

University of Strathclyde
Department of Electronic and Electrical Engineering

A PhD Thesis Entitled

**An Expert System Approach to On-line Fault
Diagnosis in Power System Networks**

by

Graeme M. Burt

September 1992

Contents

Contents	ii
List of Figures	viii
Abstract	xi
Acknowledgements	xiii
1 Introduction	1
1.1 Justification for and Introduction to Research	2
1.2 Thesis Outline	8
1.3 Associated Publications	10
2 Expert Systems	12
2.1 Introduction to Artificial Intelligence	13
2.2 The Concept and Evolution of Expert Systems	14
2.3 Realisation of Expert Systems	18
2.3.1 Structure	18
2.3.2 Knowledge Acquisition	20
2.3.3 Means of Implementation	21
2.3.4 Methods of Realisation	22
2.3.5 Techniques	25
2.3.6 Strategies	29

2.4	Dealing with Uncertainty	32
2.4.1	Sources	32
2.4.2	Techniques	33
2.4.3	Some General Points	49
2.5	Learning Facilities	50
2.6	Explanation Facilities	52
2.7	Advantages of Expert Systems	53
2.8	Applications in Power Engineering	55
3	Power Systems & Fault Diagnosis	57
3.1	Structure of Power Systems	58
3.1.1	Network	58
3.1.2	Protection Systems	59
3.1.3	Control and Monitoring Systems	62
3.2	Some Developments in Power Systems	66
3.2.1	Network	66
3.2.2	Protection	67
3.2.3	Local Control and Monitoring	68
3.2.4	Remote Control and Monitoring	70
3.3	Fault Diagnosis of Power Systems	72
3.3.1	Events	72
3.3.2	Reasons for Fault Diagnosis	72
3.3.3	Basis for Diagnosis	73
3.3.4	Reasons for a Fault Diagnosis System	74
3.3.5	An Expert System	77

3.4	Review of Other Expert Systems for Network Fault Diagnosis . . .	81
3.4.1	Automatic System Trouble Analysis	81
3.4.2	An Expert System Presented By Fukui and Kawakami . . .	83
3.4.3	An Expert System Described by Wollenberg and Sakaguchi	85
3.4.4	Distributed Diagnosis	85
3.4.5	Fault Analysis Using Simulation and Testing	87
3.4.6	An Expert System Developed by Girgis and Johns	89
3.4.7	Diagnosis of System Faults	91
3.4.8	An Expert System Developed by Jeyasurya et al	92
3.4.9	An Expert System Described by Hotta et al	94
3.4.10	GESTAL	95
3.4.11	Computer Assisted Intelligent Storm Outage Evaluator . .	96
3.4.12	An Expert System Presented by Okamoto et al	97
3.4.13	An Expert System Developed by Partanen et al	98
3.4.14	An Expert System Developed by Hsu et al	99
3.4.15	Customer Restoration And Fault Testing	100
3.4.16	AHFA	101
3.4.17	An Expert System Described by Claus et al	102
3.4.18	SEPT	103
3.4.19	An Expert System Developed by Amantegui et al	104
3.4.20	Other Systems Related to Fault Diagnosis of Power System Networks	106
3.4.21	General Comments	107
4	Diagnostic Knowledge	109
4.1	Preliminary Discussion of Implementation Issues	110

4.2	Structure of Diagnosis Results	113
4.2.1	Observations	113
4.2.2	Initial Structure	113
4.2.3	Three-Tier Hierarchical Structure	116
4.3	Structure of the Means of Diagnosis	120
4.3.1	Observations	120
4.3.2	Hierarchical Structure	120
4.3.3	Performance Monitor	123
4.3.4	Management	123
4.4	Philosophies of Diagnosis	124
4.4.1	Historian Philosophy	124
4.4.2	Wait-and-See Philosophy	124
4.4.3	Hypothesiser Philosophy	125
4.5	Wait-and-See	127
4.5.1	Gut Reaction	127
4.5.2	Central Diagnosis	129
4.6	Hypothesiser	132
4.6.1	Gut Reaction	132
4.6.2	Evolution of Hypothesising Strategies	134
4.6.3	Central Diagnosis	139
4.7	Dealing with Uncertainty	149
4.7.1	Importance in Diagnosis	149
4.7.2	Sources of Uncertainty	149
4.7.3	Problem Factors	150
4.7.4	The Developed Measure - Relative Likelihood	151
4.8	Ideas in the Deeper Diagnosis	156

5	Network Knowledge	158
5.1	Significance of Network Knowledge	159
5.2	Dynamic Data Requirements	160
5.3	Static Data Requirements and Structure	162
5.3.1	Scope	162
5.3.2	Structure of Power System Data	162
5.3.3	Plant Data	164
5.3.4	Protection and Switchgear Data	174
5.3.5	Relative Likelihood Data	180
5.3.6	Telemetry Data	181
5.4	Qualitative Simulation	183
5.4.1	Function	183
5.4.2	Reason for Being Qualitative	183
5.4.3	Representation of Activity	184
5.4.4	The Qualitative Simulation Knowledge	185
5.4.5	Explanation/Understanding Facilities	192
5.5	Loadflow	194
5.5.1	Power Infeed Search	194
5.5.2	The Original Problem	194
5.5.3	Possible Solutions	196
5.5.4	Solution Adopted	199
5.5.5	Comments	201
5.6	Modelling	203
5.6.1	Faults on Plant	203
5.6.2	Plant Internal Links	204
5.6.3	Protection Detection Logic	211
5.6.4	Protection and Switchgear Activity	221

5.7	Chapter Summary	227
6	Use of the Fault Diagnosis	228
6.1	Case Studies	229
6.1.1	Test Power System and Cases Used	229
6.1.2	Case 1	231
6.1.3	Case 2	237
6.1.4	Case 3	243
6.2	Evaluation	248
6.2.1	Some Factors	248
6.2.2	Knowledge Nature	258
6.2.3	Comparison with Other Systems	263
6.3	Anticipated Use	269
6.3.1	Short Term	269
6.3.2	Long Term	271
7	Conclusions and Future Work	275
7.1	Conclusions	276
7.2	Future Work	280
A	Introduction to Prolog	283
A.1	Convention Adopted	284
A.2	Brief Description of the Prolog Language	284
	References	289

List of Figures

- 2.1 The basic structure of expert systems 18
- 2.2 An example of a blackboard system 19
- 2.3 A semantic network 24
- 2.4 A schematic of object-oriented programming 25
- 2.5 Diagram showing backward chaining 27
- 2.6 Diagram showing forward chaining 28
- 2.7 Diagram showing backtracking 29
- 2.8 The definition of the odds multiplier λ 47

