy constraint analysis.
- *Chapter 9* provides an evaluation of the results obtained with the fuzzy inference scheme.
- *Chapter 10* provides a summary of the results obtained. Pertinent areas for further research are outlined.

Fuzzy set theory proved to be a powerful decision modelling tool because it provided:

- a method of modelling vague, incomplete and conflicting knowledge from the domain experts
- a method of encoding and using human knowledge in a way that is close to the manner in which the domain experts reasoned about the problem
- a means to model non-linear constraint relations intuitively
- the facilities necessary to break through the computational bottlenecks encountered with the binary heuristic inference scheme

This report includes a comparative analysis between the two inference schemes - emphasis is placed on the problem of uncertainty and the application of a fuzzy set methodology to resolve it.

1.3. Report structure

This remainder of this report is structured as follows:

- *Chapter 2* describes the production planning problem in general terms. The limitations of conventional quantitative and heuristic techniques are outlined. Knowledge-based planning systems are postulated as an effective approach for real-world planning problems .
- *Chapter 3* provides an overview of knowledge-based planning applications, based on a literature survey.
- *Chapter 4* describes the planning problem at the steelmaking domain, which was used as a real-world application domain. The external environment is defined with respect to (functional) domain sectors. The *primary* planning constraints are outlined as per the knowledge elicited from the domain experts in the various sectors.
- *Chapter 5* describes the knowledge representation method and general architecture of the binary heuristic inference scheme (Prototype A). Soft constraints are represented and

Soft constraint inferencing

Soft constraints were used to *prioritise* the consequent set of feasible solutions at each state-space. The motivation behind the representation of soft constraints in the prototype lay in their ability to measure the degree of satisfaction or *non-violation* of constraints.

1.2.3. Uncertainty in planning knowledge

Planning constraints generally do not have equal importance in real-world domains. It is thus reasonable to consider the relative priorities of soft constraints when instantiations are evaluated and compared i.e. it is important to incorporate the preferences of the human planner into the decision model. The preferences of the human planner are assumed to reflect the variable requirements and (constraint) specifications as defined by the external environment.

While formalising the decision process, vague or imprecise information about the preferences of the human planner may be obtained. Soft constraint inferencing is thus characterised by *uncertainty* which prohibits the attainment of a unique, objectively best solution. Consequently, the decision model should incorporate some form of *approximate reasoning* for constraint evaluation and inference. Since real-world planning constraints are often vaguely specified they lend themselves perfectly for being modelled as *fuzzy constraints*.

1.2.4. The application of fuzzy set theory

An evolutionary decision modelling approach was followed in this investigation. A binary (Boolean) heuristic inference scheme was initially designed (Prototype A). This model proved unable to resolve the uncertainty inherent in soft constraint inferencing. It provided inconsistent and unreliable decision support. Consequently, it was *not* accepted by the planning operators and supervisors. Following this a fuzzy inference scheme was designed (Prototype B). This model produced positive results in practice. Prototype B was fully accepted by the planning operators and supervisors.

At its lowest level the decision process itself may be modelled as an *I/O* activity. The inputs to the decision process are various data attributes describing the *state* of the problem space. The decision process is *controlled* by the set of planning constraints. The (inferential) mechanism of the decision process is essentially based upon a repeated comparison between instantiations or partial solutions, which ultimately leads to the generation of a complete production plan.

The author does not *explicitly* differentiate between planning constraints and objectives since:

- an objective cannot exist without there existing one or more constraints to substantiate it
- an objective is implicitly defined or *specified* by the semantic properties of the underlying constraint(s)

Various researchers appear to support this hypothesis. For example, Zimmerman [9] does not distinguish between constraints and objectives, arguing that constraint (representations) empirically model the behaviour of decision makers.

1.2.2. Constraint-directed search

By modelling the decision process as in Fig.1.1 the generation of production plans is cast as a constraint-directed activity that is influenced by *all* relevant constraint knowledge. Given the conflicting and often unstable nature of the domain constraints, the problem differs from typical constraint satisfaction problems and one cannot rely solely on propagation techniques (e.g. [10,11]) to arrive at a feasible solution. Rather, the decision model should capture the *requisite* constraint knowledge and *exploit* that knowledge to control the combinatorics of the underlying problem space.

Hard constraint inferencing

Hard constraints were used to *bound* the solution space. That is, a state-space search method was applied and hard constraints were used to remove all infeasible alternatives as well as their successors from each state-space.

At its lowest level the decision process itself may be modelled as an I/O activity. The inputs to the decision process are various data attributes describing the *state* of the problem space. The decision process is *controlled* by the set of planning constraints. The (inferential) mechanism of the decision process is essentially based upon a repeated comparison between instantiations or partial solutions, which ultimately leads to the generation of a complete production plan.

The author does not *explicitly* differentiate between planning constraints and objectives since:

- an objective cannot exist without there existing one or more constraints to substantiate it
- an objective is implicitly defined or *specified* by the semantic properties of the underlying constraint(s)

Various researchers appear to support this hypothesis. For example, Zimmerman [9] does not distinguish between constraints and objectives, arguing that constraint (representations) empirically model the behaviour of decision makers.

1.2.2. Constraint-directed search

By modelling the decision process as in Fig.1.1 the generation of production plans is cast as a constraint-directed activity that is influenced by *all* relevant constraint knowledge. Given the conflicting and often unstable nature of the domain constraints, the problem differs from typical constraint satisfaction problems and one cannot rely solely on propagation techniques ([10,11]) to arrive at a feasible solution. Rather, the decision model should capture the *requisite* constraint knowledge and *exploit* that knowledge to control the combinatorics of the underlying problem space.

Hard constraint inferencing

Hard constraints were used to *bound* the solution space. That is, a state-space search method was applied and hard constraints were used to remove all infeasible alternatives as well as their successors from each state-space.

It sacrifices the goal of rigorous optimality, concentrating instead on a credible and effective approach to production planning.

1.2.1. A structured decision model

Production planning was modelled as a *core decision process* surrounded by a multi-sectored *external environment* [Fig.1.1]. The external environment consists of various *domain experts* who collectively define the set of planning constraints and an *implicit* or *implied* set of planning objectives. All domain-specific planning knowledge is thus derived from the external environment. The planner *acquires* and *applies* this knowledge to generate production plans or sequences over particular planning horizons.

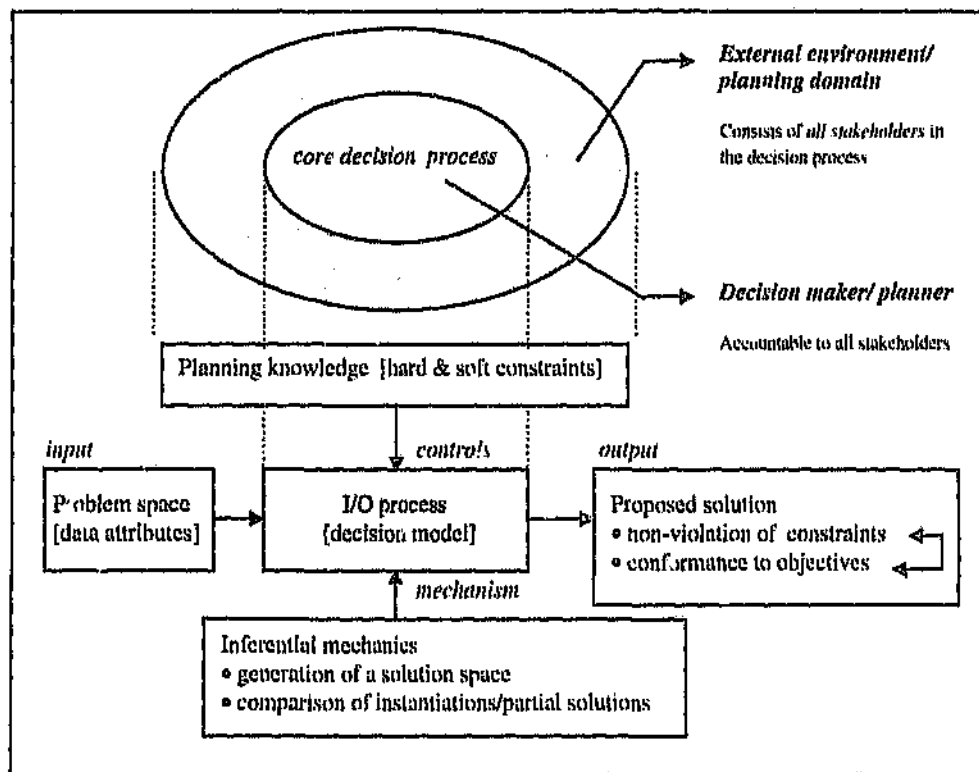


Fig.1.1. Customer oriented production planning

Table 4.3. Example of hot-rolling standards

Coffin name	Min length [m]	Max length [m]	
Simple coffin type	40 000	55 000	
Zone name	Min zone length	Max zone length	Allowable coils [mm]
Start zone	2000	4000	900 ≤ coil width ≤ 1100 3.0 ≤ coil thickness ≤ 4.0 Use soft coils Use high tolerance coils
Ascent zone	10 000	15 000	900 ≤ coil width ≤ 1500 2.5 ≤ coil thickness ≤ 4.5
Difficult dimensions zone	5000	15 000	1300 ≤ coil width ≤ 1300 1.5 ≤ coil thickness ≤ 2.5 OR 1400 ≤ coil width ≤ 1600 2.0 ≤ coil thickness ≤ 3.0
Descent zone	25 000	40 000	700 ≤ coil width ≤ 1400 2.5 ≤ coil thickness ≤ 10.0 Use no hard coils

During the generation of production sequences the urgency of any steelmaking order o_i is related to the current actual slab stock level S_i of the ordered product type p_i :

$$\Delta s_i = (S_i - \beta_i) / \beta_i \quad (\beta_i \leq S_i \leq \theta_i)$$

where θ_i = the maximum slab stock level for p_i (determined by cost and storage constraints)

β_i = the minimum slab stock level for p_i (slab buffer)

The structural integrity of the slab stock requires that $S_i \geq \beta_i$ and $S_i \leq \theta_i$ for every product type in the slab yard. As $S_i \rightarrow \theta_i$ the urgency of o_i decreases. Conversely, as $S_i \rightarrow \beta_i$ the urgency of o_i increases. The organisational goal of minimised stock levels implies that $S_i \rightarrow \beta_i$. The steelmaking planners were observed to treat θ_i and β_i as *fuzzy boundaries* - at times they violate the hard constraint barriers θ_i and β_i which results in:

- *gaps* in the slab population which makes it difficult to construct future rolling programs because of the lack of steel of the required product type in the slab yard
- long turnover times at the slab yard for stock of particular product types

The EPST is constrained by material availability at the steelmaking plant. The LPST is constrained by the required delivery date of slabs to the hot-strip mills (HISM1 and HISM2).

It is unwise to process a task too long before the LPST since this will increase the carrying costs of slab stock *and* reduce the availability of resources for other more imminent tasks. In the prototype, the delivery-date constraint is represented in terms of the deviation (in days) between the current real time t_c and the LPST⁺ of the related steelmaking order:

$$\Delta t = \text{LPST} - t_c$$

This constraint implies that

$$t_s \geq \text{EPST} \cap t_c \leq (\text{LPST} + t_p) \text{ with } \Delta t \rightarrow 0$$

where t_p = the processing time for the particular steelmaking order

t_s = the planned start time

t_e = the planned end time

The slab stock structural integrity constraint (Δs)

Most steel companies have a variety of hot rolling programmes - also called coffin plans or coffin shapes. A coffin plan defines how a rolling program must be constructed. Typically a coffin plan is characterised by starting dimensions, problematic dimensions and total length. The coffin shape is generally divided into zones. Each coffin zone is defined by slab selection criteria and a minimum-maximum range of the total zone length. An example is shown in Table 4.3. The available slab stock must correspond to the slab selection criteria of future rolling programs.

At the steelmaking domain the product mix is categorised according to slab width and steel grade. Slabs whose properties differ with respect to the above criteria are classified as different product types. The maximum and minimum stock levels for a particular product type p_i are determined by the line supervisors at the hot-strip mills.

Table 4.2. Top level breakdown of the external environment for the steelmaking domain

FUNCTIONAL DOMAIN SECTION	RELEVANT (CONSTRAINT) ISSUES
<ul style="list-style-type: none"> • sales employees • clients 	<ul style="list-style-type: none"> • delivery-date problems \leftrightarrow order prioritisation
<ul style="list-style-type: none"> • line supervisors (steelmaking plant) • production engineers (steelmaking plant) • line operators (steelmaking plant) 	<ul style="list-style-type: none"> • sequence precedence and preference issues • resource availability restrictions
<ul style="list-style-type: none"> • engineers/specialists in quality control and manufacturing/technology operations 	<ul style="list-style-type: none"> • product/production quality
<ul style="list-style-type: none"> • line supervisors (hot-strip mills) • production engineers (hot-strip mills) • line operators (hot-strip mills) 	<ul style="list-style-type: none"> • delivery-date problems • structural integrity of work-in-process (slab stock)
<ul style="list-style-type: none"> • senior management 	<ul style="list-style-type: none"> • organisational goals (maximisation of overall throughput/profit, reliability of the delivery service, etc.)

4.3. The set of planning constraints

The *primary* planning constraints for the steelmaking domain have been categorised as organisational goals, precedence rules, physical constraints, resource availability and preference constraints. Although the constraints have been explicitly categorised, there is an implicit overlap or *inter-dependence* between constraints in different categories.

4.3.1. Organisational goals

This type of constraint is generally related to production costs and to the total quality of the *delivery service* to the *clients* (steel consumers). Consequently, these constraints are defined by line supervisors, senior management and the clients.

The delivery-date constraint (Δt)

The delivery-date constraint is a *soft* temporal constraint. A task may be processed at any time between its earliest planned starting time (EPST) and its latest planned starting time (LPST) at the steelmaking plant. The EPST and the LPST for a particular task are determined by a higher-order, works-wide capacity planning system, generally referred to as *the Works Loading System*.

Table 4.1. Data attributes for each steelmaking order in the task space

ATTRIBUTE	EXAMPLE VALUE
Heat (task) number	C0784U
Customer order number	ITC 7030188
Delivery date to the hot-strip mills	1997/1/23
Grade	MC1 ¹
Slab width	950 [mm]
Slab length	9450 [mm]
Slab thickness	240 [mm]
Ordered mass [heat]	159 [tons]
Rolling process	F1 [North Mills]
Steelmaking process route	54AM

Planning generally does not occur in a uniform or structured manner. Each planning decision to be made entails side effects whose importance varies by order. A proposed sequence must be approved by persons in various departments. Each person can provide information about specific constraints which may result in sequencing alterations. The author found that the human planners spend ~20% - 30% of their time actually planning and ~70% - 80% of their time communicating with other employees - and with clients - to determine what overall constraints should influence a particular plan.

The planning decision process is thus an *iterative* and *interactive* process and it is important to structure it such that the (often unstable) requirements of the external environment - reflected by the preferences of the human planner - are effectively captured in the decision model. This was achieved by following the approach illustrated in Fig.1.1.

The planning rules are explicitly derived from the various domain experts in the external environment [Table 4.2]. The human planners were observed to violate or relax these *a priori* and to implicitly make trade-offs in ~25% of their decisions. Their reasoning was adequately modelled by using a combination of *hard constraint inferences* (expressed as statements that are either entirely TRUE or entirely FALSE) and *soft constraints inferences* (expressed as statements that have a relative degree of truth within a finite *continuum* of allowable values).

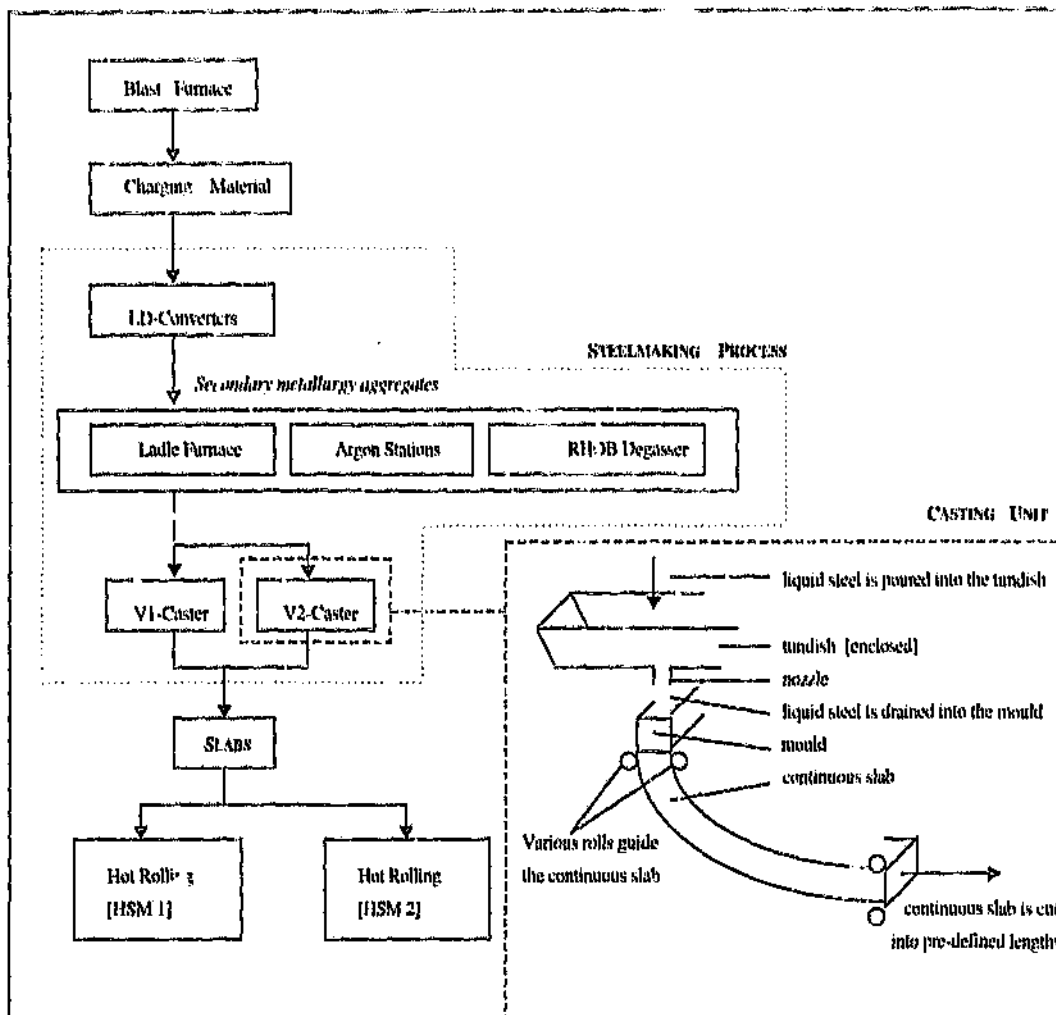


Fig.4.2. The steelmaking process

The acceptability of a particular production sequence depends on such diverse and conflicting factors as delivery-date requirements, production cost restrictions, product quality restrictions and resource availability. The generation of plans at this level is man-hour intensive, and not amenable to rapid replanning in the case of unforeseen events. The human planners generally work under permanent crisis conditions - they act as a buffer between the generally unrealistic requests of the sales staff and the often exacting demands of the production staff, who desire long production runs with minimal disturbance.