- 3.1 A simplistic illustration of the principles of (a) unit protection and
 (b) non-unit protection 61

- 4.1 The diagrammatic representation of the example scenario 115
- 4.2 Illustration of the effect of a planar disturbance area 118
- 4.3 Overlapping of protections 119
- 4.4 The search to determine the electrical status of breaker terminals 128
- 4.5 A 1 substation radius disturbance area 133
- 4.6 The disturbance area of a three zone distance protection scheme
 on a simplistic series of lines 134
- 4.7 A schematic of the first hypothesising strategy 136
- 4.8 A schematic of the second hypothesising strategy 137

4.9	A schematic of the third hypothesising strategy	139
4.10	Schematic of the relative likelihood interdependence	153
5.1	A portion of network showing a supply point busbar (at <i>ss1</i>) and a loaded busbar (at <i>ss2</i>)	165
5.2	The general template for a type of plant item	165
5.3	An example of some trip and control wiring	178
5.4	Illustration of the failure of the depth-first search	195
5.5	Schematics of loadflow nodes	200
5.6	Transformer winding arrangements with their zero sequence equiv- alents and knowledge of internal links for neutral effects	208
5.7	Illustration of the two levels of busbar model	210
5.8	Protection scheme using the differential relay type: (a) schematic; (b) knowledge base schematic	212
5.9	Protection scheme using the restricted earth fault relay type: (a) schematic; (b) knowledge base schematic	213
5.10	Illustration of a protection reach defined by cumulative impedances	215
5.11	Protection scheme using the distance relay type: (a) schematic; (b) knowledge base schematic	216
5.12	Protection scheme using the directional overcurrent relay type: (a) schematic; (b) knowledge base schematic	217
5.13	IDMT current-time characteristic	218
5.14	Protection scheme using the non-directional overcurrent relay type: (a) schematic; (b) knowledge base schematic	219
5.15	Protection scheme using the high set overcurrent relay type: (a) schematic; (b) knowledge base schematic	220
5.16	Protection scheme using the neutral overcurrent relay type: (a) schematic; (b) knowledge base schematic	220

5.17	Protection scheme using the instantaneous earth fault relay type:	
	(a) schematic; (b) knowledge base schematic	221
5.18	The abstract schematic of a two-way intertrip scheme	225
6.1	A portion of a transmission network	231
6.2	A portion of hypothetical network	237
6.3	A portion of a transmission network	244

Abstract

The increased computing power which has been installed in power system control centres coupled with the introduction of more advanced Supervisory Control And Data Acquisition (SCADA) systems has increased the volume of data presented to system operators. In certain instances, control centre staff have failed to keep track of the interpretation of system conditions following faults. There is a recognised need to provide on-line aids to the interpretation of SCADA system data, such as fault diagnosis systems.

This thesis reports on research undertaken to produce a practical fault diagnostic expert system which can cope with real-world problems encountered by power system operators such as protection/switchgear maloperations and missing alarms. The Prolog language was used to represent the knowledge in both logic rules and production rules. In view of the scale of the fault diagnosis problem, hierarchical structures were developed for diagnostic conclusions and the means of diagnosis. Two diagnostic strategies have been developed. Firstly, the “wait-and-see” strategy awaits the appearance of black-out areas and uses the intersection of related relays’ protected areas to locate faults. In contrast, the “hypothesiser” strategy compares activity obtained from qualitative simulations of fault hypotheses with received SCADA alarms. Both strategies rank competing conclusions according to a developed “relative likelihood” measure. Significant portability has been achieved through the incorporation of qualitative modelling knowledge for power and protection system equipment which allows a distinction

between the diagnostic knowledge and network knowledge to be maintained. The efficiency of the network tracing used in the expert systems was improved through the co-operative use of a loadflow routine. Both the desire to include maloperations and to contain, to a reasonable level, the number of possible conclusions considered, were found to be contrasting. The evolved hypothesising expert system is presented as a viable solution. Case studies demonstrate the beneficial information provided by the fault diagnosis. They also show the advantages of the ultimate “hypothesiser” over the “wait-and-see” strategy, which include an inherent explanation capability.

The level of knowledge in the fault diagnosis conclusions opens up potential for the expert system’s use as the basis of a restoration advisor.

Acknowledgements

I would like to thank Dr. Jim McDonald, Manager of the Centre for Electrical Power Engineering, for his guidance and help throughout the duration of my research. His time spent discussing my work was and is greatly appreciated and is unreservedly acknowledged by the author. I also thank him for his perseverance in reading and commenting on this thesis.

The help given by ScottishPower in supplying some recorded SCADA alarm listings and network diagrams is acknowledged.

I am grateful to the Science and Engineering Research Council for financing me during the course of this research. The generous help of the Isaac Sclar Scholarship Trust is also gratefully acknowledged.

I would like to thank Professor Lo for allowing me to conduct this research. And also to my colleagues in my office, thanks is due for their friendship over these couple of years. The contributions of David Young in joint project progress discussions is acknowledged as is his help with many computing queries. To my colleagues “up the spiral stairs”, I say thank you for your friendship and contribution to a good cooperative work environment.

Most importantly, I offer my deepest thanks to my family. Without their support and patience during the writing of this thesis, I could not have completed my work. In particular, I thank my mum and dad for their financial sacrifice in

giving me the opportunity to study and conduct this research at the University of Strathclyde.

Lastly, Psalm 54v4.

Chapter 1

Introduction

1.1 Justification for and Introduction to Research

Pressures on utilities to improve the management of their power systems are increasing. With continued economic development and desires for better security, many power systems have become increasingly interconnected. As a result their operation has become a more complex task. In certain areas, especially the United States of America, the development of new plant, particularly generating plant, has not kept pace with growth in demand. This condition is due to legislative pressures having been exerted by environmental limits, financial pressures and public concern over nuclear power. As a result, the maximisation of the utilisation of plant has become of increased concern. Also the increased use of imported electrical power has resulted in large power flows, and diverse system conditions arise which did not before. The pressures on utilities have been reflected on system operators charged with the operation of the network, and will continue to be so.

At the same time, technological advances have been made by which the data available at the control centre (and also in the substation) has vastly increased. Whereas in the past a lack of system data restricted the operators' potential for response to system events, this no longer applies to the same extent, especially in transmission systems. Supervisory Control and Data Acquisition (SCADA) systems have expanded to alarm many relay and switchgear status changes, and also gather current and voltage values from measuring points in the system. Computing and communication system developments have now expanded the volume of available data to levels at which data handling systems and control centre staff are becoming saturated following some disturbances. Therefore, although the potential for improved operation of the system is provided by the additional data,

that same data causes operators to be overloaded. Moreover, the interconnection of power systems has led to utility control centres being supplied with some SCADA data from neighbouring utilities.