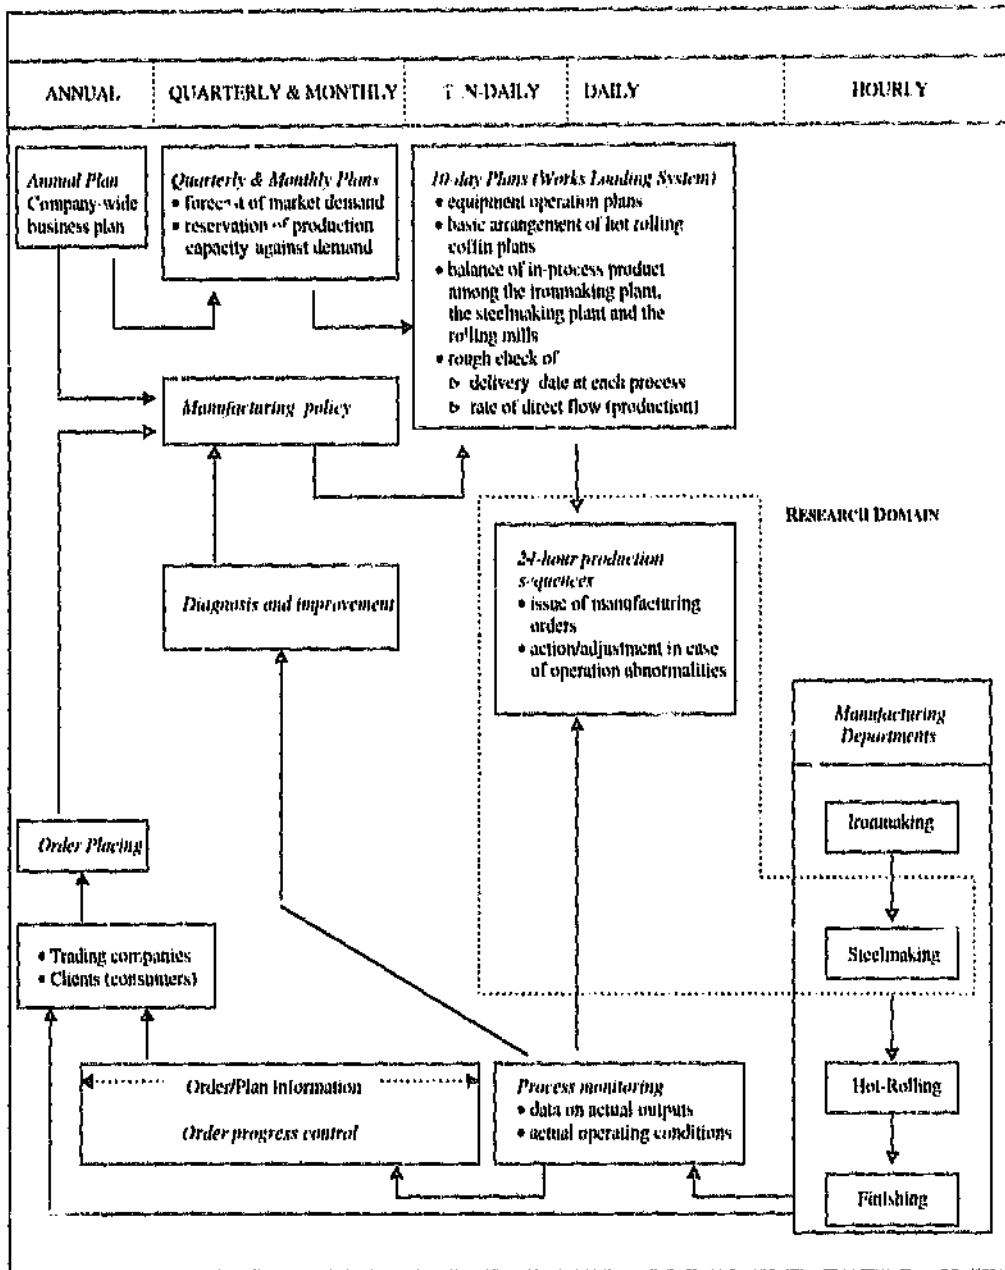


Fig.4.1. Production planning activities at an integrated steel works

4. THE STEELMAKING DOMAIN - A REAL-WORLD CASE STUDY

The steelmaking domain was used as a real-world application domain for the proposed decision modelling approach. Production planning activities at an integrated steel works occur in stages [Fig.4.1]. The different stages can generally be identified in terms of time scale and the degree of detail considered. As one moves down the planning hierarchy the degree of detail increases and the human planners are forced, due to time and complexity constraints, to make decisions based on experience and intuition rather than true cost or logistical considerations.

4.1. The steelmaking process

Steelmaking is essentially a batch type process [Fig.4.2]. Hot pig iron is delivered from the blast furnace to the LD-converters. The steel is then poured into ladles, processed further in secondary metallurgy aggregates, and finally delivered to the casters. The casters produce continuously cast strands of steel that are cut into slabs of specific length. Although the steel is produced in different aggregates in a fixed flow, the casters are the bottleneck resource of the steelmaking plant. Consequently, production planning is centred around the sequencing of tasks at the casters.

4.2. The planning decision process

A 24-hour production sequence for each caster must be established on a daily basis. Sequencing starts with a pool of several hundred steelmaking orders or *tasks* [$TS = \{\tau_1, \tau_2, \dots, \tau_n\}$] characterised by the attributes listed in Table 4.1. The steelmaking orders are prepared by the hot-strip mills. From this data the steelmaking planners have to construct casting sequences, taking into account all planning constraints.

The *mental schemata* or reasoning process followed by the planners is based upon traditional or intuitive planning rules and heuristics which they use to select feasible combinations of tasks for each planning horizon. Although experienced planners are fully aware of the planning constraints, they have neither the time nor the capability to reason coherently about thousands of instantiations in order to optimise the sequencing of operations.

the authors concluded that the solutions could be improved if constraint relaxation techniques were employed.

A production planning system (*α -vision*) was developed at the NKK Koihin works to plan the steel making process [75]. The system is based on *SCHPLAN*. It was installed to meet stringent customer requirements and to achieve a highly efficient operation using a single blast furnace. The system enhanced the competitiveness of the works by increasing throughput, reducing stock levels and expanding information on the progress of orders.

Nakayama *et al* [76] developed a KBS for planning decision support at a steelmaking plant. The KBS is based on a heuristic algorithm and discrete event simulation. W. Slany *et al* [35,77] employed fuzzy logic to resolve various planning problems at the steelmaking process, including the real-time scheduling of resources. Bernatzki [78] describes a KBS used to optimise production operations at the control stands of the Krup Stahl Huckingen steel works. The optimisation procedure is based on a fuzzy Petri net.

3.3. Current research topics in knowledge-based planning

The current research trends in knowledge based planning are exemplified by the Advanced Research Projects Agency (ARPA) research and development effort aimed at developing the next generation of knowledge-based planning tools [79]. Research topics include primarily: *reasoning under uncertainty*, constraint-based planning, decision theory, plan justification, execution monitoring & replanning support, and temporal reasoning. The ARPA research has yielded both theoretical and practical advances in knowledge-based planning technology. Stillman and Bonissone [80] provide a review of knowledge-based planning applications derived from the research.

replanning system interfaced with an MRP system to regenerate production sequences. It stabilises any impact of abnormal events by constraint relaxation.

Ben-Arieh [64] developed a Prolog based expert planning system using *SLAM II* to simulate a manufacturing environment, with Pascal sub-routines to calculate the cost of competing solutions. In one of his papers [65] he presents two methods for knowledge-based control, one at cell level and the other at workstation level. Emphasis is given to the ability of the control techniques to deal with unpredicted conditions in a dynamic environment. Bruno *et al* [66] developed an OPS5 based KBS, partially controlled by a simulation sub-system for sequencing parts in an FMS. The KBS generates a sequence and the simulation system evaluates it. Newman [67] summarises applications of KBS for production sequencing in CIM environments.

OPIS, developed by Smith [68] uses a constraint based framework to deal with reactive planning. The framework incrementally reconciles inevitable discrepancies that arise between the predictive plan and actual behaviour of the production environment. *ESS*, developed by Jain *et al* [69] uses real-time plant information to yield production sequences that do not violate resource availability constraints. The general plant knowledge is organised into several specific knowledge bases. Kerr and Ebsary [70] describe a KBS for production sequencing. The goal of the KBS is to provide *consistent* sequences. Alexander [71] developed a conceptual framework for a KBS that allows the operator to select between a number of priority rules, depending on the planning objective. Biegel and Wink [72] describe a KBS for production sequencing, based on priority rules and heuristics. A LISP based system, *OPAL* [73] used a control strategy based on a fuzzy set methodology to generate production sequences in a job shop environment.

3.2.1. Knowledge-based planning applications at integrated steel works

VAI-SchedEx [74] was installed at the LD3 steelmaking plant of VOEST-ALPINE Stahl. It used priority rules and heuristics to generate production sequences for the steelmaking plant. The KBS initially produced sub-optimal solutions; moreover, the processing time was too long (~50 minutes). After successive development iterations the solutions were improved and the processing time was reduced to ~5 minutes. Although the system produced feasible solutions,

in common with ES applications. However, the typical ES applications are generally used for classification and interpretation. Moreover, they work in a convergent fashion. On the other hand, knowledge-based planning is a synthetic task where inherently more search is involved since there are alternatives (feasible instantiations or partial solutions) to consider and compare.

Knowledge-based planning techniques generally seek to control the extent of the search. The knowledge-based planning system is charged with generating a definitive sequence of operations, which is one possible solution to the specified problem.

3.2. Knowledge-based planning applications

Several KBS applications have been developed for planning operations. The review papers by Kusiak & Chen [53] and Gupta & Chin [54] provide excellent summaries of knowledge-based planning applications. The review by Kusiak and Chen is particularly relevant because it provides specific information regarding the type of application, programming languages used, type of knowledge representation, control strategy and structure.

One of the early KBS planning applications was *ISIS* by Fox and Smith [55]. They view production sequencing as a constraint directed search and develop a search architecture capable of exploiting the constraints to improve the efficiency of the search. *ISA*, a KBS developed at DEC [56] generated a loading strategy to deal with difficult orders. Robbins [57] developed a LISP based prototype expert priority planner [*PEPS*] for solving problems at shop floor control level. Tou [58] presented a KBS structure for integrated production automation. Shaw and Winston [59] used constraint directed search to plan in a three-machine cell, with Petri-nets to model the interaction between events.

Mascot, a Prolog based production sequencing system [60] used a constraint based analysis to generate precedence relationships for conflicting resources. Steffen and Greene [61] used hierarchical planning and constraint directed heuristic search to develop a prototype planning system for a set of parallel processors. Subramanayam and Askin [62] developed a Prolog based expert planning system for part selection in an FMS. The expert rules were acquired from job shop simulation and domain experts. *SCORE*, developed by Chiodini [63] is a real-time

3. KNOWLEDGE-BASED PLANNING

3.1. Definition of knowledge-based systems

For the purpose of this investigation a *knowledge-based system* (KBS) denotes a software application that can store knowledge about a particular domain and use that knowledge to solve domain-specific problems in an intelligent manner. The typical structure of a KBS is shown in Fig. 3.1.

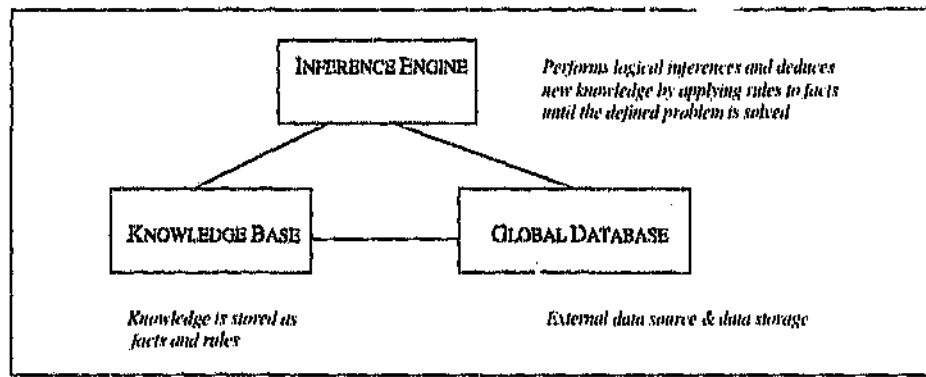


Fig.3.1. The basic structure of a knowledge-based system

3.1.1. Knowledge-based systems and expert systems

According to Duda and Shortliffe [51] KBS are also called *expert systems* (ES) if they refer to problem solving in those areas and at that level of performance that is usually achieved by human experts. A generic KBS with the above structure, but with an empty knowledge base, is called an *expert system shell* [52]. Expert system shells are suitable for building domain-specific KBS by supplementing the particular knowledge base with knowledge and inference heuristics that are specific to the domain.

3.1.2. Knowledge-based planning systems

Knowledge-based planning systems are simply KBS applications that have been developed for planning operations. Knowledge-based planning systems share many techniques and problems

According to Fogarty and Hofmann [4] almost all human planners apply a set of planning rules based upon their knowledge and expertise. The rule-driven nature of the decision process suggests that knowledge based systems are appropriate tools for production planning decision support. Since knowledge based systems emulate - or attempt to emulate - the reasoning process of the human planners, they represent less enigmatic planning tools when compared to quantitative techniques. Consequently, the application of knowledge-based systems potentially alleviates the acceptance problems associated with such techniques.

The explicit representation of the planning problem through domain-specific knowledge improves the efficiency of the decision process by pruning useless paths of investigation, ordering the search, eliminating redundancy, reducing ambiguities and exploiting knowledge from complementary and contradictory sources [8]. According to Price *et al* [50] knowledge-based systems provide a methodological approach to solve complex planning tasks that inevitably require expertise.

To date, several commercial knowledge-based tools have been developed that provide facilities to represent progressively more realistic models of real-world domains and their specific constraints. Examples of such facilities include object-oriented representations, structured natural language and interactive graphics. Besides improving interaction with the user and with existing information systems, these facilities have improved the range of production planning problems that can realistically be tackled.

The implications of the above statement are described in Table 5.1. The total range of decision variable values encompassed in the task space regions restricts the solution space as follows:

- The hard barrier $\Delta t = 0$ (TS(1•x•y•z) \rightarrow [0 \leq Δt \leq 2]) implies that the prototype will not consider a task whose LPST has expired.
- The hard barrier $\Delta t = 20$ (TS(4•x•y•z) \rightarrow [10 $<$ Δt \leq 20]) implies that the prototype will not consider a task more than 20 days prior to the LPST; if LPST - EPST $>$ 20 days and no feasible tasks are found, the prototype will selectively shift the EPST of tasks forward to equal LPST - 20 days.
- The hard barriers $\Delta w = -50$ (TS(w•14•y•z) \rightarrow [-50 $<$ Δw \leq -35]) and $\Delta w = 100$ (TS(w•15•y•z) \rightarrow [85 $<$ Δw \leq 100]) imply that the prototype will not consider a task whose slab width is more than 50mm wider, or more than 100mm narrower, than the previous task in sequence.
- The hard barrier $\Delta g = 0.40$ (TS(w•x•40•z) \rightarrow [0.39 $<$ Δg \leq 0.40]) implies that the prototype will not consider a task whose chemical analysis deviates by more than 0.4% from the chemical analysis of the previous task in sequence. Note that the maximum allowable deviation for each individual chemical element is respected. The maximum allowable deviation for certain critical elements are multiplied by a factor f ($1 < f \leq 1.5$). This factor is selectively minimised if no feasible tasks are found.
- the hard barrier $\Delta S = 1.0$ (TS(w•x•y•10) \rightarrow [0.9 $<$ ΔS \leq 1.0]) implies that the prototype will not select a task τ_i if the actual slab stock level S_i exceeds the maximum allowable stock level θ_i .

5.4. The decision model

During the first phase of development the author found that the human planners intuitively categorised the overall planning problem into distinct sub-problems, each of which corresponded roughly to a heuristic framework that could be applied repetitively to arrive at a solution. The sub-problems relate to the establishment of various *types* of partial sequences which are progressively concatenated into a complete production sequence.

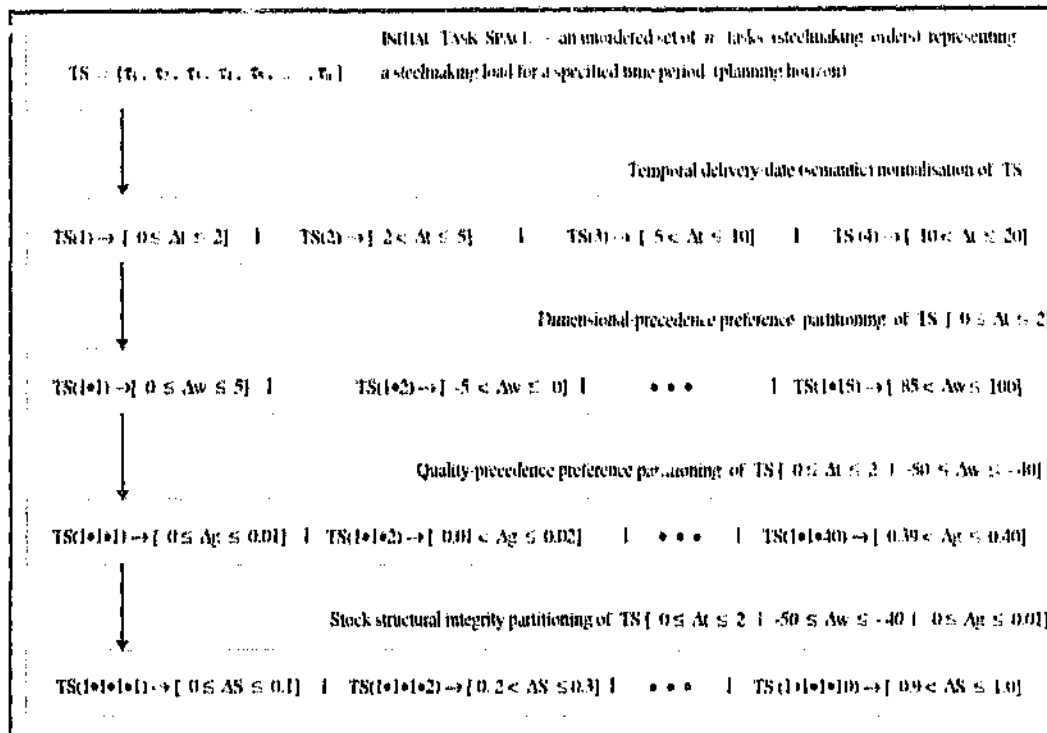


Fig.5.1. Task space normalisation schematic

the relevant constraint spaces. The tasks are implicitly *prioritised* according to their location in the normalised problem space. The normalisation of the problem space reflects the preference policy of the human planners. The total task space TS was sub-divided into 24 000 regions i.e.

$$TS = TS(1\bullet1\bullet1\bullet1) \cup TS(1\bullet1\bullet1\bullet2) \cup \dots \cup TS(4\bullet15\bullet40\bullet9) \cup TS(4\bullet15\bullet40\bullet10)$$

TS(1•1•1•1) represents the highest priority region while TS(4•15•40•10) represents the lowest priority region. Tasks were ranked according to their *membership* in a particular task space region. Thus, a task τ_i would be assigned the highest priority if the following condition were TRUE:

$$\Delta t_{\tau(i)} \in [0, \dots, 2] \cap \Delta w_{\tau(i)} \in [0, \dots, 5] \cap \Delta g_{\tau(i)} \in [0, \dots, 0.01] \cap \Delta S_{\tau(i)} \in [0, \dots, 0.1]$$

5.3.1. Hard planning constraints

These constraints are primarily related to the *special steel grades* that require extraordinary measures to avoid contamination. These measures refer to the minimum and maximum allowable number of tasks between tundish changes and casting interruptions, as well as the relative and absolute sequencing of orders.

Hard quality precedence constraints (see 4.3.2) are represented by *sequence-constraint specification (SCS)* objects [Appendix 3]. The operator may easily create, modify or delete SCS objects as dictated by the external environment. The SCS objects are named according to the special grades to which they refer. Each SCS object includes the following data:

- The start grade i.e. the steel grade that must be cast immediately after a tundish change.
- The minimum and maximum number of tasks - with the required start grade - that may be cast on a single tundish
- The follow-on grades i.e. the steel grades that may be cast after the start grade on the same tundish
- The minimum and maximum number of tasks - including all grades - that may be cast before a tundish change is required

5.3.2. Soft constraints and operator preference

Conceptually, the preferences of the human planner - or any domain expert from the external environment - relates to a bias for decision variable *values* over specific regions (semantic partitions) of the underlying soft constraint spaces. The soft constraints considered for the steelmaking domain include the temporal delivery-date constraint (Δt), the dimensional-precedence preference constraint (Δw), the quality-precedence preference constraint (Δg) and the slab stock structural integrity constraint (Δs).

In order to model the semantics of soft constraints the problem space was effectively normalised [Fig.5.1] by sub-dividing it into task space regions that correspond to the semantic partitions of

Each task τ_i must be processed within the temporal delivery-date constraint space i.e.

$$t_{start} \geq EPST \cap (t_{end} - t_{proc}) \leq LPST \quad [t_{proc}] = t_{end} - t_{start}$$

where EPST = the earliest planned start time for τ_i

LPST = the latest planned start time for τ_i

t_{proc} = the processing time for τ_i

t_{start} = the planned start time for τ_i

t_{end} = the planned end time for τ_i

In general, resource availability constraints may induce new precedence constraints at the casters due to:

- the presence of common (pre-casting) resources for particular operations
- sequencing (precedence) constraints on particular pre-casting operations

A casting sequence is feasible if no pre-casting resource availability constraints are violated and the total processing time for all tasks does not exceed the available processing time at the casters.

5.3. Knowledge representation

The decision model may be seen to consist of a core heuristic inference scheme surrounded by a *knowledge shell*. The planning knowledge is stored in files and represented graphically on various screens in a methodical and transparent manner. Any graphical element (object) on the knowledge shell may be modified or updated by the operator at any time. The knowledge shell represents a structured way of encapsulating the planning knowledge acquired from the external environment.

As the generation of sequences progresses the matrix is dynamically updated, allowing the operator to visualise the resulting changes to the state of the problem space.

5.2.2. The meta-plan generation interface

After visualising the task space the planner establishes a meta-plan for each caster. The meta-plan is a *basic indication* of what the complete sequence should look like - it represents what *types* of partial sequences must be planned and how many such partial sequences must exist between casting interruptions. A partial sequence represents a sequence of tasks between tundish changes.

Each partial sequence is represented by a graphical *partial-sequence-type (PST)* object [Appendix 3]. Casting interruptions are represented by separate graphical objects. The operator establishes a meta-plan by selecting and dragging *PST* objects with a mouse and then connecting them to other *PST* objects - or to casting interruptions - in a fixed casting direction.

Once the meta-plans have been established the prototype generates detailed production sequences that conform to the meta-plans. The prototype provides the planner with a preliminary *resource availability status* as the meta-plans are established. If resource availability constraints are violated the operator is informed via a screen. Detailed resource availability checks are carried out during the actual generation of sequences.