The result is that improved operation of power systems is desired and that operators fail to interpret system conditions from the SCADA data presented to them. This thesis suggests that a fault diagnosis system must be provided to help system operators and protection engineers, and presents research into developing such a system. Only when the genuine system conditions have been identified can effective restoration procedures be implemented to minimise the effect of the isolation of the disturbance on the system. Furthermore the recognition of system faults, protection and switchgear maloperations, and telemetry malfunctions will direct investigation teams to where work is required. The time spent testing power system equipment will be minimised. As a result, both the unnecessarily isolated plant and the faulted plant down-time could be reduced.

Expert systems, or knowledge based systems, are an attempt to replicate the reasoning of an expert in solving fairly complex real-world problems in a narrow domain. Expert systems is one of the areas involved in the field of artificial intelligence. This is a description rather than the definition of the term since many attempts have been made to specify what an expert system is. However, the techniques differ from conventional techniques in that the knowledge they contain is declarative rather than procedurally embedded. Consequently, inference and reasoning predominate rather than algorithms. The knowledge contained in the knowledge base of an expert system covers a narrow domain. However, expert system techniques have been applied to a multitude of diverse domains, and some have proved to be of great benefit. Many different means of representing knowledge have been adopted, and these have been realised by using many diverse languages, expert system “shells” and toolkits.

The application of expert system techniques to power systems was first published in 1983 [1]. However, even before then in 1969, Dy Liacco and Kraynak [2] proposed the use of logic programming in a system (ASTA–Automatic System Trouble Analysis) capable of identifying faulted plant and relay and switchgear maloperations. It can be seen that expert system techniques are suited to this problem and make implementation of a system like ASTA much more feasible than it was with conventional languages. However, it was not until 1986 that the application of these techniques to the fault diagnosis problem was reported.

In recent years therefore, utilities have become concerned about managing their systems effectively in view of the voluminous data now available, and the tools suited to a fault diagnosis system have become available. At the start of the research reported in this thesis, a few researchers had applied expert system techniques in the development of fault diagnosis systems for transmission and distribution systems. These suffered from various shortcomings including application restricted to off-line use, unsuitable performance and demands on computing resources, impractical reliance on the infallibility of the SCADA system, and a failure to recognise some of the requirements of a fault diagnosis system. There was clearly a need for research to be undertaken, not merely to attempt to use the new computing technology, but to use the new technology as a tool to solve the problems for control and protection engineers in order that the power system might be more effectively managed.

The research reported in this thesis was engaged in order to develop an expert system for the fault diagnosis of power systems which would cater for the complications involved in a real world power system. During the research duration, other authors have published their work in this same area. While some have reached similar conclusions to those described in this thesis, others have failed to recognise some of the identified problem areas.

The research was commenced with the development of a post fault diagnosis expert system based on the identification of blackout areas. An on-line strategy was then developed from this. The expert system which was created based on this, so-called, “wait-and-see” strategy was realised in Prolog and performed diagnoses based on SCADA alarms. The format of the SCADA alarms was modelled on those of the transmission system of a utility, whose network was used together with hypothetical networks for testing the knowledge. The strategy had an undesirable dependence on the receipt of all switchgear alarms. It was recognised that activity had sometimes to be identified without it having been alarmed. A “hypothesising” strategy was developed and an expert system again realised using predominantly Prolog. As shortcomings in terms of analysis time, indecisiveness, etc. were identified, adaptations were made to the strategy. Some time was spent applying the expert system to the electrical network of a nuclear power station. Experience with different networks proved valuable in verifying the desired portability of the knowledge. Significant consideration was given to the diagnosis knowledge structure, network knowledge structure and diagnosis conclusions structure. Hierarchical structures have been developed in order to deal effectively with the scale of power systems and the nature of their behaviour.

The application of expert system techniques to power system fault diagnosis per se is not novel. However, the research has resulted in several contributions to this important area of work. The depth and nature of the knowledge used in the hypothesising expert system is quite novel. Many expert systems are based on shallow knowledge, and perform their diagnosis with an abstracted view of the power system. At the other end of the scale, the expert system communicates with a numerical simulator of the system behaviour. This thesis describes the use of medium depth knowledge. Qualitative simulation knowledge has been realised which identifies the protection and switchgear response to disturbances,

and does so in a way which models the philosophy of the protection system. This uses protection, switchgear and network qualitative modelling knowledge, by which a significant level of understanding is inherent in the expert system. The presence of this “essential” or fundamental knowledge also means that the details of the network to be diagnosed form more of a definition of the network than a description of its behaviour. As a result, the expert system is more portable to other power systems. Also, this provides the scope for automatically generating a large part of the required database from existing corporate data sources such as CAD drawing files; this will become more important as utilities continue to strive towards coherent corporate data systems. Much faster implementation of the expert system can therefore be achieved.

A new method of dealing with uncertainty was developed for ranking alternative conclusions in order of their chance of being the genuine disturbance. The uncertainty measure is called “relative likelihood”. It was created in order to make use of the available information, but at the same time not require statistics which are not available. Scoring techniques, by which a certain aspect of alternative conclusions are counted, do not make use of all the available information. Conversely, the mathematically sound techniques require quantities which are presently unavailable and make assumptions which are not applicable to the domain of fault diagnosis.

A novel approach has been developed to deal with failures. A qualitative simulator was developed which could also simulate the effect of failures, and do so recognising the relationship between backup and failed activity. However, by using this knowledge to generate possible conclusions involving failures as well as system faults, the expert system generated too many possibilities. Adaptations have been developed by which the consideration of failures is suppressed until there is evidence to indicate that a failure has occurred. A robust and yet not

excessively indecisive fault diagnosis expert system has resulted.

In addition, the expert system makes novel use of a loadflow algorithm. The normal direction of load current flow is used to help direct the searching used to locate a path for fault current.

1.2 Thesis Outline

Chapter 2 of this thesis gives an introduction to expert system techniques. Commencing with the idea behind expert systems and their evolution, this chapter progresses with the means used in realising expert systems and achieving certain characteristics. An introduction to the advantages of using expert system techniques and to existing power engineering applications close chapter 2.

Chapter 3 introduces power systems in general, and their fault diagnosis in particular. It demonstrates how recent changes in power systems have both engendered a greater need for improved fault diagnosis, and have provided the basis for better identification of system conditions (in the form of system data). The provision of an automatic diagnosis system is discussed, following which the suitability of expert system techniques for this task is presented. This chapter includes an extensive critical review of other expert systems which have been developed for the fault diagnosis of power networks.

The diagnostic knowledge aspects of the developed expert system are described in chapter 4 following a brief discussion of some implementation issues. The structuring of the diagnosis conclusions and of the diagnostic process which has resulted from the research is described. A description of the philosophies of the two fault diagnosis strategies which have been developed precedes a detailed presentation of each one's diagnostic knowledge. There then follows a discussion of the problem of dealing with uncertainty and the technique developed for solving it. The chapter closes with a presentation of the ideas in the "deep" end of the diagnostic process.

Chapter 5 describes the network knowledge, and commences with an expression of the significance of this knowledge. The dynamic and static data which is required by the expert system is described and demonstrated. Again, the structure and its importance are shown. The purpose, nature and knowledge of the

qualitative simulation is described. This is followed by a description of the cooperative use of a loadflow by the expert system, and the reason behind it. Finally, the knowledge contained in the expert system to model various aspects of power system equipment behaviour is presented.

Chapter 6 deals generally with the use of the fault diagnosis expert system. Case studies are presented which demonstrate the expert system's response to disturbances. The developed expert system is then evaluated on the basis of certain performance factors, the nature of the knowledge, and a comparison with other fault diagnosis systems. A section is then given over to considering the shorter term and longer term prospects for using the developed expert system.

The final chapter expresses the conclusion that the reported research has resulted in an expert system which could, a short time after a system disturbance, convey valuable diagnostic knowledge to control and protection engineers and do so despite failures of protection equipment, switchgear equipment or the telemetry system. Some pointers for areas of fruitful future work are also given.

An appendix introducing some of the fundamentals of Prolog is included.

1.3 Associated Publications

Much interest has been expressed in the work reported in this thesis, both in person and in receipt of publications. The following lists the publications in association with this thesis:

1. J.R. McDonald, K.L. Lo, and G.M. Burt. "An expert system for post-fault diagnoses in power system networks". In *24th Universities Power Engineering Conference (UPEC)*. Belfast, September 1989.
2. J.R. McDonald, G.M. Burt, S.M. Burt, and R.F.B. McLaren. "An expert system for fault diagnosis on electrical power networks". *Distribution Developments*, pp 15–20, October 1990.
3. D. Cockburn, J.R. McDonald, G. Burt, J.R. Brailsford, J. Beaton, and K.L. Lo. "Expert systems for on-line fault diagnosis in electrical power networks". In *11th International Conference on Electricity Distribution (CIRED)*. Liège, April 1991.
4. G.M. Burt and J.R. McDonald. "Potential advantages of a diagnostic expert system for assisting operator response to system events". In *3rd International Conference on Power System Monitoring and Control*, pp 222–224. London, June 1991.
5. G.M. Burt and J.R. McDonald. "RESPONDD—An expert system for fault diagnosis in power networks". In *26th Universities Power Engineering Conference (UPEC)*. Brighton, September 1991.

6. J.R. McDonald, G.M. Burt, and D.J. Young. "Alarm processing and fault diagnosis using knowledge based systems for transmission and distribution network control". In *IEEE PES 1991 Transmission and Distribution Conference*. Dallas, September 1991.
7. J.R. McDonald, G.M. Burt, and D.J. Young. "Alarm processing and fault diagnosis using knowledge based systems for transmission and distribution network control". *IEEE Transactions on Power Systems*, Vol. PWRS-7, No. 3, pp 1292-1298, August 1992.

Chapter 2

Expert Systems

2.1 Introduction to Artificial Intelligence

Expert systems have evolved as a result of work carried out in the area of Artificial Intelligence (AI). AI activity is concerned with the production of computer systems which exhibit certain abilities natural to humans. Traditional computer science is concerned with producing systems to perform tasks which are problematic to humans, and thus prone to error; numerical computation has conventionally been the main tasks to be performed by computers. (Babbage's difference engine [3], the forerunner of computers, was designed in order to produce more accurate mathematical tables.) AI, on the other hand, covers the fields of knowledge based or expert systems, robotics, natural language and machine vision. Research in AI is said to originate from 1956 [4], while the phrase "artificial intelligence" appears to have been coined in 1958 [5]. Since the appearance of the first practical commercial expert system in 1971 (DENDRAL, a system for identifying molecular compounds), expert systems have received very widespread interest.

2.2 The Concept and Evolution of Expert Systems

Expert systems or knowledge based systems are an attempt to replicate the reasoning of an expert in solving fairly complex real-world problems in a narrow domain (or field). These problems required the intelligence of a human expert, and often relate to diagnosis and decision-making tasks. The function of an expert system for diagnosis is to infer causal, unseen occurrences or unknown items of information (called “hypotheses”), from caused, seen occurrences or known items of information ie. attainable bodies of evidence. (Decision-making expert systems infer information from evidence, and then propose a decision based on its findings; these shall not be further discussed in this thesis.) The process by which expert systems work is therefore called “inference”. The expert system methodology provides a means of capturing such scarce knowledge, and preserving, transferring and documenting it. Indeed the process of developing a work scheduling expert system for a British Alcan plant, resulted in a 10% improvement in plant operation in itself, without incorporating computing, due to the insight it gave to the working of the plant [6].

The tasks for which human experts are valued, are knowledge intensive and require scarce expertise. They are commonly described by heuristics (judgemental rules like “rules-of-thumb”), have no established theory and use information which is not always precise. Expert systems are knowledge based, and suited for such tasks. The knowledge may be acquired directly from a human expert, or from appropriate literature. Therefore an expert system would not be expected to perform better than a human expert in an “abstracted environment”. An abstracted environment is a non-real world environment, like under laboratory conditions. However it has been known for an expert system, developed using the

knowledge of many experts, to outperform any individual expert (reference cited in [7]). Moreover, as well as continuous availability, an expert system would at least provide consistent performance in stress-related circumstances, conditions which would reduce the effectiveness of a human expert. One of the motivations for the development of expert systems is to overcome the “cognitive barriers” which are encountered during stress-provoking circumstances or when conventional computer tools are used beyond their design limitations.

The knowledge incorporated into an expert system will usually be relational, fuzzy, modular and require heuristic searching. Fuzziness is the property by which something does not have clearly defined boundaries. Heuristic searching uses judgemental rules to limit the area of state space (ie. that space which contains all possible combinations of data) which is considered when searching for a solution. The knowledge in an understanding oriented, or “competent” expert system, models human reasoning. User oriented expert systems, however, require knowledge that simply gives the correct answer—they will contain an abstract view of the domain. In either case the expert system’s role may often be viewed as a distillation of knowledge—convergent, rather than divergent as many conventional programs. Clearly a very important part in creating an expert system is acquiring the knowledge, ie. “knowledge acquisition”.