The prototype reasons about resource availability as follows. A definitive sequence of l tasks must be generated for *each* continuous series of *PST* objects. Each task τ_i must be processed through a set Ω_i of operations at the steelmaking plant (Ω_i consists of both casting and pre-casting operations). The set Ω_i is defined by the *steelmaking process route* attribute of the task τ_i [Table 4.1]. Every operation ω_i ($\omega_i \in \Omega_i$) requires a processing time t_{pi} . Moreover, Ω_i is a partially ordered set in the sense that $\omega_i < \omega_j$ ($\omega_i \in \Omega_i$) since ω_i must be processed before ω_j due to general precedence constraints at the steelmaking plant.

5. A BINARY HEURISTIC INFERENCE SCHEME (PROTOTYPE A)

5.1. Software platform

The prototype was coded on G2 - a knowledge-based expert system shell developed by GENSYM [82] and distributed in South Africa by KNOWLEDGE BASED ENGINEERING [83].

5.2. The planning procedure

During execution of the prototype a *problem space* is initially generated. The problem space includes *all* steelmaking orders currently available for planning. Typically the problem space consists of ~2000 - 2500 tasks.

Once the problem space has been generated, the operator interacts with the prototype via a *graphical user interface* to establish complete 24-hour production sequences for each caster. The *operator* may be any one of three steelmaking planners. The *supervising planner* assists in the planning process where severe problems are encountered.

5.2.1. The visualisation interface

This interface consists of the *task matrix screen*. This screen provides the operator with an effective visualisation of the task space [Appendix 2]. The matrix provides the following information:

- the slab width distribution of tasks - only slab widths for which there are tasks in the task space are represented
- the steel grade distribution of tasks - only steel grades for which there are tasks in the task space are represented
- the exact number of tasks per slab width and steel grade
- the relative importance of tasks - per width and grade - with respect to the temporal delivery date constraint

Quality-precedence preference (Δg)

Although tasks may legally be cast in sequence as long as the chemical analysis deviations between subsequent heats are within the limits shown in Fig.4.3, the technology engineers and line supervisors prefer that the analysis deviations be minimised. There are costs associated with the required sequencing operations between subsequent tasks with different steel grades - the greater the deviation between chemical analysis, the longer the length of slab that must be cut off and scrapped. The deviation in chemical analysis is defined as:

$$\Delta g_{ij} = \sum_{E} \varepsilon_j - \varepsilon_i$$

where ε_i and ε_j = the percentage composition of a particular chemical element in tasks τ_i and τ_j respectively

E = the set of relevant chemical elements defined by the product technology engineers [Fig.4.3]

This constraint implies that $\Delta g_{ij} \rightarrow 0$ and $\varepsilon_j - \varepsilon_i \leq e_{max}$ where e_{max} is the maximum allowable deviation for each chemical element e ($e \in E$).

Dimensional-precedence preference (Δw)

Tasks may legally be cast in sequence as long as the deviation in slab width between subsequent heats does not exceed the allowable range i.e. 100mm narrower and 50 mm wider. However, the line supervisors prefer that the slab width deviations between subsequent heats be minimised in order to minimise the tapering effect on the slab during continuous casting, particularly with the higher casting speeds [Table 4.4]. The deviation in slab width between subsequent tasks τ_i and τ_j is defined as:

$$\Delta w_{ij} = w_j - w_i$$

where w_j = the slab widths for τ_j

w_i = the slab widths for τ_i

resource is prepared for the next sequence of ~7 tasks. This constraint implies that the sequencing of operations at the casters must be compatible with the degassing procedures at the RHOB unit, thus, in effect, imposing new precedence constraints at the casters.

Tundish utilisation

The tundish, a part of the casting unit, has to be maintained after approximately 240 minutes. The tundish maintenance takes roughly 100 minutes. Therefore, a second tundish is used on each caster while the first is maintained. In order to ensure the availability of tundishes at all times, the following constraints have been imposed by the production engineers and line supervisors:

- an average of at least 6 tasks between tundish changes is required over any 24-hour period
- a *minimum* of 3 tasks must be processed before a tundish change may occur, with the exception of certain steel grades where less than 3 tasks per tundish are permitted due to quality precedence rules
- no more than 3 consecutive tundishes with ≥ 3 tasks each is permitted

4.3.5. Preference constraints

These constraints may be viewed as abstractions or extensions of other constraints. The reason for a particular preference may be due to cost or quality factors. Sufficient information generally does not exist to derive actual costs. By their very nature, these constraints are vague and imprecise. They may be modelled *implicitly* within the semantic representations of other constraints i.e. they relate to a preference for decision variable values within particular semantic partitions of the associated (soft) constraint spaces. Two important preference constraints exist at the steelmaking domain: quality-precedence preference and dimensional-precedence preference.

4.3.4. Resource availability

These are essentially temporal constraints and refer primarily to the availability of the casters (casting units) and the pre-casting resources i.e. the LD-Converters and the secondary metallurgy aggregates.

The capacity of the casters

The casters are the bottleneck resource of the steelmaking plant. The maximum casting speed (measured in m/min) depends on the steel grade and slab dimensions [Table 4.4]. Clearly, the product mix affects the capacity of the casters since different products imply different processing times. Moreover, the average sequence length has a direct impact on the capacity of the casters, since each casting interruption requires a complete set-up of the particular caster. The set-up times vary between ~80 minutes to ~240 minutes depending on the required set-up procedures.

Table 4.4. Casting speed vs slab width at the casters

Slab width [mm]	Casting speed [m/min]
950 - 1000	1.1
1025 - 1200	1.0
1225 - 1400	0.9
1425 - 1600	0.8
1625 - 1750	0.7
1775 - 2000	0.6

Degassing procedures

Degassing operations restrict the availability of the RHOB Degasser [Fig.4.2]. All heats with ultra-low carbon (ULC) and low-carbon (LC) steel grades must be treated at this resource. At the present time a maximum of ~7 tasks may be treated consecutively at the resource. Following this, a mandatory set-up time of 4 - 5 hours is required, during which time the

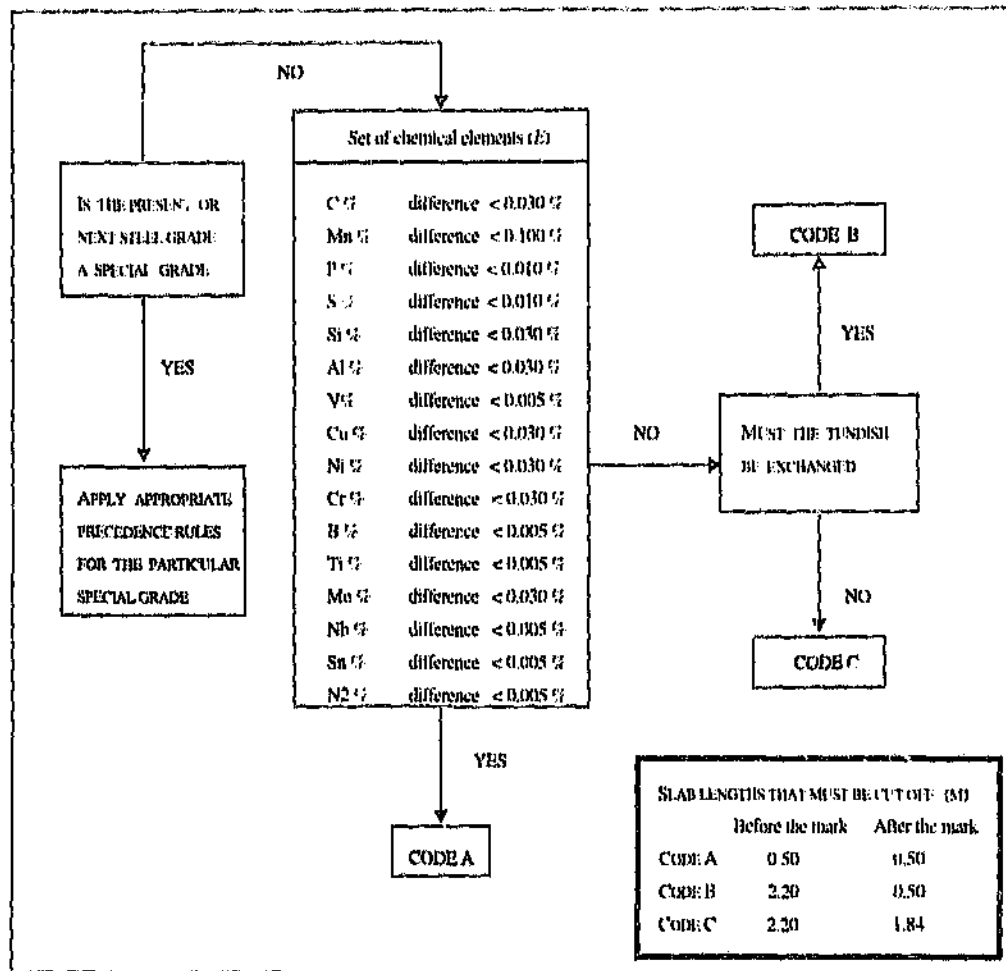


Fig.4.3. Quality precedence constraint specific to steel grades

casting units and are imposed by the production engineers. At the present time, the maximum width adjustments permitted on the casters are 100 mm narrower or 50 mm wider between consecutive heats. Width deviations outside these limits require a complete caster set-up.

4.3.3. Physical constraints

These constraints result from the limitations of the casting units and are imposed by the production engineers. For example, the V1-caster may only produce slabs of thickness ~210mm and the V2-caster may only produce slabs of thickness ~240mm.

Tasks may be consecutively cast in sequence as long as they do not violate the product quality constraints imposed by the technology engineers and line supervisors. The general precedence rules in this category are shown schematically in Fig.4.3. The *special steel grades* defined in Fig.4.3 require extraordinary measures to avoid contamination:

- certain steel grades may only be cast immediately after a complete caster set-up i.e. after a casting interruption
- certain steel grades may only be cast immediately before a complete caster set-up i.e. before a casting interruption
- certain steel grades must be cast singularly between casting interruptions i.e. casting must be interrupted before and after these grades are cast and no other grades may be cast in the same sequence
- certain steel grades may only be cast immediately after a tundish change
- certain steel grades may only be cast immediately after a tundish change *and* only a limited number of tasks may be cast using the same tundish; however, certain other steel grades may be subsequently cast using the same tundish
- certain steel grades must be cast singularly in a tundish i.e. the tundish must be changed before and after these grades are cast and no other grades may be cast using the same tundish

For example, consider the MC1 steel grades [MC1⁰¹, MC1⁰², ... MC1³¹]. These grades must be cast immediately after a tundish change. Moreover, they must be cast immediately after a casting interruption i.e. a complete caster set-up. A minimum of 3 and a maximum of 5 MC1 tasks may be processed with a single tundish. Any of the following - but *only* the following - steel grades *must follow* the MC1 steel grade on the same tundish - AC2, AC3, MC2, MC3. These grades may follow the MC1 grades in any order. A minimum of 6 and a maximum of 8 tasks - including all the above grades - may be processed before a tundish change is required.

Dimensional precedence rules

Tasks may be consecutively cast in sequence as long as they do not violate the allowable subsequent deviation in slab dimensions. These constraints result from the limitations of the

The sequence length (l)

The sequence length is defined as the number of tasks planned between casting interruptions. To optimise production the casters should ideally process steel without interruption. This implies that the average sequence length should be maximised. However, sequencing operations have associated costs and precedence rules. The most expensive operation is to stop the caster and then set it up from scratch - this is necessary in order to:

- produce certain problematic steel grades i.e. adhere to certain precedence rules that impose casting interruptions
- perform either planned or unplanned maintenance on the caster

A survey undertaken at various steel companies [81] concluded that the factors restricting the on-time delivery of orders included primarily the average sequence length at the casters.

4.3.2. Precedence constraints

This type of constraint relates to compatibility criteria between subsequent tasks in a sequence. Two categories of precedence rules exist for the steelmaking process

- precedence rules that relate to product quality constraints
- precedence rules that relate to dimensional constraints

Quality precedence rules

Steel manufacturing is a make-to-order process. Within the steelmaking milieu, every task is of a specific steel grade. The steel grade refers to the chemical composition of the steel, which in turn is related to the specific application for which the steel was ordered by the client. Client specifications, such as size, strength and toughness, generally differ by order.

represents a cognitive distance between the mental schemata of the domain experts and the representation of their knowledge in the decision model. The author found that the human planners and domain experts view the imminence of an order as a *continuously* changing variable. The true semantics of this constraint thus infer some kind of *gradual transition* from absolute non-membership or partial membership [$\mu\{T\} < 1$ at $T = EPST$] to absolute membership [$\mu\{T\} = 1$ at $T = LPST$].

The typical linguistic quantifiers - used by the human planners and domain experts to express the imminence of an order - point to an exponential increase in $\mu\{T\}$ as $T \rightarrow LPST$ [Fig.6.2]. The *extent* to which an order is imminent - as a function of time - is only precisely quantified at the extreme right boundary of the constraint space ($\mu\{T\} = 1$ when $T = LPST$). Thus, the temporal due-date constraint surface $\mu\{T\}$ is essentially a *fuzzy* non-linear function.

6.2. Rule-base evolution

Fig.6.3 illustrates the increase in the size of the rule-base throughout the evaluation period. The most interesting aspect of Fig.6.3 is the *scalloped effect*. The shape of the curve indicates that the process of accumulating knowledge and extending the knowledge base evolved through a sequence of *stages*. The beginning of each stage generally incorporated the acquisition of core knowledge about a particular aspect of the planning decision process, and the rest of the stage involved extending the core - acquiring additional *related* knowledge in the form of new or modified rules - so that the heuristic algorithm could be applied to new situations as they arose.

Three criteria explained the behavioural deficiencies in the prototype:

- the domain experts neglected to express rules to cover all the special cases that arose
- the rules did not produce correct conclusions because they made erroneous assumptions
- some of the rules or advice from the domain experts was overlooked or incorrectly implemented

Consider the temporal delivery-date constraint. Given a particular time t , the human planners reason about the *imminence* of an order based on their own judgement or on particular demands from the external environment. In classical set theory - which forms the foundation of binary heuristic inference schemes - an arbitrary point in time t , must be defined to discriminate between *imminent* and *non-imminent* orders. Since the boundaries between what is in the set and what is outside the set are sharp, these types of constructs are called *crisp sets*. A characteristic function for such a set appears as:

$$\mu[T] = \{ \text{TIME} \geq t_i \}$$

The discriminant or characteristic function for this set reflects its Boolean nature. Transgressing the constraint space along the TIME axis [Fig.6.2] the membership (truth) of the *imminence* decision variable remains FALSE [0] until $T = t_i$ at which point it immediately becomes TRUE [1] - the line connecting membership and non-membership in the membership transition graph is *dimensionless*.

The normalisation of the problem space is based on a *multiple* partitioning of the soft constraint spaces. In effect, this represents a *multiple-valued logic* approach [Fig.6.2]. However, the task space regions remain essentially crisp sets. Such *rigid* partitioning of the constraint spaces




Fig.6.2. The semantics of the temporal delivery-date constraint

Since the binary heuristic inference scheme enforces a logic that is based on a *finite* set of truth values, each task space region must, by necessity, be associated with *crisply defined* constraint sub-spaces (partitions). The ambiguity problem thus remains largely unsolved. Consider two tasks τ_i and τ_j whose input decision variable values differ within the range

$$\Delta t \in [0, \dots, 2] \cap \Delta w \in [0, \dots, 5] \cap \Delta g \in [0, \dots, 0.01] \cap \Delta f \in [0, \dots, 0.1]$$

Although the objective (cost) function assigns individual cost values to τ_i and τ_j these values may not be realistic determinants since $\tau_i \in TS(1 \bullet 1 \bullet 1 \bullet 1)$ and $\tau_j \in TS(1 \bullet 1 \bullet 1 \bullet 1)$ - the relative priorities of constraints have *not* been adequately defined.

Relative ratios

The tuneable weighting parameters $[\rho_1, \rho_2, \rho_3, \rho_4]$ in Prototype A (see 5.4.2) are essentially relative ratios. They were introduced into the cost function in an attempt to resolve the above problem. However, this approach failed because the relative priorities of constraints were required to be *explicitly* and *absolutely* defined *a priori*. Moreover, this approach relied on an assignment of ratings *outside* the decision model itself. Consequently, it did not address the intrinsic relationship between constraints.

6.1.2. Vagueness

Vagueness is associated with an imprecise human perception of the overall problem. More specifically, it relates to the inability of the domain experts to provide precise (soft) constraint information. Zadeh [88] comments on vagueness as follows - *When a point is reached where the cardinality of sub-classes exceeds the information-handling capacity of the human brain, the boundaries of sub-classes are forced to become imprecise.*

The human planners often have to contend with imprecise and unstable soft constraint semantics *and* with a solution space that is combinatorially explosive. Consequently, they perceive the problem from an intuitive or subjective perspective.

was attributed to *subjective factors* involved in soft constraint inferencing. The inherent *uncertainty* in many of the (soft) inference rules expressed by the domain experts presented a major difficulty in the formulation of the decision process. Two main types of uncertainty were encountered. These are similar to the uncertainty categories defined by Klir and Folger [87] i.e. *ambiguity* and *vagueness*.

6.1.1. Ambiguity

Ambiguity results from situations where the choice between two or more instantiations is not clearly defined. Each instantiation in the set A_i has a degree of desirability or utility depending on the *relative importance* (preference) assigned to the different constraints by the human planners and domain experts. The decision on which instantiation to select is a fuzzy one, because each decision relates to an action whose consequence cannot be exactly quantified. It is clear that no instantiation can completely *satisfy* all the constraints. However, the different feasible instantiations *partially* satisfy the constraints to various degrees.

Several approaches have so far been proposed in the literature to model and remove ambiguity [35]. For example, Yager argues that ranking or weighting of objectives can be achieved by linear ordering of objectives, intervals, relative ratios or absolute ratings. These are progressively more difficult to obtain from the domain experts. Forcing the domain experts to provide such information may yield incorrect answers if they simply cannot provide it *a priori*.

Partial ordering of objectives

The implicit prioritisation of tasks through the normalisation of the problem space (see 5.3.2) represents a partial ordering of objectives. The set of soft constraints $[\Delta t, \Delta w, \Delta g, \Delta S]$ was ranked through a depth-first penetration of the task space regions [Fig.5.1]. Thus, the search was propagated - and the set of feasible instantiations was populated - in the following order:

TS(1•1•1•1) → TS(1•1•1•2) → ... TS(4•15•40•10)




Fig.6.1. Degree of human intervention (Prototype A)

The set of z feasible instantiations [$A_i = \{\lambda_{i1}, \lambda_{i2}, \dots, \lambda_{ir}\}$] at any node i in the search was bounded as follows:

$$A_i = A_Q \cap A_D \cap A_R$$

where $A_Q =$ the set of feasible instantiations with respect to hard quality precedence constraints

$A_D =$ the set of feasible instantiations with respect to the *hard barriers* of the soft constraint spaces [$\Delta t, \Delta w, \Delta g, \Delta S$]

$A_R =$ the set of feasible instantiations with respect to hard resource availability constraints

The consistent reduction in human intervention throughout the first half of the evaluation series resulted primarily from modifications or expansion to the set of binary heuristic rules that perform the above operation. However, since $z > 1$ almost without exception, the sequencing problem is essentially *underconstrained*. To differentiate between the z feasible instantiations the *soft* constraints must be additionally considered.

During the evaluation a point was reached where the domain experts could no longer provide consistent and reliable information regarding the relative priorities of constraints. This problem

C. EVALUATION (PROTOTYPE A)

During a tentative implementation period of 8 weeks the steelmaking planners generated 60 production sequences using Prototype A. After every third plan was generated, the model behaviour was analysed and the knowledge base was adjusted or extended to resolve undesired behaviour. The results over the full implementation period are presented here.

6.1. Degree of human intervention

The degree of human intervention is defined here as the ratio of tasks replaced or added by the operator to the total number of tasks sequenced by the prototype, in order to finalise a production sequence. It is a relative measure of the *consistency* and *reliability* of decision support provided by the prototype.

The degree of human intervention reduced consistently as the knowledge base was expanded up to a point where the addition of binary heuristic rules did not seem to have a significant effect on the model behaviour i.e. there was a perceived *stabilisation* in system functionality [Fig.6.1]. This does not mean that the added rules were not necessary - they were necessary to handle new cases and maintain planning integrity.

All prototype-generated sequences were *heuristically TRUE* in that they were legal representations of the knowledge acquired from the external environment. However, human intervention was required in order to reflect real-world scenarios for which the existing rule-base had not made provision. In effect, the decision model *underspecified* the problem.

The binary heuristic inference scheme employed a combination of crisp *AND*, *OR*, *NOT* logical expressions i.e. the inference scheme was based on the premise that any planning decision can be derived from events that are either *TRUE* or *FALSE*, with no intermediate state. Such logical expressions were effectively used to bound the solution space through *hard constraint inferences*.

conditions, constraints must be *selectively relaxed* in order to allow the generation of a feasible solution. Constraint relaxation may occur in a number of ways:

- the operator may modify the graphical meta-plan by selecting alternate partial sequence types or by increasing the $\kappa_{min(\gamma)}$ attribute value of the existing partial sequence type(s)
- the operator may increase the value of the decision threshold coefficient γ_r
- the hard barriers of soft constraints - the allowable range of decision variable values - may be selectively widened

Δs_k = the deviation between the actual stock level of the associated product type p_k and the minimum stock level for p_k at the slab yard (as a ratio)

$\rho_1, \rho_2, \rho_3, \rho_4$ = tuneable parameters in the range [0, 1] representing weights for each decision variable

An additional pruning parameter was applied - in the form of a *decision threshold coefficient* (y_d) - to truncate all paths for which $c_{jk} > y_d$. Guan [86] describes a similar approach. In this investigation y_d was defined as a tuneable parameter linked to a graphical slider on the operator interface. This allowed the operator to dynamically set a band-width on the solution space, corresponding to a shift towards higher or lower constraining values on the decision process.