The origins of expert systems lie in a different artificial intelligence philosophy. After Turing [8] first set the philosophical tone in 1950, the first significant achievement was the General Problem Solver (GPS) in 1957. This program was designed to tackle problems (like theorem proving) by using domain-independent problem solving methods. The “paradigm” (pattern or programming style) can therefore be called “power based”, since it was an attempt to capture all-encompassing reasoning power. GPS was, however, unsuccessful in achieving its goals. The shift from this paradigm appeared in the next decade,

part way through the DENDRAL project, when it was decided to incorporate AI techniques. With the incorporation of specialist expertise, DENDRAL became a powerful program which could identify molecular compounds from analytical data—a narrow and highly specialised domain. Retrospectively it can be seen that this was an important shift from power based to knowledge based systems. What followed, in the growth of knowledge based or expert systems, gives weight to the worth of the knowledge based approach.

The most commonly referenced expert systems are MYCIN and PROSPECTOR. MYCIN [9] is an expert system for infectious blood disease diagnosis and therapy, and was developed in 1973. PROSPECTOR [10] is an aid to evaluating the favourability of an exploration region for ore deposits of particular types, and was developed in 1979.

The early expert systems were all consultant systems; they operated in isolation, with their only input being through the user. Such stand alone arrangements are not suitable for applications in many domains, including power engineering. However, many more areas gained from their adoption when it was realised that expert systems need not operate in isolation and could be integrated with other systems, like data acquisition systems. The Electric Power Research Institute's (EPRI) first development using expert system ideas monitored 22 sensors [4]. The system, for performing vibration monitoring and malfunction diagnosis of large pumps in nuclear power stations, became operational in 1978, development having started in 1976.

Whereas some (eg. in the medical field) remain sceptical of expert systems, some professions have accepted the technology. In 1981 a cement manufacturer replaced an ineffective model-based controller with an expert system [11]. This expert system forms a closed loop controller and works automatically.

The first important application of expert system technology to power system

operation occurred in 1983 when Sakaguchi and Matsumoto [1] presented an expert system for power system restoration. Since then, the interest in expert systems has continued to grow.

“The design of expert systems requires support for knowledge acquisition, knowledge representation, knowledge based inference, and search based computation, along with facilities for explanation, probabilistic reasoning, and the development of expert system shells”[12]. With the exception of the final point, which is beyond the remit of this application-directed thesis, these points will be briefly discussed in the following sections. By way of explanation, an expert system shell is developed from an expert system by removing the domain specific knowledge. The remaining reasoning knowledge can then be paired with a different domain’s knowledge, and the problem solving in the new domain facilitated by the same techniques.

2.3 Realisation of Expert Systems

2.3.1 Structure

Expert systems have generally used a standard structure. This structure contains four major parts: knowledge base, inference engine, user interface and designer utilities. Their arrangement is as shown in Figure 2.1.

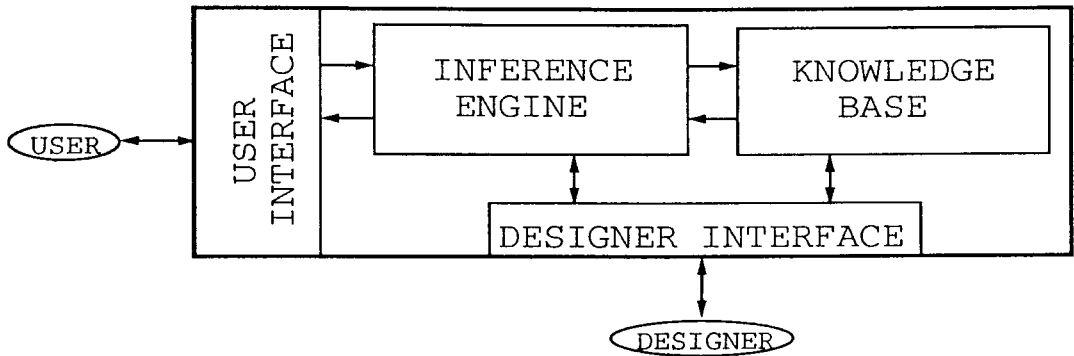


Figure 2.1: The basic structure of expert systems

It is the knowledge base which contains the expert knowledge; the knowledge can be represented by various means as seen in the following paragraphs. The knowledge base often consists of two identifiably separate parts: a knowledge source and a database. The knowledge source contains only the expert knowledge. This usually corresponds with the knowledge which is specific to the domain of expertise but is independent of the particular system, case, etc. to which the expert system is being applied. A particular type of knowledge often included is “meta-knowledge”; meta-knowledge is knowledge about knowledge and knowledge about how to reason to obtain implicit knowledge. As an illustration, meta-knowledge which uses voltage level alarms to imply the status of breakers whose status is not known directly but which should have tripped to isolate a fault is expressed in English:

```

If explanation includes fault
  and explanation includes breaker
  and breaker is not wired into data acquisition system,
then apply knowledge to determine breaker status from voltage
  level alarms.

```

Knowledge of the application of the expert system (be it system, equipment, or case) is contained in the database which will change with the inclusion of new data or results derived from operation of the expert system.

There need not be only one knowledge base. Where identifiable sub-tasks exist or where abstractions (different levels of detail) are used, the knowledge can be distributed accordingly among several knowledge bases. Often, these knowledge bases communicate via a database called a “blackboard”; these systems are called blackboard systems. Figure 2.2 illustrates a detailed fault diagnosis blackboard system, which uses various depths of diagnosis knowledge base (KB) and numerical programs.

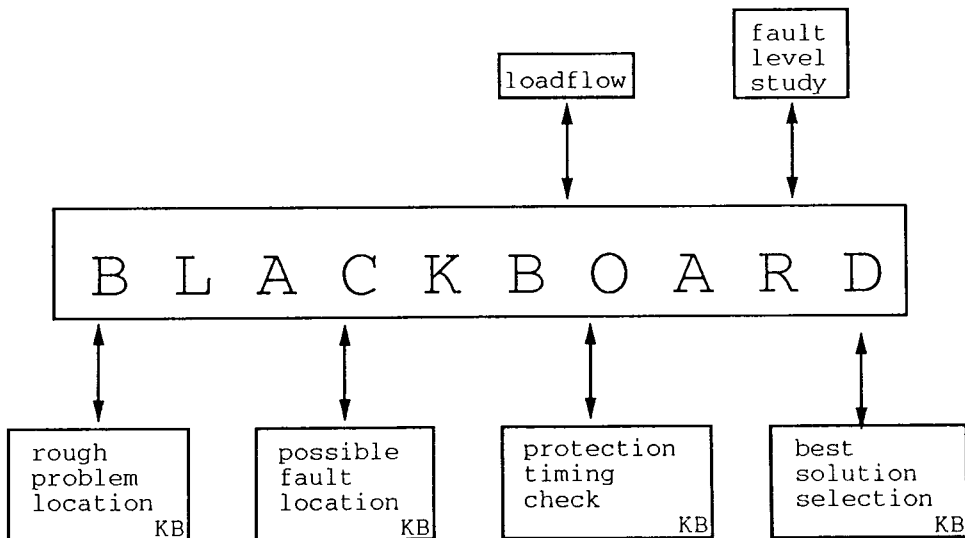


Figure 2.2: An example of a blackboard system

The inference engine uses the knowledge in the knowledge source. The function of the inference engine is to apply the knowledge (ie. reason with the knowledge) in order to draw inference from it and the available body of evidence.

Consequently a reasoning strategy is required. An explanation facility may be provided by the inference engine, or instead, a separate explanation expert system may be included with the core expert system. The intelligent behaviour of expert systems is not due to the control strategy of the inference engine. Reference [7] quotes Feigenbaum thus: "...in the knowledge is the power. The interesting action arises from the knowledge base, not the inference engine."

The user interface is a critical part of any expert system. A successful expert system will not only fulfill its claim to perform tasks which previously required human expertise, but also it will be accepted by the ultimate users. A user friendly interface is therefore very important. Indeed, for consultant expert systems, which are dependent on the user for their data, a good user interface is imperative. This is at least as critical as good communication skills in a human expert.

The designer interface provides the expert system designer with the tools to modify the knowledge in the knowledge base. The facilities may range from a simple editor, to complex consistency-checking tools. The designer interface element is also called the builder utilities or knowledge engineering interface.

In addition to the above four main elements, the importance of an external interface is increasing. Many applications require data from electronic systems and so some interface facilities are required. Also, by integrating expert systems with analytical tools, expert systems will be suitable for applications in many more problem areas (eg. [13]).

2.3.2 Knowledge Acquisition

There are two main sources of knowledge: human experts and literature. Where they exist, appropriate books and guidelines/codes of practice are very useful. Often these do not exist, however, so knowledge has to be elicited from experts. The most popular method, though not the most effective, is by unstructured interview.

Structured interviews can be more efficient in obtaining the desired knowledge. Experts are, however, often not aware of, or able to describe, the knowledge they use. In such cases, knowledge can be acquired by performance evaluation of the expert at work on familiar, tough, limited-information or constrained-processing tasks which are either genuine or simulated. In this way experts are used for what they are good at, namely generating or scrutinising examples. Often, at least in the power engineering domain, it has been people functioning within the domain who have developed expert systems. This has the advantage that many of the main concepts are already understood. The reasoning underlying theories, and simulations are thus extensively used, as well as expert consultation.

Engelmore [14] proposed an automatic approach for acquiring knowledge for a particular task. A general purpose model of the domain is developed, and is operated on by developed knowledge compilation techniques to transform the general knowledge into task specific knowledge.

2.3.3 Means of Implementation

The early expert systems were written in specialised language interfaces built upon languages such as FORTRAN and Lisp. Conventional languages have been used; EPRI's first system was realised in FORTRAN IV. Such algorithmic languages could lead to the most efficient systems in terms of processing time and memory on current hardware, but at a very high human development cost, being less naturally suited to such AI applications. The languages which have had lasting appeal are Lisp, Prolog and OPS (Official Production System). Lisp and Prolog provide for much creativity which can result in an efficient implementation.

As well as languages, so called expert system shells have been used to develop expert systems. In concept, a shell is an expert system with its expert knowledge removed. What is left includes inference mechanisms, knowledge representation

structures, explanation facilities, means of dealing with uncertainty, and development and debugging facilities. Consequently using a shell means the designer has less to implement in order to realise an expert system. But, as the design methodology is the best suited to only that domain for which a shell's source expert system was developed, a shell is only advantageous where the domain closely matches that particular class of problems.

Expert system development toolkits offer a range of facilities of more general relevance. Toolkits, like Knowledge Engineering Environment (KEE), Automatic Reasoning Tool (ART) and Lisp Object Oriented Programming System (LOOPS), offer a variety of methods which can be put together to produce the knowledge representation and reasoning process that best suits the task. Although toolkits are costly in computing resources, they have the advantage of being able to support multiple "paradigms". Yazdani [15] gives a more detailed review of the different implementation techniques.

2.3.4 Methods of Realisation

Such means of producing expert systems support mainly four paradigms: logic programming, rule based programming, object-oriented programming and access-oriented programming [16]. Before going on to define these paradigms, the knowledge representation methods of which they make use will be described. Knowledge is variously represented as rules, frames, taxonomies and semantic (or associative) networks. Other means are available, like scripts and ISA hierarchies, but these are less commonly used. A rule is a statement which relates an antecedent (or premise) to a consequent (or conclusion). The antecedent can be one condition or several conditions (logically AND-ed together). Where the consequents relate to actions to be performed, each may likewise be made up of several actions. Production rules are particularly common. The following example, in which variables

are surrounded by < ... >, illustrates the form of such:

```
if relay is <Relay>
  and relay:state is 'operated'
  and relay:type is 'unit'
  and relay:cts is <CTs>
then fault is area_between_cts(<CTs>).
```

A frame represents a list of properties specific to a particular object. Each frame contains standard slots which define state and behaviour. A slot may contain a value or a set of rules which define how to get the value. Thus, they can be of the form shown by the following example:

```
relay
  name      : r1
  type      : unit
  cts       : [ct1,ct2]
  status    : operated
  protects: unknown, method:area_between_cts
```

Taxonomies organise knowledge into hierarchies. They use the structure of frames, and are also able to describe the relationship between frames. This ability to point to other frames is now commonly assumed to be a standard feature in frames, and the term taxonomies is seldom used.

Semantic networks also describe both objects and their relationships to other objects. They do not, however, use the frame structure. Semantic networks are usually illustrated with “nodes” representing objects and “arcs” representing relationships as shown by the example in Figure 2.3.