Constructive backtracking

A backtracking feature - applied in the generation of both partial and complete sequences - handles situations where the cost of the provisional optimum instantiation λ_{jk} is higher than the pre-established threshold y_d . If this happens, it effectively means that the trial task associated with the provisional optimum instantiation cannot be chosen as the j th task in the sequence i.e. the path is effectively blocked. This implies that the provisional optimum instantiation λ_{ij} of the *previous* node - the instantiation with the lowest cost at node i - is not feasible.

In this case, the search returns to the previous node (node i) and the instantiation λ_{ij} with the next lowest cost in the set $\{\lambda_{j(t)}, \dots, \lambda_{j(z)}\}$ of z feasible instantiations is chosen as the provisional optimum λ_{ij} at node i . However, if this instantiation has a cost $c_{ij} > y_d$ - or subsequently leads to $c_{jk} > y_d$ - then it too must be discarded. This procedure is repeated until a (no longer provisional) optimum instantiation is found - for every node in the search - which does not exceed the threshold y_d .

If no such instantiation exists and $i < l$ it effectively means that the state of the constraint set $\{C = \{c_1, c_2, \dots, c_l\}\}$ prohibits the attainment of a solution to the problem i.e. the set of feasible solutions $\{S = \{s_1, \dots, s_l\}\}$ is empty since only a partial solution is possible. Under such

A set of heuristic monitoring rules subjects I to the following hard constraints

- $\kappa_{min(y)} \leq \eta_j \leq \kappa_{max(y)}$

where: $\kappa_{min(y)}$ = the minimum allowable number of tasks in partial sequence y

$\kappa_{max(y)}$ = the maximum allowable number of tasks in partial sequence y

Both $\kappa_{min(y)}$ and $\kappa_{max(y)}$ are conditions imposed by the external environment and depend upon the *type* of partial sequence (see 4.3.2)

- $$\sum_{i=0}^n \sum_{v=0}^m t_{iv} + \sum_{p=0}^r t_p \leq t$$

where: t_{iv} = the processing time for task η_{iv}

t_p = the set-up time for casting interruption p

t = the total available processing time at the caster

The search procedure for generating both partial and complete sequences was based upon the best-first search method [84,85]. Best-first search generates all *feasible* combinations of decision variable values at each node j in the search. Each feasible combination of values - each instantiation λ_{jk} between two successive tasks τ_j and τ_k in sequence - along with the cost of adding the task τ_k to the sequence is then placed in a list for subsequent comparison.

The instantiation that leads to the *lowest cost* - the provisional optimum λ_{jk} - is selected from the list and the associated task τ_k is added to the sequence. This procedure is repeated until j is equal to the total length of the sequence. The cost of a particular instantiation was defined as the weighted sum of the decision variable values:

$$c_{jk} = \rho_1 \Delta t_k + \rho_2 \Delta w_{jk} + \rho_3 \Delta g_{jk} + \rho_4 \Delta s_k$$

where: c_{jk} = the cost of an instantiation λ_{jk}

Δt_k = the deviation between the current real time and the LPST of task τ_k (in days)

Δw_{jk} = the slab width deviation between τ_k and τ_j (in mm)

Δg_{jk} = the deviation in chemical analysis between τ_k and τ_j (as a percentage)

In general the number of potential (partial or complete) sequences is:

$$n = m + m^2 + \dots + m^l$$

where: m = the number of alternatives at each state

l = the sequence length

Although the solution space is combinatorially explosive with respect to the sequence length, the number of potential (partial and complete) sequences evaluated was significantly reduced by applying the set of constraints $[C = \{c_1, c_2, \dots, c_l\}]$. In effect, hard constraint inferences remove all infeasible alternatives as well as their successors from each state-space. Heuristic prioritisation procedures (soft constraint inferences) rule out numerous alternatives by identifying promising paths.

5.4.2. The search method

The initial problem space consists of an unordered set of n tasks representing a steel making production load for a specified time period T . The prototype is expected to generate a definitive sequence of l ($l \ll n$) tasks over a period t ($t < T$) for each caster. The goal is to generate a sequence of tasks that conforms to the meta-plan established by the human planner.

The solution consists of repeatedly searching the problem space until a satisfactory complete sequence is found. The total sequence length l is defined as:

$$l = \sum_{x=0}^n \sum_{y=0}^m \eta_{xy}$$

where n = the number of uninterrupted casting sequences in the meta-plan

m = the number of heuristically defined partial sequences in a particular uninterrupted casting sequence

η_{xy} = the number of tasks in partial sequence y of casting sequence x

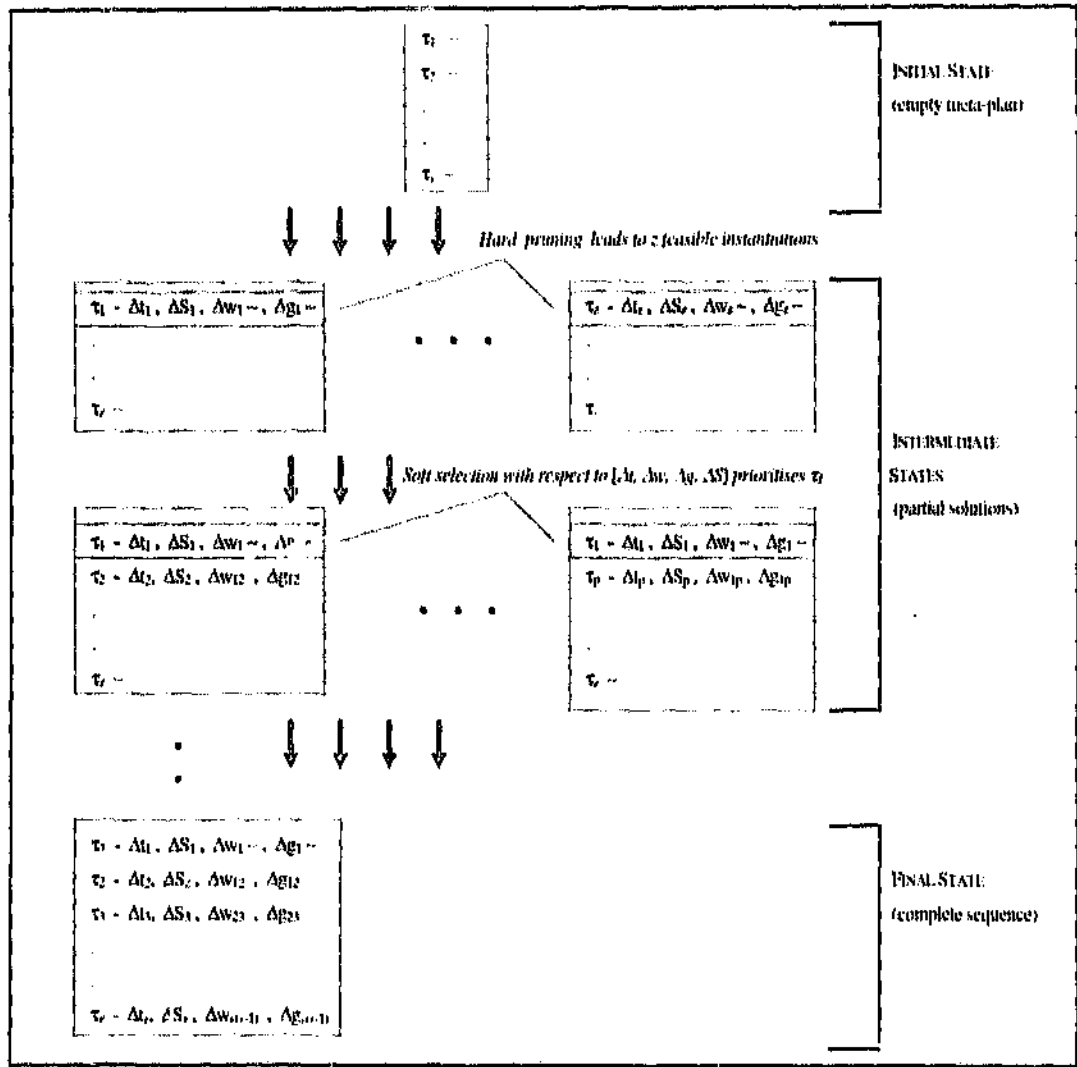


Fig. 5.2. State-space schematic

State-space size

The combinatorial nature of the state-spaces is a function of the number of input decision variables, the range of allowable decision variable values and the number of tasks in sequence.

Table 5.1. The implications of $\tau_i \in TS(\bullet|\bullet|\bullet)$

Input decision Variable	Semantic partition (range of values)	Unit of measure	Implication
Δt	[0, 2]	days	The LPSST of τ_i is less than 2 days into the future; this implies that τ_i has maximum priority with respect to the delivery date constraint.
Δw	[0, 5]	mm	τ_i is favourable with respect to the dimensional-precedence preference constraint, since the slab width of τ_i does not deviate, or deviates negligibly ($< 5\text{mm}$) from the slab width of the previous task in sequence.
Δg	[0, 0.01]	percentage	τ_i is favourable with respect to the quality-precedence preference constraint, since the chemical analysis of τ_i does not deviate, or deviates negligibly ($< 0.01\%$) from the chemical analysis of the previous task in sequence.
ΔS	[0, 0.1]	ratio	The actual slab stock level S_i for product type p_i is approaching the minimum stock level β_i set by the external environment i.e. $(S_i - \beta_i) / \beta_i \leq 0.1$. This implies that τ_i has maximum priority with respect to the stock structural integrity constraint.

The decision model was correspondingly decomposed into two primary reasoning components. A database of heuristically defined partial sequences was initially created, from which a central heuristic algorithm drew to establish complete production sequences. The generation of both partial sequences and complete sequences was based on the principle of *state-space search*.

5.4.1. State-space representation

Fig.5.2 shows a portion of the state-space for a typical sequencing problem. Both the formulation of heuristically defined partial sequences and the establishment of complete sequences followed the procedure shown in Fig.5.2.

The *composite moments or centroid* defuzzification method is the most widely used technique since it has several desirable properties [90]:

- the defuzzified values tend to move smoothly around the output fuzzy region i.e. changes in the fuzzy set topology from one model frame to the next usually results in smooth changes in the expected value
- it is relatively easy to calculate and does not consume prohibitive amounts of processing time.




Fig.7.2. Min-Max method of implication with an output fuzzy region

region to the *minimum* of the predicate truth. The output fuzzy region is updated by taking the *maximum* of these minimised fuzzy sets. These steps are outlined in equations 7.1 and 7.2.

$$\mu_{cfs}[x_i] \leftarrow \min (\mu_{pt}, \mu_{cfs}[x_i]) \quad [7.1]$$

Equation 7.1 indicates that the consequent fuzzy set (*cfs*) is modified before it is used. This modification sets each truth function element to the minimum of either the truth function or the truth of the proposition antecedent.

$$\mu_{sfs}[x_i] \leftarrow \max (\mu_{sfs}[x_i], \mu_{cfs}[x_i]) \quad [7.2]$$

Equation 7.2 indicates that the solution fuzzy set (*sfs*) is updated by taking - for each truth function value - the maximum of either the truth value of the *sfs* or the fuzzy set that was correlated in equation 7.1. These steps result in reducing the effective height of the solution fuzzy set to equal the maximum truth of the predicate and then, using the modified fuzzy region, applying it to the output fuzzy region by using the *OR* (union) operator. When all the propositions have been correlated, the output fuzzy region reflects the contribution from each proposition [Fig.7.2].

7.3.3. Method of defuzzification

To find the corresponding scalar value for the utility *u* of a particular instantiation the value that best represents the information contained in the output fuzzy region *UTILITY* [Fig.7.2] must be found. This process is called *defuzzification* - it yields the expected value of the solution variable for the particular execution of the fuzzy decision model.

The defuzzification methods are a compromise between the need to find a single point result and the loss of information such a process entails. Such information loss is a natural consequence of a reduction in the representational dimensionality of the output fuzzy region.

degree of truth (membership). All those that have some truth contribute to the final state of the output fuzzy region.

The functional relationship between the degree of truth in related fuzzy regions is called the *method of implication*. The functional relationship between fuzzy regions and the expected (output) value of a set point is called the *method of defuzzification* - these methods together constitute the backbone of approximate reasoning.

7.3.1. Fuzzy propositions

The fuzzy decision model contains a series of conditional fuzzy propositions of the form

IF w is Z THEN x is Y

where w and x are scalar values - such as Δt or Δs - and Z and Y are linguistic variables - such as *HIGH* or *URGENT*. The proposition following the *IF* term is the antecedent or predicate. The proposition following the *THEN* statement is the consequent. The statement x is Y is conditional on the truth of the predicate - it may be interpreted as follows

x is a member of Y to the extent that w is a member of Z

That is, the consequent is correlated with the truth of the predicate. The fundamental proposition can be extended with fuzzy connectors

IF $(w$ is $Z) \circ (y$ is $W) \circ \dots \circ (u$ is $S)$ THEN x is Y

where the connector \circ represents (some form of) the *AND* or *OR* operator.

7.3.2. Method of implication

The current literature describes various implication mechanisms, the most common appearing to be the *Min-Max* implication method. This method involves restricting the consequent fuzzy

7.2.2. Decision variables

A decision variable is described in terms of its total constraint space (generally referred to in the literature as the *universe of discourse*). This space is composed of multiple, overlapping fuzzy sets, each fuzzy set describing a semantic partition of the underlying constraint space. The constraint space constitutes the allowable range of decision variable values.

In process engineering models, which are currently the most common applications of fuzzy logic, the conventional space representations are generally either trapezoidal or triangular. The trapezoid usually maps membership functions at the domain extremes while the triangles slice up the fuzzy space into a series of smaller - but well defined - fuzzy regions. The fundamental idea is to capture the linguistic nature of each sub-region - the compatibility of the fuzzy region increases from left to right across the domain until it reaches unity; after this it begins to fall off to zero.

To convert a series of individual fuzzy sets into a continuous surface, each fuzzy set must, to a certain degree, overlap the neighbouring set(s). There is no precise algorithm for determining the minimum or maximum degree of overlap. The interference pattern depends on the semantic nature of the constraint and the intrinsic degree of imprecision associated with the two neighbouring semantic partitions.

7.3. Fuzzy inferencing

Unlike the binary heuristic inference scheme (Prototype A) where statements are executed serially, the principal reasoning protocol in a fuzzy inference schemes is a parallel processing paradigm. In Prototype A pruning algorithms and heuristics were applied to reduce the number of rules examined. In a fuzzy inference scheme *all* the rules are fired. However, some have no degree of truth in their premise and thus do not contribute to the outcome.

The root mechanism in a fuzzy model is the *proposition* (inference rule). These are statements of relationship between the decision variables and one or more output fuzzy regions. When the fuzzy decision model is triggered, a series of conditional fuzzy propositions is evaluated for

The degree of membership is also known as the *membership function or truth function* since it establishes a one-to-one correspondence between an element of the constraint space x and its truth value indicating its degree of membership in the set. The surface of the fuzzy set F is mapped through a *truth generating function* of the form

$$\mu[x] \leftarrow f(x \in F)$$




Fig.7.1. The decision variable *IMMINENCE* and the fuzzy set *LOW*

A fuzzy set encodes the imprecision or fuzziness associated with the semantic constraint description through its control surface. The shape of the curve represents the semantics of the actual constraint, as defined by the domain experts and surmised by the human planners.

The degree of membership is also known as the *membership function* or *truth function* since it establishes a one-to-one correspondence between an element of the constraint space x and its truth value indicating its degree of membership in the set. The surface of the fuzzy set F is mapped through a *truth generating function* of the form

$$\mu[x] \leftarrow f(x \in F)$$

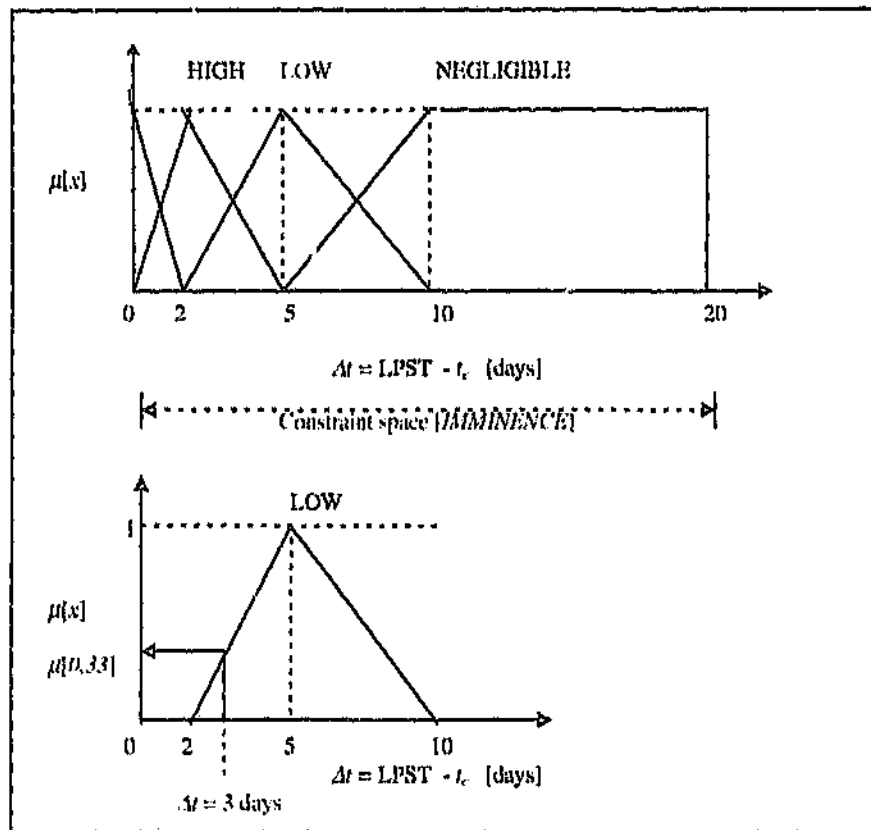


Fig.7.1. The decision variable *IMMINENCE* and the fuzzy set *LOW*

A fuzzy set encodes the imprecision or fuzziness associated with the semantic constraint description through its control surface. The shape of the curve represents the semantics of the actual constraint, as defined by the domain experts and surmised by the human planners.

7. A FUZZY SET METHODOLOGY FOR SOFT CONSTRAINT INFERENCE

7.1. Approximate reasoning

A great deal of work has been done on the theory and application of models which simulate human *approximate* reasoning. Some of the most significant progress has resulted from the application of fuzzy set theory. Fuzzy set theory deals with a set of events that do not have crisply defined membership as in ordinary (conventional) set theory.

7.2. Representing soft planning constraints as fuzzy constraints

The linguistic expressions or quantifiers used by the domain experts to define the semantics of soft planning constraints may effectively be modelled as *linguistic variables*. At its root, a linguistic variable is the name of a *fuzzy set*, directly representing a specific region in the underlying constraint space. Fuzzy models manipulate linguistic variables.

Linguistic variables were found to encapsulate the properties of approximate or imprecise constraint semantics in a systematic and computationally useful manner. They reduce the apparent complexity of real-world planning knowledge by matching *semantic labels* to the underlying constraint spaces.

7.2.1. Fuzzy sets

A fuzzy set consists of three components [Fig.7.1]

- a horizontal axis of monotonically increasing real numbers that constitute the population of the fuzzy set
- a vertical axis between 0 and 1 indicating the degree of membership in the fuzzy set
- the surface of the fuzzy set itself that connects an element (a decision variable *value* in the constraint space) with a degree of membership in the set

Uncertainty in planning knowledge

A knowledge base is a repository of human knowledge. Since most human knowledge is imprecise in nature it generally follows that the knowledge base component of a knowledge-based system is a collection of rules and facts which are often neither totally certain nor totally consistent. In conventional knowledge based systems uncertainty is dealt with through a combination of predicate and probability-based methods [96]. However, there are setbacks to their use in real-world planning domains:

- consistent application of these methods is not computationally feasible
- generally, the necessary inter-dependencies and probability distributions are not known

Moreover, these methods generally cannot resolve the *inherent* uncertainty of the information in the knowledge base. As a result, they are mostly ad-hoc in nature.

The remainder of this report describes an alternative approach to the management of uncertainty in soft constraint inferencing. The approach is based on the use of fuzzy set theory, which forms a theoretical basis for the representation of *approximate reasoning*.

6.5. Interpretation of results

In the current literature one often finds counter claims regarding the (previously published) success of knowledge based planning applications. For example, Kathawala *et al* [89] claim that the results of the *ISIS* [55] application have been mixed - the system works in that reasonable plans are generated, but they are often less satisfactory than was originally hoped. Typically, unsuccessful implementations result from a failure to meet the anticipated final performance requirements. In some cases, the system is not accepted by the end-users or is not used for its intended purpose. Fox *et al* [55] postulate that automated planning - or planning decision support - is often reduced to periodic runs whose outputs provide guidance for future loadings.

Soft constraint inferencing

As with Prototype A unsuccessful implementations often result - at least partially - from the inability of the decision model to effectively resolve the *uncertainty* inherent in soft constraint inferencing. Soft constraint inferencing appeared to depend upon *imprecise* and *unstable* criteria. The related part of the decision process was thus considered as *not definitively repeatable*. Consequently, evaluating the truth or falsity of an *IF* clause according to the Aristotelian binary logic and executing the *THEN* clause in a correspondingly crisp manner proved to be an ineffective approach for soft constraint inferencing. Consider the following binary heuristic inference rule:

IF $0 \leq \Delta t_j < 1 \cap 10 \leq \Delta w_{ij} < 15 \cap 0 \leq \Delta g_{ij} \leq 0.01 \cap 0.2 \leq \Delta S_i < 0.3$
THEN $\tau_j \in TS(1 \cdot 5 \cdot 1 \cdot 3)$

This rule implicitly prioritises a task τ_j by virtue of its membership in the relevant task space region. This occurs prior to the assignment of a cost c_{ij} to the instantiation λ_{ij} . The above rule is based upon a *rigid* semantic partitioning of the soft constraint spaces $\{\Delta t, \Delta w, \Delta g, \Delta S\}$. Coding knowledge in this manner results in a cognitive distance between the decision model and the domain experts as it forces the domain experts to dichotomise their thought patterns at essentially artificial boundaries. In order to effectively model ambiguity and vagueness, an *infinite-valued* logic approach is required.

6.4. Conformance to selected planning objectives

Three quantifiable indicators from the set of planning objectives were used to measure the impact of the prototype on the quality of the planning decision process:

- *quality-precedence cost factor* i.e. the average planned ratio of slab not regarded as prime material [O(1) - approximate improvement ratio ~0.17]
- *sequence factor* i.e. the average number of tasks between casting interruptions [O(2) - approximate improvement ratio ~0.09]
- *the delivery-date achievement ratio* i.e. the ratio of tasks planned on-time to the total number of tasks in sequence [O(3) - approximate improvement ratio ~0.12]

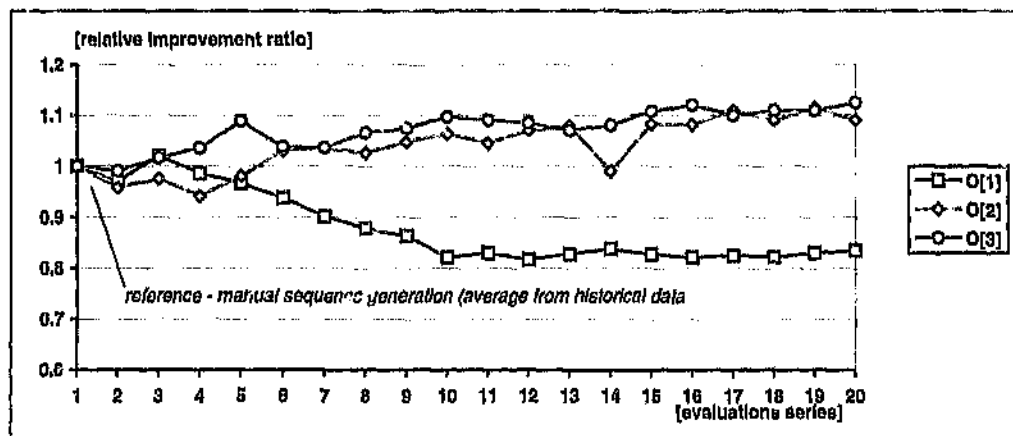


Fig.6.5. Conformance to planning objectives

Moderate gains on all three indicators were realised [Fig.6.5]. In effect, the prototype served to *guide* the human planners in a specific direction, although a large part of the decision process was left entirely to the planners. Thus, although moderate improvements in the decision process were realised, the system was not accepted by the planning operators and supervisors.

The rules were divided into two categories: those of a relatively general nature and those that were specific to the steelmaking domain. The evolution of the rule base showed a general trend towards increasingly domain-specific rules, since it generally followed recurring attempts to deal with unforeseen or unexpected domain-specific irregularities and special cases.

6.3. Sequence generation time

The average time required by the planners to generate production sequences reduced significantly, as shown in Fig.6.4. Prototype A proved to be inefficient with respect to both memory and CPU requirements, primarily as a result of the extensive data search and manipulation cycles employed with the two-phase reasoning approach (see 5.4). The average sequence generation time stabilised at approximately 30 to 40 minutes.

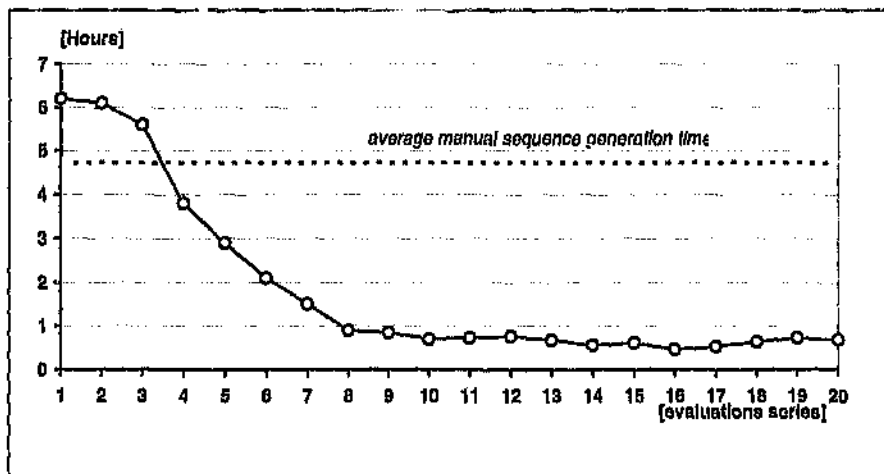


Fig.6.4. Average sequence generation time (Prototype A)

Note that Prototype A initially required more time to generate solutions than the manual sequence generation process itself. As the knowledge base was expanded, pruning algorithms and heuristics were included so that the number of rules fired was reduced, thus, in effect, improving the efficiency of the search procedure. However, the exponential increase in the size of the knowledge base meant that it became increasingly complex and difficult to maintain.

Repeated attempts were made to resolve the behavioural deficiencies:

- new knowledge was progressively added to the knowledge base in efforts to handle additional special cases
- the domain experts modified the original knowledge to correct errors in it




Fig.6.3. Rule-base evolution (Prototype A)

Another interesting aspect of the rule-base evolution was the *utility* of the knowledge for each required improvement in system functionality. That is, the extent to which the rule-base had to be extended to incorporate the new functionality, given the occurrence of new conditions requiring new or modified rules.

The trend was clearly towards an increasingly larger amount of rules required to effect minor improvements in system functionality. As an attempt was made to reduce the level of human intervention - to improve the consistency and reliability of the decision support function - an exponentially increasing set of binary heuristic rules was required to deal with an expanding *uncertain decision zone*.




Fig.9.1. Application of bias (constraint preference) on the Δw and Δt soft decision variables

The two discrepancies in Table 9.1 resulted from *contradictory knowledge* i.e. the reasoning applied by the human planner during the evaluation did not coincide with the way in which the rules had been explicitly defined in the prototype. The rules in question were isolated and modified. The prototype responded by re-assigning the highest u_i to the instantiation selected by the human planner. In both cases the model was rectified within a matter of minutes. The high level of planner-prototype coherence achieved with Prototype B *under controlled test conditions* significantly increased operator confidence in the system.

9.1.1. Response to operator preference

Prototype B allows the operator to place a relative bias on specific planning objectives, via tuneable bias parameters on the operator interface (see 8.3.2). Fig.9.1 shows the results of a sensitivity analysis on the application of bias to decision variables in the prototype. The prototype responds by attempting to reflect the bias in the generated solutions.

The extent to which the prototype incorporates the bias depends upon the characteristic nature of the problem space i.e. no planning constraints are violated, although particular constraints may be relatively compromised or relaxed to reflect the bias entered by the operator.

9.1.2. Sensitivity analysis with respect to fuzzy set morphology

Determining the shape of fuzzy sets is an important part of the modelling process since the fuzzy set morphology determines the correspondence between the decision model and the semantic descriptions of the soft constraints - as specified by the domain experts. If the fuzzy sets are not properly shaped, then the degree of membership may either be too high or too low.

It is reasonable to assume that the semantics of real-world soft constraints are never precisely specified since their specification arises from human thought and cognitive processes.

9. EVALUATION (PROTOTYPE B)

9.1. Evaluations prior to implementation

A series of pre-implementation tests were carried out on the fuzzy decision model to determine how accurately the model emulates the approximate reasoning of the human planners and to calibrate the model where discrepancies occurred. In each test a seed task τ_i and a random series of instantiations were generated. The objective of the test was to determine which instantiation λ_{ij} represented the best follow-on task τ_j for the designated seed task τ_i . First, the human planner selected a follow-on task, after *careful inspection* of the available series of instantiations. The fuzzy decision model was then executed, allocating an utility value u_{ij} to each instantiation. The results are shown in Table 9.1.

Planner-prototype coherence was defined as an exactly matched selection of τ_j between the prototype and the planner i.e. where the prototype allocated the highest u_{ij} to the instantiation selected by the human planner. The same tests were carried out on Prototype A. Clearly, the binary heuristic inference scheme presents a departure from the mental schemata followed by the human planners. In contrast, the fuzzy inference scheme emulated the reasoning process of the human planners with unexpectedly high accuracy and consistency.

Table 9.1. Planner-prototype coherence in the prioritisation of instantiations

Planner-prototype coherence ratio <i>(Evaluation data for tests 1 to 30 in Appendix 8)</i>	Prototype A		Prototype B	
	0.47		0.98	
Discrepancies (Prototype B)	Human Planner		Prototype B	
1. Test 3 <i>(refer to Appendix 8)</i>	Δr_i	1	Δr_i	0
	Δw_{ij}	-50	Δw_{ij}	100
	Δg_{ij}	0.00 <i>(same grade)</i>	Δg_{ij}	0.05
	ΔS_j	0.37	ΔS_j	1.00
2. Test 11 <i>(refer to Appendix 8)</i>	Δr_i	3	Δr_i	1
	Δw_{ij}	0	Δw_{ij}	100
	Δg_{ij}	0.09	Δg_{ij}	0.19
	ΔS_j	0.10	ΔS_j	0.25

The generation of samples was based on a Monte-Carlo simulation of available data. The sample size for the seed tasks (τ_i) was 100 with 15 feasible instantiations (λ_{ij}) per seed task.

8.3.2. Responding to the preferences of the operator

The operator may emphasise a particular preference by tuning the bias parameters on the planning parameter interface. One bias parameter exists for every soft constraint. Each bias parameter has an integer value in the range [1, 2]. The bias parameter is manipulated by the operator via a simplistic graphical slider. The fuzzy decision model accommodates changes in bias dynamically and incorporates these into the selection of tasks for sequencing.

Each bias parameter is associated with a *bias coefficient* δ_f for every fuzzy set f on the related constraint space. The bias coefficient δ_f provides a modified preference measure for the associated semantic partition of the constraint space. When the model is initialised each bias coefficient has a default value of [1.0] and its future value is determined by proportional scaling of the bias parameter on the operator interface. During execution of the fuzzy decision model each singleton is adjusted by the truth of the proposition predicate and scaled by the bias coefficient for each proposition that specifies a particular point in the output space:

$$S_i \leftarrow \left(\prod_{k \in k} \delta_f \right) p_N S_i$$

where $k =$ the set of relevant semantic partitions over the soft constraint spaces $\Delta w, \Delta t, \Delta s, \Delta g$.

8.4. The search method

Prototype B applies essentially the same search method as Prototype A (see 5.4.2). However, only one state-space is considered i.e. partial and complete sequences are concurrently generated. Moreover, the cost c_{ij} of each instantiation λ_{ij} is based upon a fuzzy constraint analysis. The cost c_{ij} is effectively the scalar output utility value u_{ij} determined by executing the fuzzy decision model for the particular instantiation λ_{ij} . The most optimal instantiation is the one associated with the minimum fuzzy mismatch between τ_j and τ_i i.e. the instantiation with the highest u_{ij} .

The degree of utility U is derived from the intersection of all the fuzzy objectives o_k :

$$U = o_1 \cap o_2 \cap \dots \cap o_k$$

or

$$\mu_U = \min \{ \mu_1, \mu_2, \dots, \mu_y \}$$

where μ_U = the degree of membership in the *consequent* fuzzy region

μ = the degree of membership in the relevant semantic region (fuzzy set) of the related fuzzy constraint spaces

y = the total number of semantic regions defined linguistically by the domain experts

Thus, the consequent fuzzy region is restricted to the minimum of the predicate truth. The output fuzzy region is updated by taking the maximum of the minimised singletons:

$$S_k = \max \{ \mu_{U(1)}, \dots, \mu_{U(n)} \} \text{ in the interval } [0, 1]$$

where S_k = the output singleton value

n = the total number of propositions

This is the well-known *min-max* implication method. When all the propositions have been correlated the output space contains a fuzzy region that reflects the contribution from each proposition.

The output fuzzy region is then defuzzified to return a crisp numerical utility value u_{ij} for λ_{ij} .

The centroid method of defuzzification is applied in this investigation:

$$u_{ij} = \frac{\sum_{i=1}^n d_i \mu_i S_i}{\sum_{i=1}^n \mu_i S_i}$$

8.3. The decision model

Prototype B is essentially a fuzzy multi-criteria decision model [91]. Each instantiation λ_j has the same q soft decision variables ($q = [\Delta t, \Delta w, \Delta s, \Delta q]$). For every instantiation, the decision variable values of the node task τ_i are compared with those of another task designated as a trial task τ_j which is randomly selected from the set of z feasible instantiations at that node. The allowable range of each soft decision variable - defined by the hard barriers of the related soft constraint - serves as the universe of discourse over which the fuzzy sets are specified.

A particular crisp (input) value of a decision variable - such as Δw_{ij} - is fuzzified by means of the respective membership function, producing a fuzzy number (possibility measure) μ_{ij} in the interval $[0,1]$. In turn, the procedure is repeated for all q variables.

8.3.1. Inference mechanism

The structure of the fuzzy inference scheme is described as follows [91]:

Let $\lambda = \{\lambda_j | j = 1, \dots, z\}$ be a finite set of feasible alternative instantiations to be compared at any search node t

Let $O = \{o_a | a = 1, \dots, k\}$ be a finite set of implicit planning objectives - derived from a finite sub-set of soft constraints [$C = \{c_b | b = 1, \dots, q\}$] - according to which the desirability or *utility* of a decision is judged.

The goal is to determine the degree of utility u_{ij} for each alternative λ_j with respect to the set of planning objectives. The attainment of a fuzzy objective o_a by alternative λ_j is expressed by the degree of membership $\mu_j [v_j]$ in the related fuzzy set f (v_j is the decision variable value associated with trial task τ_j of instantiation λ_j).

The representation of operator preferences

The human planners may communicate their preference policy to the prototype via a planning parameter interface. By modifying parameter values on the interface, the planners effectively calibrate the decision process according to their actual preference.

8.2.2. The fuzzy planning propositions

Since the fuzzy propositions or rules support the inferential mechanics of the fuzzy decision model, their clarity and comprehension is very important. The fuzzy propositions of the form

IF x is Y AND v is U AND ... AND w is Z THEN S_i

are represented in a tabular format [Appendix 5]. The table may be updated on-line by the operator. The fuzzy propositions are defined by the human planner and approved by the external environment.

8.2.3. The output fuzzy region

The output fuzzy region is represented as a singleton geometry space consisting of 10 sparse singletons [Appendix 6]. In effect, every proposition is given a relative weighting from 1 to 10. This approach pre-empts an intuitively clear comparison between propositions on the part of the planners and the external environment. For example, the human planner can easily relate the two fuzzy propositions

IF Δg is SMALL AND Δt is EXTREME AND Δw is WIDE AND Δs is LOW THEN UTILITY is S₉

IF Δg is SMALL AND Δt is HIGH AND Δw is WIDE AND Δs is LOW THEN UTILITY is S₆

by assigning the relative values of S₉ and S₆ to them.

The representation of operator preferences

The human planners may communicate their preference policy to the prototype via a planning parameter interface. By modifying parameter values on the interface, the planners effectively calibrate the decision process according to their actual preference.

8.2.2. The fuzzy planning propositions

Since the fuzzy propositions or rules support the inferential mechanics of the fuzzy decision model, their clarity and comprehension is very important. The fuzzy propositions of the form

IF x is Y AND v is U AND ... AND w is Z THEN S_i

are represented in a tabular format [Appendix 5]. The table may be updated on-line by the operator. The fuzzy propositions are defined by the human planner and approved by the external environment.

8.2.3. The output fuzzy region

The output fuzzy region is represented as a singleton geometry space consisting of 10 sparse singletons [Appendix 6]. In effect, every proposition is given a relative weighting from 1 to 10. This approach pre-empted an intuitively clear comparison between propositions on the part of the planners and the external environment. For example, the human planner can easily relate the two fuzzy propositions

IF Δg is SMALL AND Δt is EXTREME AND Δw is WIDE AND Δs is LOW THEN UTILITY is S_9

IF Δg is SMALL AND Δt is HIGH AND Δw is WIDE AND Δs is LOW THEN UTILITY is S_6

by assigning the relative values of S_9 and S_6 to them.

8. A FUZZY INFERENCE SCHEME (PROTOTYPE B)

8.1. The planning procedure

The planning procedure on Prototype B follows the same approach as on Prototype A (see 5.2).

8.2. Knowledge representation

The knowledge representation method for Prototype B is based on the same principle as for Prototype A (see 5.3). However, various interactive knowledge representation screens were added to the prototype. These screens allow the operator to modify or update any element of knowledge on-line.

8.2.1. Soft constraints and operator preference

All soft constraints are represented as fuzzy constraint spaces. The fuzzy spaces are represented graphically on-screen [Appendix 4]. The operator may modify the shape of any fuzzy set by selecting the appropriate fuzzy control surface with a mouse and dragging the surface across the screen. The prototype does not allow the operator to create illegal fuzzy set shapes.

The planning objectives are implied from the semantic representation - the fuzzy set shapes - of the related constraints and the fuzzy proposition table [Appendix 5]. For example, consider the temporal delivery-date constraint. As the current time approaches the LPST of any task τ_j the urgency of the task is perceived to increase.

For instance, if the LPST is less than two days into the future the fuzzy set EXTREME will have a relative degree of truth i.e. $u[T] > 0$ ($T < 2$). The reader will note - in Appendix 5 - that all the propositions for which EXTREME is an antecedent have a relatively high singleton value [S_5 - S_{10}]. Thus, the instantiation λ_j related to a trial task τ_j will have a relatively high utility value u_j even if other criteria - such as stock integrity and dimensional precedence factors - do not contribute to its utility.

7.4. Probability and Possibility

At the mathematical level, fuzzy (truth) values are commonly misunderstood to be probabilities, or fuzzy set theory is interpreted as a new way of handling probabilities. This is not the case. A minimum requirement of probabilities is additivity i.e. probabilities must add up to 1 or the integral of their density curves must be 1. This does not hold in general with membership grades (truth functions).

Semantically, the distinction between fuzzy set theory and probability theory translates to the notions of probability and a degree of membership. Probability statements are concerned with the likelihood of outcome of well-defined events. With fuzziness (imprecision) one cannot express unequivocally whether an event occurred or not - the objective is to model the *extent* to which an event occurred. In effect, fuzzy set theory is concerned with the imprecision that is an *intrinsic* part of the event or concept - it attempts to characterise imprecision or undecidability *within the control structure itself* rather than in the outcome of the model. Moreover, fuzzy systems can represent a more complex *possibility* space than probabilities [90] for the following reasons:

- their cumulative distributions can sum to less than 1 or more than 1 - this is due to the interpretation of a fuzzy set as a possibility rather than a probability
- they are independent of *a priori* statistical frequencies; they provide a more intuitive method of expressing concepts for which probability distributions are unknown or unattainable
- they are able to reflect the imprecision in probabilities themselves since fuzzy models can reduce contradictory (conflicting) solution states to a fuzzy surface
- they can more nearly represent decision variables; they provide a mathematically sound and semantic-based modelling capability at a high level of abstraction where changes in the system are reflected through linguistic modifications

Since piecewise linear interpolation between the singleton points is used to construct an output fuzzy set representation, defuzzification in singleton and fuzzy set based models is essentially equivalent. In fact, the centroid defuzzification method is simplified:

$$u = \frac{\sum \mu S}{\sum \mu}$$

Singleton representations generally provide much faster defuzzification - because the entire area need not be integrated - but since the output is represented by a set of sparse points, the precision can often be low.




Fig.7.3. Min-Max method of implication with an output singleton geometry

The centroid technique effectively finds the *balance* point of the output fuzzy region by calculating its weighted mean. Arithmetically, this is formulated as:

$$u = \frac{\sum_i d_i \mu(d_i)}{\sum_i \mu(d_i)}$$

where d_i is the i th point of the solution space (fuzzy region) and $\mu(d_i)$ is the truth membership value for that point - in effect, the centroid technique finds the centre of gravity of the output fuzzy region.