These knowledge representation means are variously used by the common paradigms. Logic programming involves the use of logical rules (or statements). By far, the most common means of implementing this paradigm is Prolog. Prolog is a general purpose language which is based on first order predicate logic. A rule

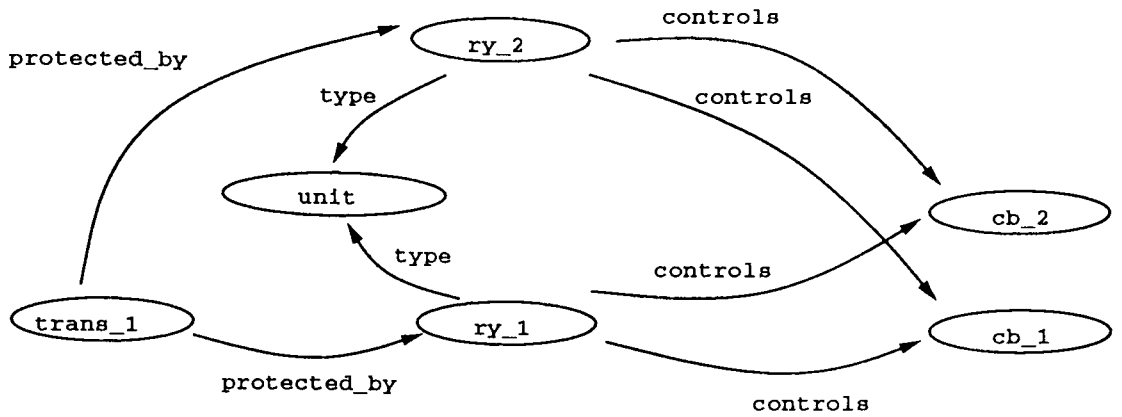


Figure 2.3: A semantic network

is represented as shown in the following example (variables commence with upper case letters):

```

maybe_trans_fault(Fault):-
    just_operated(Relay),
    protects(Relay,Faults),
    member(Fault,Faults),
    faulted_plant(Fault,Location),
    a_transformer(Location).
  
```

Prolog is extremely flexible; it can support knowledge representation in production rules, frames and semantic networks, and can support aspects of the other paradigms.

Rule based programming is much like logic programming. The knowledge is represented as production rules for this paradigm.

Object-oriented programming involves objects which are grouped into classes of objects with the same structure and behaviour. Knowledge is represented as taxonomies: each object has a frame with slots, each of which may contain values, pointers to other objects or sets of rules. An object or “instance” of a sub-class inherits all the slots of the classes to which the sub-class belongs, while the objects’ information is manipulated by “methods”. A typical scheme is shown in Figure 2.4.

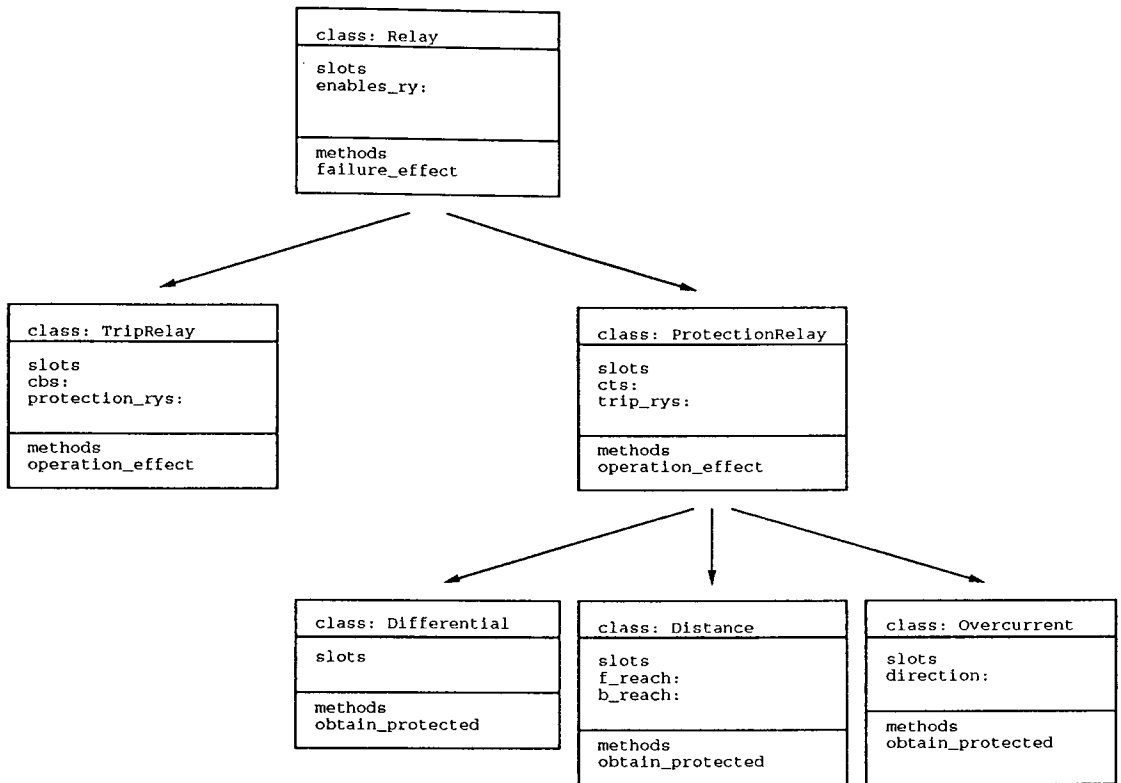


Figure 2.4: A schematic of object-oriented programming

Access-oriented programming also uses frames, but in addition, facilitates the separation of monitoring program events from the processes that can cause those events.

2.3.5 Techniques

As the application reported in this thesis is fault diagnosis, it is aspects particularly relevant to diagnosis systems which shall be discussed. In expert systems following the various paradigms, several techniques recur in isolation or in combination.

2.3.5.1 Searching

A fundamental technique in expert systems is searching of the state space. The state space holds all possible combinations of variables in the problem domain. In

order to find an acceptable combination of variables (a solution), the state space must be searched (or navigated). Searching techniques can be broadly classified as either width-first or depth-first. A width-first search will progress on several routes to a solution simultaneously, with each at a similar stage of development. A depth-first search, on the other hand, will first attempt to reach one result.

To obtain one solution, width-first searches generally take longer and have greater memory requirements. However, it is often necessary to obtain the best solution (whatever the measure of “best” is). Using the depth-first or the width-first techniques, an exhaustive search can be performed to obtain this optimum solution. This can be very time consuming. Therefore, a heuristic search strategy is often preferable.

A heuristic search makes use of heuristics to select routes to take first where an alternative exists. This approach allows the “best” solution to be reached without every alternative being compared. Clearly, the complete elimination of all but one of the solution routes, is only applicable in well-structured domains in which the best solution consists of the best partial solution at every level. However this strategy is sometimes adopted in other domains, and is used to find one of the better solutions—the time saved from not having to do an exhaustive search compensates the lack of guaranteed optimality. An example of such a task is the reconfiguration of electrical networks. As any reconfiguration proposals must be checked by numerical analysis, the use of heuristics to guide the search for a way of restoring supply to disconnected customers [17] can save time, and hence money.