7.3.4. Singleton geometry output spaces

The output fuzzy regions are generally represented as a series of fuzzy sets. However, a singleton geometry output space may be used instead. In such a case, the terms associated with the solution fuzzy sets are represented as single vertical points instead of fuzzy set membership functions [Fig.8.3]. The propositions are expressed in a manner identical to the fuzzy set based model

IF (x is Z) \circ (y is W) \circ ... \circ (u is S) THEN S_i

where S_i is any of the singleton support terms. When the model is initialised, each singleton point has a value [1.0] and its future value is determined by proportional scaling due to the predicate truth function

$$S_i \leftarrow p_n S_i$$

where p_n is the predicate truth value. In effect, the topology of the output fuzzy region is determined by a set of scaled vertical singletons. Each singleton is adjusted by the truth of the proposition predicate for each proposition that specifies a particular point in the output space [Fig.7.3].

Table 10.1. Comparative analysis (Prototype A vs Prototype B)

CRITERIA	PROTOTYPE A	PROTOTYPE B
Conformance to selected planning objectives: <ul style="list-style-type: none"> • quality precedence cost factor • sequence factor (average sequence length) • delivery-date achievement ratio 	<i>(improvement ratio on historic average)</i> 0.17 0.09 0.12	<i>(improvement ratio on historic average)</i> 0.31 0.23 0.42
Degree of human intervention <i>(to finalise production sequences)</i>	~25 %	~5 %
Average sequence generation time	~27 minutes	~7 minutes
Number of inference rules or propositions <i>(excluding problem space normalisation inferences in Prototype A)</i>	182	144
Rule-base evolution	Exponential	Stable
Hard constraint representation	Graphical objects and binary heuristic rules	Retained from Prototype A
Soft constraint representation and inferencing	<ul style="list-style-type: none"> • problem space normalisation • binary heuristic prioritisation rules 	<ul style="list-style-type: none"> • fuzzy sets • fuzzy propositional rules of inference
Acceptability / Credibility <i>(refer to Appendix 1)</i>	LOW (not accepted by planning operators and supervisors)	HIGH (fully accepted by planning operators and supervisors)
Maintainability	LOW (complex and opaque knowledge base)	HIGH (transparent and intuitively clear knowledge base)
Planner-prototype coherence ratio <i>(degree of coherence in the prioritisation of instantiations under controlled test conditions)</i>	0.47	0.98

10.3. Additional benefits (fuzzy decision modelling)

Apart from improvements in the prototype behaviour, various additional benefits were accrued with the fuzzy decision modelling approach.

10.3.1. Reduced model complexity

The fuzzy inference scheme proved to be well suited to modelling the highly complex, non-linear problem spaces encountered at the steelmaking domain. The fuzzy inference scheme

The knowledge base had to be consistently updated and expanded as new and unexpected conditions arose. In effect, the decision model *underspecified* the problem. The author concluded that a prohibitively large number of binary heuristic inference rules would be required to secure consistent and reliable decision support.

In addition, the human planners were found to resist a model on which they could not effectively impose their variable decision criteria. The rigidity of the binary heuristic inference approach did not allow for such a capability within the realm of a reasonably contained and transparent knowledge base. Every modification to soft constraint semantics or inferential propositions required code changes to the rule base.

10.1.2. Prototype B

The fuzzy inference scheme emulated the reasoning process of the human planners with unexpected accuracy. Moreover, the knowledge base was found to be highly transparent and maintainable. The fuzzy decision model required fewer rules and the knowledge representation method was intuitively clear to the operators. Under simulated conditions, the system matched the decisions of the human planners with almost absolute precision. Under practical conditions, the system provided consistent and reliable planning decision support. Prototype B was fully accepted by the planning operators and supervisors.

The superior performance of Prototype B [Table 10.1] was attributed to fuzzy representational and prioritisation techniques for reasoning with soft planning constraints. The robustness - the consistence and reliability - of the inference scheme made it possible to introduce almost any constraint - or constraint semantic property - that was deemed appropriate by the domain experts.

10.2. Comparative analysis (Prototype A ↔ Prototype B)

Table 10.1, compares the results obtained with Prototype A and Prototype B.

10.1.1. Prototype A

Although moderate improvements in the decision process were realized [Table 10.1] the prototype was *not* accepted by the end-users. The seemingly paradoxical nature of these results was attributed to the inability of the binary heuristic inference scheme to model uncertainty in soft constraint inferencing.

The following problems were encountered during the formulation of the decision process:

- The planning process was virtually undocumented which caused a high degree of dependency on the *private* knowledge of the domain experts. This knowledge is held in a largely intuitive, undefined format. It was generally communicated in *linguistic* expressions or, at the very least, linguistic quantifiers.
- The external environment consisted of *multiple* domain experts from various sectors. Consequently, the planning knowledge was often conflicting, inconsistent and unreliable.
- Various complex interrelationships existed. Knowledge concerning specific aspects of the production process was dispersed and not easily clarified by the relevant domain experts

Prototype A failed primarily because the binary heuristic inference scheme enforced a *rigid* semantic partitioning of soft constraint spaces. The domain experts were required to explicitly and absolutely define the partition boundaries. The consequent fragmentation of the planning knowledge resulted in:

- a collection of inference rules that provided neither consistent nor reliable results
- an unnecessary multiplication of (inference) rules

The binary heuristic inference rules presented a cognitive distance between the *mental schemata* of the domain experts and the manner in which their knowledge was represented in the decision model. According to Suh *et al* [93] most of the problems concerning unsuccessful developments of knowledge based systems stem from non-technical issues such as cognitive problems, rather than from purely technical issues.

10. SUMMARY AND CONCLUSIONS

Using the steelmaking domain as a real-world case study, the author observed that the human planners spend a large part of their time consulting with various persons in the external environment to determine *how* particular soft constraints should affect the selection of tasks for sequencing.

The external environment essentially consists of all domain experts with whom the planner consults during the generation of production sequences. These persons impose particular conditions on the planning decision process and may moreover provide real time information that could alter a production sequence proposed by the planner. The external environment for the steelmaking domain was found to consist of:

- sales employees and clients - with whom the planner consults on delivery-date problems and order prioritisation issues
- line supervisors, production engineers and line operators at the steelmaking plant - with whom the planner consults on sequencing precedence and preference issues as well as resource availability restrictions
- engineers and specialists in quality control and manufacturing operations - with whom the planner consults on production quality issues.
- line supervisors at downstream processing units - with whom the planner consults on delivery-date problems and the structural integrity of work-in-process (slab stock)
- senior management - who impose specific organisational goals on the planning decision process

10.1. Acceptance of the prototype

Heizer and Rencher [92] suggest that the planning operators and supervisors want to understand how and why the models on which they are basing important decisions work. Vollmann et al [3] indicate that the logic of the applied method or technique must be transparent to gain acceptance. These factors provided the motivation for the knowledge-based approach followed in this investigation.

variables [Δd , Δw , Δs , Δg] and the utility or desirability of the trial task τ . The certainty factor is simply tacked on to the consequent.

In effect, both predicate and probability-based methods rely on an assignment of certainty measures *outside* the model itself. In contrast, fuzzy logic represents uncertainty and imprecision as an *intrinsic* part of the decision model. In this respect, the fuzzy decision model provides a better, more consistent and more mathematically sound method of managing uncertainty in planning knowledge. The fuzzy decision model produces an estimated *utility* for each instantiation λ_j with a degree of membership (truth) in the relevant consequent fuzzy region. Note that the bias coefficients in prototype B (sec. 8.3.2) serve only to increase the degree of membership in the relevant consequent fuzzy region.

The degree of membership represents the compatibility of the decision model with the belief in the implication function between the decision variable values (associated with λ_j) and the output utility value u_j . No such compatibility exists with certainty measures. While the fuzzy decision model predicts and generates an answer, certainty measures are applied to an answer that must already be known or anticipated.

A single part of the decision model, which may have only a small effect on the output, could be tested so that it was the only part affecting the output - this was how all pre-evaluation tests were conducted.

Improved handling of uncertainty

The handling of uncertainty in knowledge-based decision support systems is an area of continual debate. The methods employed by conventional knowledge-based systems are Bayesian probabilities or some form certainty (confidence) factors. These methods imply that the domain experts must supply the *prior* conditional probability - or certainty measure - that particular decisions or inferences will be observed when a specific instantiation is encountered. Usually, these values are not known *a priori* and the domain experts thus have to assume them.

Probability distributions cannot generally be applied in real-world planning decision support - it was certainly not possible in this investigation - because we are concerned with the degree of truth and *not* with the probability of outcome. Moreover, the statistical data required to define distributions generally does not exist or cannot be trusted in real-world domains.

Alternatively, while certainty factors proved their usefulness in early knowledge-based applications [90] they are an essentially ad-hoc approach to belief management, often subject to unpredictable interpretations by the domain experts and human planners. In addition, this approach may lead to a prohibitively complex decision model since a large number of certainty measures have to be defined and maintained. Consider the hypothetical rule for a particular instantiation λ_j :

IF $0 \leq \Delta d < 1 \cap 10 \leq \Delta w < 15 \cap 0 \leq \Delta g \leq 0.0' \cap 0.2 \leq \Delta r < 0.3$
THEN $\tau_j \in \text{TS}(1.05 \cdot 1.03)$ [CF = 0.75]

As with the relative ratios $[\rho_1, \rho_2, \rho_3, \rho_4]$ applied in Prototype A, the certainty factor of 0.75 says little about the intrinsic relationship between the set of soft constraints or *decision*

This had two important side benefits

- the decision model could generally be modified with fewer induced errors.
- the relative simplicity of the model meant that logical or structural problems could be located and fixed in a minimum amount of time.

The same ease of maintenance and understandability also meant that the model could be validated with greater precision and for a wider variety of input cases. This significantly increased operator confidence in the model. Moreover, the prototype allowed the operators to modify virtually *all* planning knowledge *on-line*. The operators were thus able to *effectively* impose their variable decision criteria on the model.

The ability to model conflicting knowledge

In the current literature, there is almost always an unstated assumption that one expert exists or that all the experts in the field are in complete agreement. In the real world of decision modelling, this is not the case. Real-world planning decisions have no simple solution and involve *conflicting views* from domain experts in the various sectors of the external environment. The fuzzy inference scheme is well suited to representing and reasoning with conflicting knowledge. This was a major contributing factor in:

- the improved consistency and reliability of prototype generated solutions
- the reduced complexity of soft constraint representation and inferencing

Improved control of the decision process

The parallel processing structure of the fuzzy decision model - the evaluation of propositions in parallel - proved advantageous from a prototype development, coding and calibration perspective. The fuzzy propositions were easily understood by the human planner and the domain experts, because they are formulated on the basis of intuitive reasoning. The planner could easily interpret the effect or outcome of each proposition. The inferences associated with each proposition could be tested individually.




Fig.9.5. Conformance to planning objectives (Prototype B)

9.3. Interpretation of results

A significant benefit of the fuzzy inference scheme was the ability to encode knowledge directly in a manner that is close to the way that the domain experts and human planners themselves reason about the decision process. This proved to be the primary failure of the binary heuristic inference scheme - the domain experts were forced to deconstruct their expertise into fragments of knowledge. This process led to an unnecessary multiplication of rules. Moreover, it severely undermined the ability of the human planners to effectively articulate a solution to complex sequencing problems.

With the ability to directly model imprecise information, the fuzzy inference scheme reduced the overall cognitive distance in the modelling process. The knowledge acquisition process was easier, more reliable, and less prone to unrecognised errors or ambiguities.

Reduced model complexity and improved system maintainability

The fuzzy inference scheme required fewer rules than the binary heuristic inference scheme and these rules are closer to the way knowledge is expressed in natural language.

containing one element for every possible combination of decision variable values. The matrix is updated by executing the fuzzy decision model for every one of these instantiations. During actual sequence generation the utility value u_{ij} for any particular instantiation λ_{ij} is directly located via an implicit transfer function i.e. the array-index for the correct u_{ij} is a function of the decision variable values defining λ_{ij} .




Fig.9.4. Average sequence generation time

9.2.3. Conformance to planning objectives

The fuzzy inference scheme improved the conformance of planning decisions to the requirements of the external environment [Fig.9.5]. Three indicators from the set of planning objectives were used for this evaluation:

- *quality precedence cost factor* i.e. the average planned ratio of slab not regarded as prime material [O1] - approximate improvement ratio ~0.31]
- *sequence factor* i.e. the average number of tasks between casting interruptions [O2] - approximate improvement ratio ~0.23]
- *the delivery-date achievement ratio* i.e. the ratio of tasks planned on-time to the total number of tasks in sequence [O3] - approximate improvement ratio ~0.42]

9.2. Evaluations after implementation

The planners generated 60 production sequences using Prototype B. The results over the full evaluation period are presented here.

9.2.1. Degree of human intervention

The level of human intervention reduced significantly and remained relatively *consistent* throughout the evaluation period - Fig.9.3 clearly shows an increase in the consistency and reliability of prototype generated solutions with respect to Prototype A.

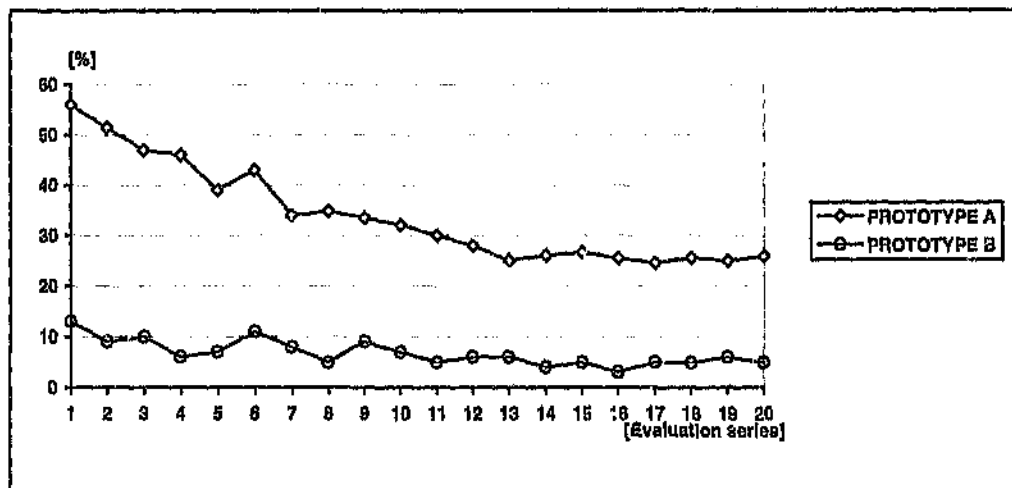


Fig.9.3. Degree of human intervention

9.2.2. Sequence generation time

The fuzzy inference scheme completes processing without involved calculations and with a significantly reduced rule-base. Consequently, the sequence generation time was significantly reduced [Fig.9.4].

During the evaluation period the processing time was further reduced through the design and implementation of a *fuzzy decision matrix*. This matrix is effectively a multi-dimensional array




Fig. 9.2. Subjective variations on the semantics of Δt

Consequently, there are *subjective factors* involved in fuzzy specification and various domain experts may moreover disagree on the exact semantic nature or partitioning of the constraint spaces.

Prototype B applied either trapezoidal or triangular fuzzy set shapes [Appendix 4]. A *subjective* fuzzy set elicitation technique was used, in the sense that the domain experts were allowed to define the fuzzy set co-ordinates based on an intuitive understanding of the constraint space semantics. More advanced techniques may be used to elicit fuzzy set shapes. For example, a neural network may provide a sophisticated non-linear separability analysis on large quantities of historical data. Neural systems have been used to find natural membership functions in data and thus directly create fuzzy regions. In order to apply this technique with confidence, the underlying data must evidently be *reliable* and *well-documented* - this was not the case at the steelmaking domain.

Alternatively, if there is reason to believe that a complete (or partial) mathematical model is applicable for a particular decision variable, then mathematical surface sampling may be used. By simulating the decision process and randomly sampling the control surface of the constraint space, a relationship between the perceived constraint semantics and one or more fuzzy sets may be determined. Techniques such as the root locus method, frequency response plots, Bode diagrams and polar plots provide sampling of the active constraint space surfaces. Although such techniques may provide a mathematically structured approach for representing constraint space semantics, there is no reason to believe that they will significantly improve the functionality of the decision model.

Fig.9.2 shows the results of a sensitivity analysis on the prototype, with respect to the fuzzy set morphology. The results show a high tolerance for fuzzy shapes that are not precisely drawn. This contributes to the fundamental *robustness* of the fuzzy decision model i.e. it contributes to the inherent ability of the model to provide consistent and reliable decision support even where the semantics of soft constraints are imprecisely specified.



To whom it may concern.



Iscor Steel
Flat Products
Vanderbijlpark Work
Deifos Boulevard
Vanderbijlpark 1920
PO Box 2
Vanderbijlpark 1900
Republic of South Africa
Tel: (016) 889 9111
Fax: (016) 889 2097

We have been assisting Mr. LM Besteiro to develop a decision support system for the development of steelmaking plans on a daily basis.

Although every effort was made (by ourselves and all our colleagues involved with steel-making planning process) to provide the correct planning information, the old system did not meet our performance requirements. There were several infrastructural problems and shortcomings in the system which did not make it a practical system for us to plan on.

The upgrading of the old system to the new system has resolved these problems. We fully support the new system. It is a flexible system and we are able to enforce changing and real situations on it. It is also a visual and effective planning tool.



Supervising Planner
Steelmaking Planning

970109
Date.

Flat Products
Vanderbijlpark Works



Iscor Limited
Reg No 89 02164 06

Directors
HJ Smith
(Executive Chairman)
BC Alberts
WAH Clewlow
TL de Beer
CT Fenton
JJ Geldenhuys
JH Kahn
Dr D Konar
M Macdonald
RWV Robertson
JL van den Berg
LL van Niekerk

APPENDIX 1: RESPONSE FROM THE PLANNING SUPERVISOR

88. L.A. Zadeh (1973) *Outline of a New Approach to the Analysis of Complex Systems and Decision Processes*. IEEE Transactions on Systems, Man and Cybernetics Vol.3(1), 28-45.
89. Y. Katiawala and W.R. Allen (1993) *Expert Systems and Job Shop Scheduling*. International Journal of Operations and Production Management, Vol.13(2), 23-35.
90. E. Cox (1995) *The Fuzzy Systems Handbook*. Academic Press, London.
91. I.S. Shaw and M. Besteiro (1996) *Production Scheduling by means of a Multi-attribute Decisionmaking Model*. Internal Report No. END-363, Rand Afrikaans University, Johannesburg.
92. J. Heizer and B. Render (1988) *Production Operations Management*. Allyn and Bacon, Boston.
93. K. Suh *et al* (1993) *Using Human Factor Guidelines for Developing Expert Systems*. Expert Systems Vol.10, 151-156.
94. D. Dubois and H. Prade (1980) *Fuzzy Sets and Systems: Theory and Applications*. Academic Press, New York.

74. K. Stohl *et al* (1991) *Development of an Expert System for a Steel Plant*. IFAC' Expert Systems in Mineral and Metal Processing, 39-44.
75. H. Hamada *et al* (1992) *Construction of a Works Production System Using AI*. NKK Technical review No.65, 10-19.
76. S. Nakayama *et al* (1991) *Combinatorial Scheduler: Simulation and Optimisation Algorithm*. Proceedings of the 1991 Winter Simulation Conference (ed. B.L. Nelson *et al*), 280-288.
77. W. Slany *et al* (1994) *Vague Data Management in Production Process Scheduling Applied to High-Grade Steelmaking*. Proceedings of the First International Conference on AI Planning Systems. Morgan Kaufmann Publishers, San Mateo.
78. K.P. Bernatzki (1994) *Operations Planning and Time Management in the Huckingen Steelworks With Fuzzy Technology* (German: Disposition und Zeitwirtschaft im Stahlwerk Huckingen unter Einsatz von Fuzzy-Technologie). Stahl und Eisen Vol.114(5), 89-95.
79. N. Fowler *et al* (1995) *The ARPA-Rome Knowledge-Based Planning and Scheduling Initiative*. IEEE Expert [February 1995], 4-9.
80. J. Stillman *et al* (1995) *Developing New Technologies for the ARPA-Rome Planning Initiative*. IEEE Expert [February 1995], 10-16.
81. D. Ameling (1990) *The Present State of Production to Order With Reference to Flat Products* (German: Stand der Auftragsabwicklung am Beispiel der Fenbleckfertigung). Stahl und Eisen Vol.110(4), 67-74.
82. *Gensym Corporation*, 125 Cambridge Park Drive, Cambridge, MA 02140 (Tel. 617-547-2500; Fax. 617-547-1962)
83. *Knowledge Based Engineering*. PO Box 786748, Sandton 2146, South Africa (Tel. 011-884-0510; Fax. 011-884-0471)
84. A. Barr and E.A. Feigenbaum (1981) *The Handbook of Artificial Intelligence (Vol.1)*. Addison-Wesley, New York.
85. E. Rich (1983) *Artificial Intelligence*. McGraw-Hill, New York.
86. Q. Guan (1994) *Extending Constraint Satisfaction Problems With Fuzzy Set Theory*. Ph.D Thesis, Technical University of Vienna.
87. G.J. Klir and T.A. Folger (1988) *Fuzzy Sets, Uncertainty and Information*. Prentice Hall, New Jersey.

61. M.S. Steffen and T.J. Greene (1986) *A Prototype System for Scheduling Parallel Processors*. Proceedings of the 1986 IEEE International Conference on Robotics and Automation, 910-917.
62. S. Subramanayam and R.G. Askin (1986) *An Expert Systems Approach to Scheduling in Flexible Manufacturing Systems*. FMS: Methods and Studies (ed. A. Kusiak). Elsevier, New York.
63. V. Chiodini (1986) *An Expert System for Dynamic Manufacturing Rescheduling*. Symposium on Real Time Optimisation in Automated Manufacturing Facilities. National Bureau of Standards Publications 724, 357-372.
64. D. Ben-Arieh (1986) *Knowledge-Based Control system for Automated Production and Assembly*. Modelling and Design of Flexible Manufacturing Systems (ed. A. Kusiak). Elsevier, New York.
65. D. Ben-Arieh *et al* (1989) *Knowledge-Based Scheduling Under Unpredicted Conditions: Two Approaches*. International Journal of Production Research Vol.27(5), 869-882.
66. G. Bruno *et al* (1986) *A Rule Based System to Scheduling Production*. IEEE Computers Vol.19, 32-40.
67. P.A. Newman (1987) *Scheduling in CIM Systems*. Artificial Intelligence: Computer Integrated Manufacturing (ed. A. Kusiak). IFS, New York.
68. S.F. Smith (1987) *A Constraint-Based Framework for Reactive Management of Factory Schedules*. Proceedings of the International Conference on Expert Systems and Leading Edge in Production Planning and Control, South Carolina.
69. S. Jain, K. Barber and D. Osterfeld (1990) *Expert Simulation for On-Line Scheduling*. Communications of the ACM, Computers & Industrial Engineering Vol.33, 55-60
70. R.M. Kerr and R.V. Ehsary (1988) *Implementation of an Expert System for Production Scheduling*. European Journal of Operational Research Vol.33, 17-29.
71. S.M. Alexander (1987) *An Expert System for Selection of Scheduling Rules in a Job Shop*. Computers & Industrial Engineering Vol.12(3), 167-171.
72. J.E. Bingle and L.J. Wink (1989) *Expert Systems Can Do Job Shop Scheduling: An Exploration and a Proposal*. Computers & Industrial Engineering Vol.17(1-4), 347-362.
73. E. Bensana *et al* (1986) *An Expert Systems Approach in Industrial Job Shop Scheduling*. Proceedings of the 1986 IEEE International Conference on Robotics and Automation, 1645-1650.

48. H. Muller, S.D. Samblancx and D. Matthys (1987) *The Expert System Approach and the Flexibility-Complexity Problem in Scheduling Production Systems*. International Journal of Production Research Vol.25(11), 1659-1670.
49. K. Kempf, C.L. Pape, S.F. Smith and B.R. Fox (1990) *Issues on the design of AI-Based Schedulers: A Workshop Report*. AI Magazine Special Issue, 37-55.
50. J.D. Price, J.C. Malley and P.W. Balsmeier (1994) *Expert Systems: Application to Inventory Control and Production Management*. Information Management [September 1994], 26-30
51. R.O. Duda and E.H. Shortliffe (1983) *Expert Systems Research*. Science [April 1983], 261-268.
52. J.Kriz (1987) *Knowledge-based Expert Systems in Industry*. Ellis-Horwood Limited (Wiley), New York.
53. A. Kusiak and M. Chen (1988) *Expert Systems for Planning and Scheduling Manufacturing Systems*. European Journal of Operations Research Vol.34, 113-130.
54. Y.P. Gupta and D.C.W. Chin (1989) *Expert Systems and their Applications in Production and Operations Management*. Computers and Operations Research Vol.16(4), 363-378.
55. M.S. Fox and S.F. Smith (1984) *ISIS: A Knowledge-Based System for Factory Scheduling*. Expert Systems Journal Vol.1(1), 25-49.
56. D.E. O'Connor (1984) *Using Expert Systems to Manage Change and Complexity in Manufacturing*. Artificial Intelligence Applications for Business (ed. W. Reitman). Norwood, New Jersey.
57. J.H. Robbins Jr. (1985) *PEPS: The Prototype Expert Priority Scheduler*. Technical Paper MM-1093, Society of Manufacturing Engineers.
58. J.T. Tou (1985) *Design of Expert Systems for Integrated Production Automation*. Journal of Manufacturing Systems Vol.4(2).
59. M.J.P. Shaw and A.B. Winston (1986) *Applications of AI to Planning and Scheduling in Flexible Manufacturing*. FMS: Methods and Studies (ed. A. Kusiak). Elsevier, New York.
60. J. Erschler and P. Esquirol (1986) *Decision-Aid in Job Shop Scheduling: A Knowledge-Based Approach*. Proceedings of the 1986 IEEE International Conference on Robotics and Automation, 1651-1656.

34. J.D. Wiest (1969) *A Management Guide to PERT/CPM: With GERT/PDM/DCMP and Other Networks*. Prentice-Hall, New York.
35. W. Slany (1996) *Scheduling as a multiple criteria optimisation problem*. *Fuzzy Sets and Systems* Vol.78(2), 197-222.
36. E.W. Davis and J.H. Patterson (1975) *A Comparison of Heuristic and Optimal Solutions in Resource-Constrained Scheduling*. *Management Science* Vol.21, 944-955.
37. J.P. Ignizio (1980) *Solving Large-Scale Problems: A Venture Into a New Dimension*. *Journal of the Operational Research Society* Vol.31, 217-225.
38. T. Yang, J.P. Ignizio and J. Song (1989) *An Exchange Heuristic Algorithm for Scheduling With Limited Resources*. *Engineering Optimisation* Vol.14, 189-205.
39. T. Yang, J.P. Ignizio and E.E. Deal (1989) *An Exchange Heuristic for generalised Job Shop Scheduling*. *Engineering Optimisation* Vol.15, 83-96.
40. F.F. Boctor (1990) *Some Efficient Multi-Heuristic Procedures for Resource Constrained Scheduling*. *European Journal of Operations Research* Vol.49, 3-13.
41. M.M. Khattab and F. Choobineh (1991) *A New Approach for Project Scheduling With Limited Resources*. *International Journal of Production Research* Vol.29(1), 185-198.
42. R.W. Eglese (1990) *Simulated Annealing: A Tool for Operational Research*. *European Journal of Operational Research* Vol.46, 217-281.
43. M. Widmer (1991) *Job Shop Scheduling With Tooling Constraints: A Tabu Search Approach*. *Journal of the Operational Research Society* Vol.42(1), 75-82.
44. K.D. Schnelle and R.S.H. Mah (1994) *Product Quality Using a Real-Time Expert System*. *ISIJ International* Vol.34(10), 815-821.
45. G.A. Cleveland and S.F. Smita (1989) *Using Genetic Algorithms to Schedule Flow Shop Releases*. *Proceedings - International Conference on Genetic Algorithms*, 160-169.
46. U. Canzian, G. Guida, W. Poloni and S. Pozzi (1989) *CRONOS-II: A Knowledge-Based Scheduler for Complex Manufacturing Environments*. *IEEE 2nd International Conference on Data and Knowledge Systems for Manufacturing and Engineering*, 76-83.
47. E.E. Deal, J.G. Chen, J.P. Ignizio and V. Jeyakumar (1990) *An Expert System Scheduler: Some Reflections on Expert System Development*. *Computers and Operations Research* Vol.17(6), 571-580.

15. H. Emmons (1986) *Scheduling and Sequencing Algorithms*. Production Handbook (edited by J.A. White). Wiley, New York.
16. V.B. Godin (1978) *Interactive Scheduling: Historical Survey and State of the Art*, AIIE Transactions Vol.10, 331-337.
17. S.C. Graves (1981) *Review of Production Scheduling*. Opns. Res. Vol.29, 646-675.
18. J.M. Moore and G.L. Wilson (1967) *A Review of Simulation Research in Job Shop Scheduling*. Production Inventory Management Vol.8, 1-10.
19. S. Panwalker and W. Iskander (1977) *A Survey of Scheduling Rules*. Operations Research, Vol.25, 45-61.
20. J.F. Muth and G.L. Thompson (1963) *Industrial Scheduling*. Prentice-Hall, New Jersey.
21. R.W. Conway *et al* (1967) *Theory of Scheduling*. Addison-Wesley, Massachusetts.
22. S. Eilon and J.R. King (1967) *Industrial Scheduling Abstracts 1950-1966*. Oliver & Boyd, London.
23. S.E. Elmaghraby (1973) *Symposium on the Theory of Scheduling*. Springer, Berlin.
24. K.R. Baker (1974) *Introduction to Sequencing and Scheduling*. Wiley, New York.
25. J.K. Lenstra (1976) *Sequencing of Enumerative Methods*. Mathematisch Centrum, Amsterdam.
26. A.H. Rinnoy (1976) *Machine Scheduling Problems: Classification, Complexity and Computations*. Nijhoff, The Netherlands.
27. E.G. Coffman Jr. (1976) *Computer and Job Shop Scheduling Theory*. Wiley, New York.
28. S. French (1982) *Sequencing and Scheduling*. Wiley, New York.
29. H.M. Wagner (1959) *An Integer Programming Model for Machine Scheduling*. Naval Research Logistics Quarterly Vol.6, 131-140.
30. J. Carlier and E. Pinson (1989) *An Algorithm for Solving the Job Shop Problem*. Management Science Vol.35(2), 164-176.
31. E. Ignall and L.E. Sebrage (1965) *Application of the Branch and Bound Technique to Some Flow-Shop Problems*. Operations Research Vol.13, 400-412.
32. J.P. Stinson, E.W. Davis and B.M. Khumawala (1978) *Multiple Resource-Constrained Scheduling Using Branch and Bound*. AIIE Transactions Vol.10, 252-259.
33. M. Held and R.M. Karp (1962) *A Dynamic Programming Approach to Sequencing Problems*. Journal SIAM Vol.10, 196-210.

REFERENCES

1. T.B. Green, W.B. Newsom and R.Jones (1977) *Survey of the application of quantitative techniques to production/operations management in large corporations*. Academic Management Journal Vol.20, 669-676.
2. G. Thomas and J. Da Costa (1979) *A Sample Survey of Corporate Operations Research*. Interfaces Vol.9, 102-111.
3. T.E. Vollmar, W.L. Berry and D.C. Whybark (1988) *Manufacturing Planning and Control Systems*. Irwin, New York.
4. D.W. Fogarty and T.R. Hoffman (1983) *Production and Inventory Management*. South-West Publishing, Cincinnati.
5. R. Peterson and E.A. Silver (1979) *Decision Systems for Inventory Management and Production Planning*. Wiley, New York.
6. J.D. Little (1970) *Models and Managers: The Concept of Decision Calculus*. Management Science Vol.16, 466-485.
7. R.M. O'Keefe, V. Belton and T. Ball (1986) *Experiences with Using Expert Systems in OR*. Journal of the Operational Research Society Vol.37, 657-668.
8. F. Hayes-Roth and N. Jacobstein (1994) *The State of Knowledge-Based Systems*. Communications of the ACM Vol.37(3), 27-38.
9. H.J. Zimmerman (1992) *Fuzzy Mathematical Programming* (ed. S.C. Shapiro) Encyclopaedia of Artificial Intelligence Vol.1, 521-528, Wiley, New York.
10. M. Steffik (1981) *Planning with constraints (MOLGEN: Part 1)*. Artificial Intelligence Vol.16, 111-140.
11. G.J. Sussman and G.L. Steele Jr. (1980) *CONSTRAINTS - a language of expressing almost-hierarchical descriptions*. Artificial Intelligence Vol.14, 1-40.
12. J. Blasewitz, J. Leistra and A.H. Rinnooy Kan (1983) *Scheduling Subject to Resource Constraints: Classification and Complexity*. Discrete Applied Mathematics Vol.5, 11-24.
13. Y. Zong, T. Yang and J. Ignizio (1993) *An Expert System Using an Exchange Heuristic for the Resource-Constrained Scheduling problem*. Expert Systems with Applications Vol.6, 327-348.
14. M.R. Garey, D.S. Johnson and R. Sethi (1976) *Complexity of Flow Shop and Job Shop Scheduling*. Mathematics and Operations Research. Vol.1.

The fuzzy inference scheme was repetitively applied to determine the relative *fitness* of sequences during the search. In order to preserve the absolute order of tasks in sequence during crossover, the Precedence Preservative Crossover operator was applied.

The results of preliminary tests showed that the GA stabilised after ~20 - 30 iterations (generations) with a cumulative utility value u_0 approximately 5% higher (on average) than with the straightforward constructive backtracking approach. Although these tests are preliminary, the author is of the opinion that the collaboration between a fuzzy inference scheme and a genetic search algorithm provides a powerful tool for constraint based planning.

10.4.2. Expansion of the fuzzy knowledge base

Four soft constraints were considered in this investigation [Δt , Δw , Δg , ΔS]. Although Prototype B produced positive results in practice, the functionality of the system may be improved by expanding the size of the fuzzy knowledge base i.e. by modelling a larger sub-set c ($c \in C$) of planning constraints as fuzzy constraints. In particular, resource availability constraints may be modelled on the basis of *fuzzy time windows*. This approach is described by Dubois [94].

10.4. Pertinent areas for further research and development

10.4.1. Improvement of the search method

This investigation was concerned with the development of an inference scheme that *exploits* the constraint knowledge in order to effectively guide the search towards a solution. The inference scheme was supported by a state-space search method employing constructive backtracking.

State-space search is based on the notion of *hill-climbing* - the search thus is local in scope. The concept is illustrated hypothetically in Fig.10.1. Starting the search in the neighbourhood of the lower peak will cause the higher peak to be missed. The search is directed toward the hill with the steepest gradient at its base region. Even where backtracking occurs, the search is re-directed in the neighbourhood of the current point.




Fig.10.1. Hypothetical schematic of a multi-peaked problem space

The quality of prototype-generated solutions may be improved by increasing the *explorative capability* of the applied search method. In this respect *genetic algorithms* combine both exploration and exploitation at the same time in an optimal way. The author experimented with a genetic algorithm (GA), in collaboration with a knowledge consultant from the firm Knowledge Based Engineering [33].

required significantly fewer rules. Moreover, the rules were closer to the way knowledge is expressed by the domain experts. The fuzzy rule-base was essentially static while the binary rule-base showed a clear tendency towards exponential expansion as new rules were added to handle new and unexpected situations.

10.3.2. Improved system maintainability

The knowledge representation scheme was transparent and intuitively clear to the human planner and the domain experts. This meant that logical or structural errors in the model could be fixed within a minimum amount of time *and* the model could be modified with fewer induced errors. Moreover, the model could be validated with greater precision and for a wider variety of input cases - *this greatly increased confidence in the model.*

10.3.3. The ability to model conflicting knowledge

Since the external environment invariably consists of multiple domain experts, the planning knowledge is often conflicting. The fuzzy inference scheme was capable of representing and reasoning with conflicting knowledge. This contributed to the improved consistency and reliability of the prototype and to the reduction in model complexity.

10.3.4. Improved handling of uncertainty

Uncertainty proved to be the major problem in the formulation of the planning decision process. Conventional knowledge-based systems generally deal with uncertainty through a combination of predicate and probability-based methods. There are setbacks to the use of these methods in real-world planning domains, primarily because they rely on the assignment of certainty measures or values *outside* the decision model itself. In contrast, the fuzzy decision model represents uncertainty *intrinsically*. Consequently, it provides a superior method for modelling uncertainty.

[Relative utility values for varying bias on the Δt soft decision variable: all other soft decision variable values kept equal ($\epsilon_w = 0$, $\Delta y = 75$, $\Delta g = 0$)]

Δt (days)	$\delta = 1.00$	$\delta = 1.25$	$\delta = 1.50$	$\delta = 1.75$	$\delta = 2.00$
1	1.088	0.966	0.844	0.722	0.600
2	0.750	0.688	0.625	0.563	0.500
3	0.517	0.480	0.442	0.405	0.368
4	0.353	0.334	0.316	0.297	0.279
5	0.188	0.188	0.188	0.188	0.188
6	0.180	0.180	0.180	0.180	0.180
7	0.160	0.160	0.160	0.160	0.160
8	0.140	0.140	0.140	0.140	0.140
9	0.120	0.120	0.120	0.120	0.120
10	0.100	0.100	0.100	0.100	0.100
11	0.100	0.100	0.100	0.100	0.100
12	0.100	0.100	0.100	0.100	0.100
13	0.100	0.100	0.100	0.100	0.100
14	0.100	0.100	0.100	0.100	0.100
15	0.100	0.100	0.100	0.100	0.100
16	0.100	0.100	0.100	0.100	0.100
17	0.100	0.100	0.100	0.100	0.100
18	0.100	0.100	0.100	0.100	0.100
19	0.100	0.100	0.100	0.100	0.100
20	0.100	0.100	0.100	0.100	0.100

**APPENDIX 9: EVALUATION DATA (SENSITIVITY ANALYSIS WITH RESPECT TO BIAS ON THE
SOFT DECISION VARIABLES Δw AND Δc)**

[Relative utility values for varying bias on the Δw soft decision variable; all other soft decision variable values kept equal ($\Delta r = 0$, $\Delta s = 75$, $\Delta g = 0$)]

Δw_i (mm)	$\delta = 1.00$	$\delta = 1.25$	$\delta = 1.50$	$\delta = 1.75$	$\delta = 2.00$
100	0.500	0.500	0.500	0.500	0.500
95	0.510	0.574	0.638	0.765	0.702
90	0.520	0.585	0.650	0.780	0.715
85	0.530	0.596	0.663	0.795	0.729
80	0.540	0.608	0.675	0.810	0.743
75	0.550	0.619	0.688	0.825	0.756
70	0.560	0.630	0.700	0.840	0.770
65	0.570	0.641	0.713	0.855	0.784
60	0.580	0.653	0.725	0.870	0.798
55	0.588	0.662	0.735	0.882	0.809
50	0.596	0.664	0.740	0.895	0.818
45	0.598	0.741	0.884	1.169	1.026
40	0.607	0.752	0.897	1.187	1.043
35	0.616	0.763	0.911	1.205	1.058
30	0.625	0.775	0.924	1.223	1.074
25	0.634	0.786	0.938	1.241	1.089
20	0.643	0.797	0.951	1.259	1.105
15	0.652	0.808	0.964	1.277	1.121
10	0.661	0.819	0.978	1.295	1.136
5	0.671	0.833	0.995	1.319	1.157
0	0.688	0.860	1.032	1.376	1.204
-5	0.671	0.827	0.984	1.296	1.140
-10	0.654	0.795	0.936	1.218	1.077
-15	0.636	0.761	0.887	1.138	1.012
-20	0.618	0.728	0.838	1.057	0.947
-25	0.600	0.694	0.788	0.977	0.883
-30	0.582	0.661	0.739	0.896	0.818
-35	0.564	0.627	0.690	0.816	0.753
-40	0.546	0.594	0.641	0.648	0.618
-45	0.528	0.558	0.588	0.648	0.618
-50	0.500	0.500	0.500	0.500	0.500

**APPENDIX 8: EVALUATION DATA (PLANNER-PROTOTYPE COHERENCE IN THE
PRIORITISATION AND SELECTION OF INSTANTIATIONS)**

Test No.	λ_i selection (from the same set of 15 feasible instantiations)											
	Human Planner				Prototype A				Prototype B			
	Δt_i	Δw_{ij}	Δv_i	Δg_{ij}	Δt_i	Δw_{ij}	Δv_i	Δg_{ij}	Δt_i	Δw_{ij}	Δv_i	Δg_{ij}
1	0	0	11	0.096	2	0	37	0.180	0	0	11	0.096
2	1	0	94	0.047	1	0	92	0.306	1	0	94	0.047
3	1	-50	37	0.00	0	50	40	0.252	0	100	75	0.090
4	0	0	6	0.078	0	0	6	0.078	0	0	6	0.078
5	3	0	19	0.023	3	0	14	0.003	3	6	19	0.023
6	1	100	9	0.125	1	100	37	0.331	1	100	9	0.125
7	1	50	41	0.086	1	50	71	0.254	1	50	41	0.086
8	5	0	45	0.130	5	0	90	0.311	5	0	45	0.130
9	0	50	17	0.021	0	50	17	0.021	0	50	17	0.021
10	2	50	42	0.259	2	50	84	0.388	2	50	42	0.259
11	0	-50	97	0.049	4	-50	34	0.125	0	-50	97	0.049
12	0	100	69	0.188	0	100	57	0.359	0	100	69	0.188
13	1	100	27	0.154	1	100	55	0.262	1	100	27	0.154
14	1	0	39	0.143	1	0	39	0.143	1	0	39	0.143
15	1	0	54	0.044	1	0	67	0.162	1	0	54	0.044
16	0	0	17	0.002	0	0	17	0.002	0	0	17	0.002
17	3	50	88	0.005	3	50	96	0.151	3	50	88	0.005
18	1	100	5	0.104	1	100	49	0.139	1	100	5	0.104
19	2	-50	54	0.173	2	-50	54	0.173	2	-50	54	0.173
20	1	-50	28	0.166	1	-50	2	0.261	1	-50	28	0.166
10	2	50	58	0.339	2	50	32	0.042	2	50	58	0.339
11	3	0	10	0.009	0	-50	75	0.253	1	100	25	0.188
12	0	100	100	0.257	0	100	63	0.326	0	100	100	0.257
13	1	0	96	0.073	1	100	88	0.267	1	0	96	0.073
14	3	0	32	0.062	5	0	85	0.292	3	0	32	0.062
15	1	100	47	0.223	1	100	47	0.223	1	100	47	0.223
16	0	0	28	0.104	0	0	28	0.104	0	0	28	0.104
17	3	50	49	0.281	3	50	76	0.299	3	50	49	0.281
18	1	100	47	0.345	1	100	78	0.040	1	100	47	0.345
19	1	-50	50	0.390	7	-50	46	0.243	1	-50	50	0.390
20	1	-50	32	0.326	1	-50	78	0.148	1	-50	32	0.326

- Shading for Prototype A highlights *coherence*
- Shading for Prototype B highlights *incoherence*

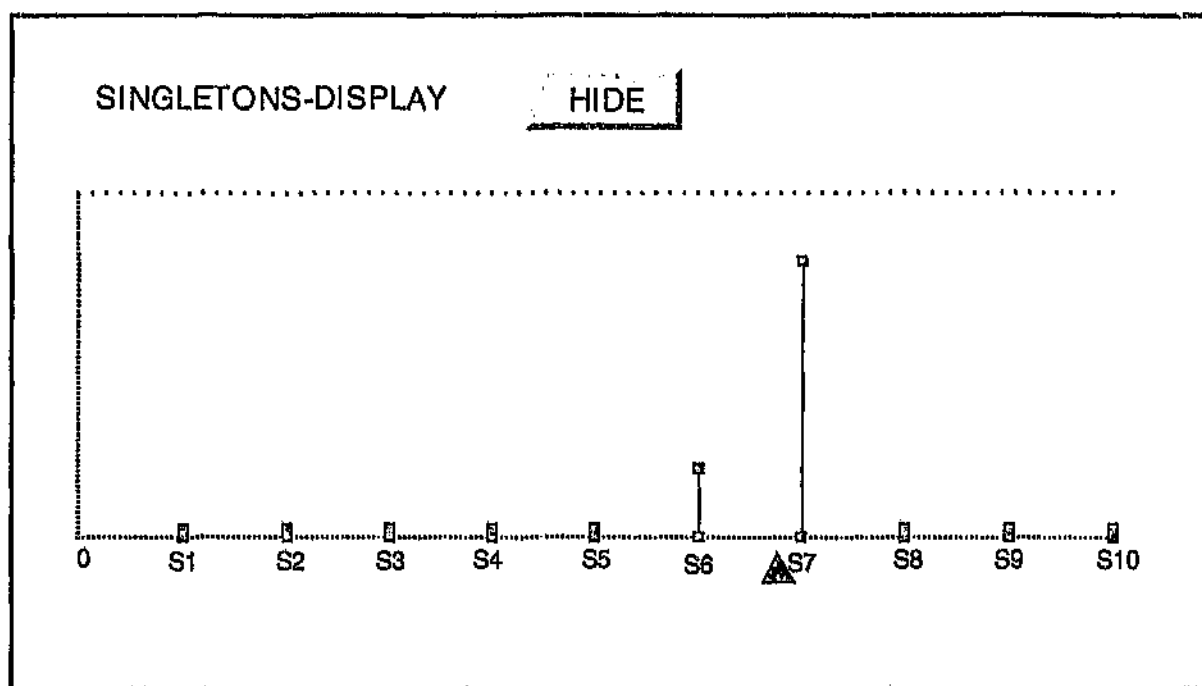
APPENDIX 7: EVALUATION DATA (PROTOTYPE CONFORMANCE TO SELECTED PLANNING OBJECTIVES)

- SELECTED OBJECTIVE: [O1] the average planned ratio of slab cut off or not regarded as prime material
 [O2] the average number of tasks between casting interruptions
 [O3] the delivery-date achievement ratio

Reference - manual sequence generation (average from historical data)

Test	PROTOTYPE A						PROTOTYPE B					
	O1		O2		O3		O1		O2		O3	
	Relative	Actual	Relative	Actual	Relative	Actual	Relative	Actual	Relative	Actual	Relative	Actual
1	1.000	0.19	1.000	6.50	1.000	0.57	1.000	0.19	1.000	6.50	1.000	0.57
2	0.970	0.18	0.958	6.23	0.990	0.56	0.740	0.14	1.120	7.28	1.370	0.78
3	1.020	0.19	0.975	6.34	1.016	0.58	0.680	0.13	1.200	7.80	1.450	0.83
4	0.986	0.19	0.942	6.12	1.035	0.59	0.701	0.13	1.150	7.48	1.470	0.84
5	0.967	0.18	0.980	6.37	1.088	0.62	0.725	0.14	1.270	8.26	1.480	0.84
6	0.940	0.18	1.030	6.70	1.038	0.59	0.693	0.13	1.200	7.80	1.320	0.75
7	0.901	0.17	1.043	6.78	1.035	0.59	0.745	0.14	1.250	8.13	1.390	0.79
8	0.877	0.17	1.026	6.67	1.065	0.61	0.716	0.14	1.026	6.67	1.270	0.72
9	0.863	0.16	1.047	6.81	1.074	0.61	0.725	0.14	1.070	6.96	1.420	0.81
10	0.822	0.16	1.063	6.91	1.097	0.63	0.654	0.12	1.170	7.61	1.320	0.75
11	0.831	0.16	1.045	6.79	1.090	0.62	0.696	0.13	1.230	8.00	1.270	0.72
12	0.817	0.16	1.070	6.96	1.090	0.62	0.724	0.14	1.190	7.74	1.280	0.73
13	0.829	0.16	1.079	7.01	1.070	0.61	0.752	0.14	1.160	7.54	1.370	0.78
14	0.839	0.16	0.989	6.43	1.080	0.62	0.720	0.14	1.180	7.67	1.210	0.69
15	0.828	0.16	1.082	7.03	1.107	0.63	0.711	0.14	1.210	7.87	1.450	0.83
16	0.821	0.16	1.082	7.03	1.120	0.64	0.700	0.13	1.200	7.80	1.460	0.83
17	0.826	0.16	1.110	7.22	1.100	0.63	0.697	0.13	1.130	7.35	1.400	0.80
18	0.823	0.16	1.090	7.09	1.118	0.64	0.723	0.14	1.160	7.54	1.435	0.82
19	0.830	0.16	1.115	7.25	1.090	0.62	0.685	0.13	1.260	8.19	1.450	0.83
20	0.835	0.16	1.090	7.09	1.125	0.64	0.692	0.13	1.210	7.87	1.420	0.81

APPENDIX 6: SINGLETON GEOMETRY OUTPUT SPACE



HIDE

UNDESIRABLE	EXTREME	WIDE	INSIGNIFICANT	7
UNDESIRABLE	EXTREME	WIDE	NON-URGENT	8
UNDESIRABLE	EXTREME	WIDE	SEMI-URGENT	5
UNDESIRABLE	EXTREME	WIDE	URGENT	5
UNDESIRABLE	EXTREME	VERY-NARROW	INSIGNIFICANT	8
UNDESIRABLE	EXTREME	VERY-NARROW	NON-URGENT	7
UNDESIRABLE	EXTREME	VERY-NARROW	SEMI-URGENT	8
UNDESIRABLE	EXTREME	VERY-NARROW	URGENT	5
UNDESIRABLE	EXTREME	NARROW	INSIGNIFICANT	7
UNDESIRABLE	EXTREME	NARROW	NON-URGENT	6
UNDESIRABLE	EXTREME	NARROW	SEMI-URGENT	5
UNDESIRABLE	EXTREME	NARROW	URGENT	4
UNDESIRABLE	HIGH	WIDE	INSIGNIFICANT	4
UNDESIRABLE	HIGH	WIDE	NON-URGENT	3
UNDESIRABLE	HIGH	WIDE	SEMI-URGENT	2
UNDESIRABLE	HIGH	WIDE	URGENT	2
UNDESIRABLE	HIGH	VERY-NARROW	INSIGNIFICANT	8
UNDESIRABLE	HIGH	VERY-NARROW	NON-URGENT	8
UNDESIRABLE	HIGH	VERY-NARROW	SEMI-URGENT	4
UNDESIRABLE	HIGH	VERY-NARROW	URGENT	3
UNDESIRABLE	HIGH	NARROW	INSIGNIFICANT	8
UNDESIRABLE	HIGH	NARROW	NON-URGENT	4
UNDESIRABLE	HIGH	NARROW	SEMI-URGENT	3
UNDESIRABLE	HIGH	NARROW	URGENT	3
UNDESIRABLE	LOW	WIDE	INSIGNIFICANT	1
UNDESIRABLE	LOW	WIDE	NON-URGENT	1
UNDESIRABLE	LOW	WIDE	SEMI-URGENT	1
UNDESIRABLE	LOW	WIDE	URGENT	1
UNDESIRABLE	LOW	VERY-NARROW	INSIGNIFICANT	3
UNDESIRABLE	LOW	VERY-NARROW	NON-URGENT	2
UNDESIRABLE	LOW	VERY-NARROW	SEMI-URGENT	1
UNDESIRABLE	LOW	VERY-NARROW	URGENT	1
UNDESIRABLE	LOW	NARROW	INSIGNIFICANT	2
UNDESIRABLE	LOW	NARROW	NON-URGENT	1
UNDESIRABLE	LOW	NARROW	SEMI-URGENT	1
UNDESIRABLE	LOW	NARROW	URGENT	1
UNDESIRABLE	NEGLECTIBLE	WIDE	INSIGNIFICANT	1
UNDESIRABLE	NEGLECTIBLE	WIDE	NON-URGENT	1
UNDESIRABLE	NEGLECTIBLE	WIDE	SEMI-URGENT	1
UNDESIRABLE	NEGLECTIBLE	WIDE	URGENT	1
UNDESIRABLE	NEGLECTIBLE	VERY-NARROW	INSIGNIFICANT	1
UNDESIRABLE	NEGLECTIBLE	VERY-NARROW	NON-URGENT	1
UNDESIRABLE	NEGLECTIBLE	VERY-NARROW	SEMI-URGENT	1
UNDESIRABLE	NEGLECTIBLE	VERY-NARROW	URGENT	1
UNDESIRABLE	NEGLECTIBLE	NARROW	INSIGNIFICANT	1
UNDESIRABLE	NEGLECTIBLE	NARROW	NON-URGENT	1
UNDESIRABLE	NEGLECTIBLE	NARROW	SEMI-URGENT	1
UNDESIRABLE	NEGLECTIBLE	NARROW	URGENT	1


HIDE

LARGE	EXTREME	WIDE	INSIGNIFICANT	8
LARGE	EXTREME	WIDE	NON-URGENT	7
LARGE	EXTREME	WIDE	SEMI-URGENT	6
LARGE	EXTREME	WIDE	URGENT	8
LARGE	EXTREME	VERY-NARROW	INSIGNIFICANT	9
LARGE	EXTREME	VERY-NARROW	NON-URGENT	8
LARGE	EXTREME	VERY-NARROW	SEMI-URGENT	7
LARGE	EXTREME	VERY-NARROW	URGENT	6
LARGE	EXTREME	NARROW	INSIGNIFICANT	8
LARGE	EXTREME	NARROW	NON-URGENT	7
LARGE	EXTREME	NARROW	SEMI-URGENT	6
LARGE	EXTREME	NARROW	URGENT	5
LARGE	HIGH	WIDE	INSIGNIFICANT	5
LARGE	HIGH	WIDE	NON-URGENT	4
LARGE	HIGH	WIDE	SEMI-URGENT	3
LARGE	HIGH	WIDE	URGENT	3
LARGE	HIGH	VERY-NARROW	INSIGNIFICANT	7
LARGE	HIGH	VERY-NARROW	NON-URGENT	6
LARGE	HIGH	VERY-NARROW	SEMI-URGENT	6
LARGE	HIGH	VERY-NARROW	URGENT	4
LARGE	HIGH	NARROW	INSIGNIFICANT	8
LARGE	HIGH	NARROW	NON-URGENT	8
LARGE	HIGH	NARROW	SEMI-URGENT	4
LARGE	HIGH	NARROW	URGENT	4
LARGE	LOW	WIDE	INSIGNIFICANT	2
LARGE	LOW	WIDE	NON-URGENT	1
LARGE	LOW	WIDE	SEMI-URGENT	1
LARGE	LOW	WIDE	URGENT	1
LARGE	LOW	VERY-NARROW	INSIGNIFICANT	4
LARGE	LOW	VERY-NARROW	NON-URGENT	3
LARGE	LOW	VERY-NARROW	SEMI-URGENT	2
LARGE	LOW	VERY-NARROW	URGENT	1
LARGE	LOW	NARROW	INSIGNIFICANT	8
LARGE	LOW	NARROW	NON-URGENT	2
LARGE	LOW	NARROW	SEMI-URGENT	1
LARGE	LOW	NARROW	URGENT	1
LARGE	NEGLECTIBLE	WIDE	INSIGNIFICANT	1
LARGE	NEGLECTIBLE	WIDE	NON-URGENT	1
LARGE	NEGLECTIBLE	WIDE	SEMI-URGENT	1
LARGE	NEGLECTIBLE	WIDE	URGENT	1
LARGE	NEGLECTIBLE	VERY-NARROW	INSIGNIFICANT	2
LARGE	NEGLECTIBLE	VERY-NARROW	NON-URGENT	1
LARGE	NEGLECTIBLE	VERY-NARROW	SEMI-URGENT	1
LARGE	NEGLECTIBLE	VERY-NARROW	URGENT	1
LARGE	NEGLECTIBLE	NARROW	INSIGNIFICANT	1
LARGE	NEGLECTIBLE	NARROW	NON-URGENT	1
LARGE	NEGLECTIBLE	NARROW	SEMI-URGENT	1
LARGE	NEGLECTIBLE	NARROW	URGENT	1


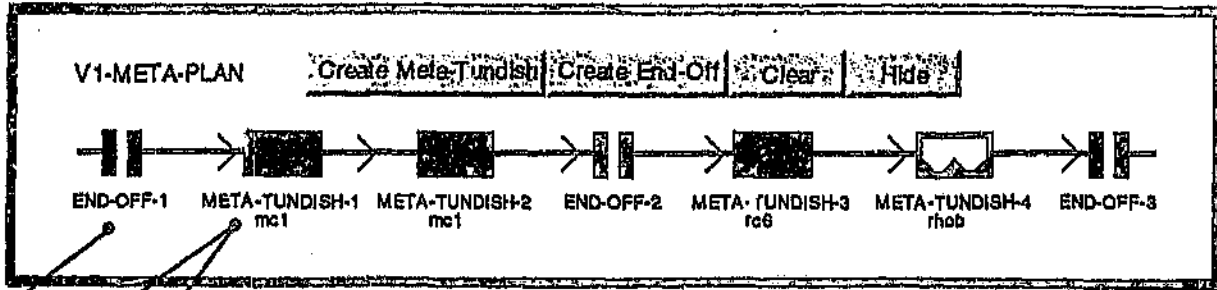
APPENDIX 5: FUZZY PROPOSITIONAL TABLE (INTERACTIVE SCREEN)

				HIDE
SMALL	EXTREME	WIDE	INSIGNIFICANT	8
SMALL	EXTREME	WIDE	NON-URGENT	8
SMALL	EXTREME	WIDE	SEMI-URGENT	7
SMALL	EXTREME	WIDE	URGENT	7
SMALL	EXTREME	VERY-NARROW	INSIGNIFICANT	10
SMALL	EXTREME	VERY-NARROW	NON-URGENT	8
SMALL	EXTREME	VERY-NARROW	SEMI-URGENT	6
SMALL	EXTREME	VERY-NARROW	URGENT	7
SMALL	EXTREME	NARROW	INSIGNIFICANT	8
SMALL	EXTREME	NARROW	NON-URGENT	8
SMALL	EXTREME	NARROW	SEMI-URGENT	7
SMALL	EXTREME	NARROW	URGENT	6
SMALL	HIGH	WIDE	INSIGNIFICANT	6
SMALL	HIGH	WIDE	NON-URGENT	6
SMALL	HIGH	WIDE	SEMI-URGENT	4
SMALL	HIGH	WIDE	URGENT	4
SMALL	HIGH	VERY-NARROW	INSIGNIFICANT	8
SMALL	HIGH	VERY-NARROW	NON-URGENT	7
SMALL	HIGH	VERY-NARROW	SEMI-URGENT	6
SMALL	HIGH	VERY-NARROW	URGENT	5
SMALL	HIGH	NARROW	INSIGNIFICANT	7
SMALL	HIGH	NARROW	NON-URGENT	8
SMALL	HIGH	NARROW	SEMI-URGENT	8
SMALL	HIGH	NARROW	URGENT	5
SMALL	LOW	WIDE	INSIGNIFICANT	3
SMALL	LOW	WIDE	NON-URGENT	2
SMALL	LOW	WIDE	SEMI-URGENT	1
SMALL	LOW	WIDE	URGENT	1
SMALL	LOW	VERY-NARROW	INSIGNIFICANT	6
SMALL	LOW	VERY-NARROW	NON-URGENT	4
SMALL	LOW	VERY-NARROW	SEMI-URGENT	3
SMALL	LOW	VERY-NARROW	URGENT	2
SMALL	LOW	NARROW	INSIGNIFICANT	4
SMALL	LOW	NARROW	NON-URGENT	3
SMALL	LOW	NARROW	SEMI-URGENT	2
SMALL	LOW	NARROW	URGENT	2
SMALL	NEGLECTIBLE	WIDE	INSIGNIFICANT	2
SMALL	NEGLECTIBLE	WIDE	NON-URGENT	1
SMALL	NEGLECTIBLE	WIDE	SEMI-URGENT	1
SMALL	NEGLECTIBLE	WIDE	URGENT	1
SMALL	NEGLECTIBLE	VERY-NARROW	INSIGNIFICANT	3
SMALL	NEGLECTIBLE	VERY-NARROW	NON-URGENT	2
SMALL	NEGLECTIBLE	VERY-NARROW	SEMI-URGENT	1
SMALL	NEGLECTIBLE	VERY-NARROW	URGENT	1
SMALL	NEGLECTIBLE	NARROW	INSIGNIFICANT	2
SMALL	NEGLECTIBLE	NARROW	NON-URGENT	1
SMALL	NEGLECTIBLE	NARROW	SEMI-URGENT	1
SMALL	NEGLECTIBLE	NARROW	URGENT	1

APPENDIX 4: THE FUZZY CONSTRAINT SPACES (INTERACTIVE SCREEN)



APPENDIX 3: THE META-PLAN GENERATION INTERFACE



TUNDISH CONFIGURATION

Tundish Name

Minimum Length

Maximum Length

Start Grade Name

Minimum Quantity

Maximum Quantity

Icon Colour

FOLLOW ON GRADES

CANCEL DELETE ACCEPT

CREATE META-TUNDISH

Partial-sequence-type (PST) objects

Casting interruption objects

Sequence-constraint-specification (SCS) objects

FOLLOW ON GRADES FOR MC1

Follow-on Grade Name

ACCEPT APPEND DELETE HIDE

AC2 AC3 MC2 MC3

APPENDIX 2: THE VISUALISATION INTERFACE

VISUALISATION TABLE		OK 2153 HEATS																	
GRADE	970	975	1000	1025	1050	1075	1100	1125	1150	1175	1200	1225	1250	1300	1400	1500	1600	1700	1800
AC1-32	1						2				1								
AC1-33			2				10				15			40	25	15	10	10	
AC1-34			1								1	1	1						
AC1-35																			
AC2-31																			
AC2-32			4	6	10	10						4	4	4	4	4	4	4	
AC2-33																			
AC2-34																			
AC2-35																			
AC2-36											1								
AC2-37																			
AC2-38																			
AC2-39																			
AC2-40																			
AC2-41																			
AC2-42																			
AC2-43																			
AC2-44																			
AC2-45																			
AC2-46																			
AC2-47																			
AC2-48																			
AC2-49																			
AC2-50																			
AC2-51																			
AC2-52																			
AC2-53																			
AC2-54																			
AC2-55																			
AC2-56																			
AC2-57																			
AC2-58																			
AC2-59																			
AC2-60																			
AC2-61																			
AC2-62																			
AC2-63																			
AC2-64																			
AC2-65																			
AC2-66																			
AC2-67																			
AC2-68																			
AC2-69																			
AC2-70																			
AC2-71																			
AC2-72																			
AC2-73																			
AC2-74																			
AC2-75																			
AC2-76																			
AC2-77																			
AC2-78																			
AC2-79																			
AC2-80																			
AC2-81																			

Each cell indicates the number of tasks for each product type p_i .

Product types are categorised according to steel grade (down) and slab width (across)

Author: Besteiro M.

Name of thesis: Customer oriented production planning at an integrated steel works.

PUBLISHER:

University of the Witwatersrand, Johannesburg

©2015

LEGALNOTICES:

Copyright Notice: All materials on the University of the Witwatersrand, Johannesburg Library website are protected by South African copyright law and may not be distributed, transmitted, displayed or otherwise published in any format, without the prior written permission of the copyright owner.

Disclaimer and Terms of Use: Provided that you maintain all copyright and other notices contained therein, you may download material (one machine readable copy and one print copy per page) for your personal and/or educational non-commercial use only.

The University of the Witwatersrand, Johannesburg, is not responsible for any errors or omissions and excludes any and all liability for any errors in or omissions from the information on the Library website.