An alternative technique is bi-directional searching. However this entails the risk of routes passing in the middle and not meeting to form a complete route. Consequently it appears to be seldom adopted.

2.3.5.2 Chaining

Except in the simplest of cases, in coming to a solution, an expert system needs to chain through its knowledge. This is a particular form of searching. There are two techniques, backward chaining and forward chaining, and these are of greatest relevance to the rule based paradigm. Backward chaining inference is also called goal directed and was used in MYCIN. Commencing with a possible conclusion (a goal or final “hypothesis”), the inference chains through the knowledge in order to reach and check for the available data or evidence. Figure 2.5 illustrates this. In the figure, the unshaded nodes ($E1, \dots, E3$) represent evidence, straight




Figure 2.5: Diagram showing backward chaining

lines are lines of inference, shaded nodes represent inferred knowledge, label $FH1$ indicates a final hypothesis or conclusion, and the broken lines show the path of reasoning.

A rule based expert system for on-line valve monitoring and diagnosis [18] uses the backward chaining technique. The expert system has knowledge of 25 possible causes of sensor signal deviations from the normal. As these possible conclusions are static and of a reasonable number, standard backward chaining is suitable. By backward chaining, evidence is gathered for each possible malfunction. The level

of supportive evidence then indicates the likelihood that a malfunction exists.

In contrast, the forward chaining inference is data directed, and is often called such. Existing data is used to infer new, derived data which is used in subsequent inference. Conclusions are thus reached as illustrated in Figure 2.6, which follows the same convention as Figure 2.5.




Figure 2.6: Diagram showing forward chaining

A rule-based expert system for controlling a cement-kiln [11] effectively uses the forward chaining technique. Signals from temperature and gas sensors are assigned LOW, OKAY and HIGH values. Based on these, and the known status of the kiln, controller actions are proposed by the rules. These proposed actions are merged to resolve adjustments to controller set-points. The total number of possible combinations of set-point adjustments is extensive, and so a data driven approach is preferable.

The first chaining route attempted will not always lead to a desired termination. Alternative knowledge can be used by backtracking, a common technique in expert systems. Backtracking is the process of going in reverse along an already traversed path or procedure, in order to find an alternative solution. For example, Figure 2.7 illustrates how, with all evidence except *E1* available, forward chaining

and backtracking finds conclusion *FH2* after the reasoning fails to conclude on *FH1*. Again the same symbols as Figure 2.5 are used.

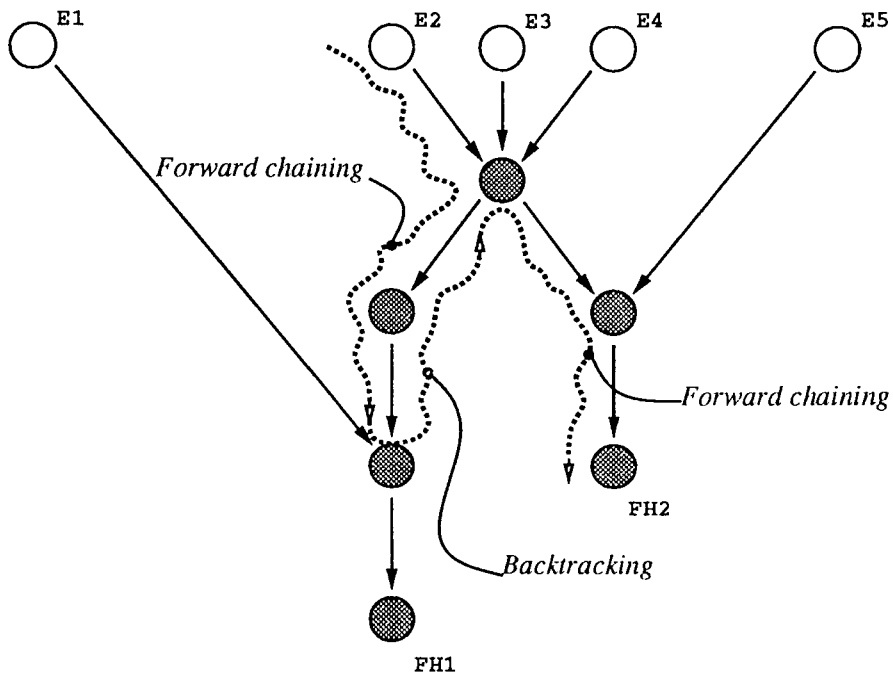


Figure 2.7: Diagram showing backtracking

2.3.5.3 Pattern matching

Pattern matching is another widely used technique. This technique compares a combination of attributes with another combination, looking for equality. It is used to match rule antecedents with known facts (forward chaining), and to match rule consequents with a set goal (backward chaining). (Using the principles of fuzzy logic, an exact match is not required; see section 2.4.2.3.)

2.3.6 Strategies

The basic generate and test strategy (or “synthesis” based strategy) uses backward chaining. Having generated a hypothesis or goal, the knowledge is applied to derive the expected observable effects, and these are compared with the known

evidence. The correspondence between the expected and actual evidence has implications for the match between the hypothesis and the actual situation.

The “analysis” based strategy attempts to resolve from the available evidence what could have produced it; it thus uses forward chaining. Some expert systems, like PROSPECTOR, can select appropriately between forward and backward chaining.

The complexity of tasks to which expert systems have been applied, has resulted in various strategies being adopted in order to make problems more manageable. Such strategies include: the space transformation strategy; classification model; planning strategy; least-commitment principle and constraint propagation.

The space transformation strategy involves changing the domain representation by problem decomposition and abstraction. Hence a large complex problem is converted into a collection of smaller sub-problems. A blackboard system could be the means of realising this strategy: the blackboard would contain the emerging partial solutions to the problems at different levels of abstraction.

The classification model is a framework for organising reasoning according to classification as findings, intermediate hypotheses, and conclusion hypotheses. Finding-to-(inferred)finding, finding-to-hypothesis and hypothesis-to-hypothesis rules are applied based on the initial observations, until a final conclusion is reached.

In domains where irrecoverable or catastrophic conditions exist, the planning strategy is particularly appropriate. According to this strategy, possible actions are simulated in a domain model, and those actions which would be irrecoverable in the real world are withdrawn.

According to the least-commitment principle and constraint propagation, the strategy is to suspend a particular line of reasoning and only restart it on com-

pletion of another goal. The idea is to postpone decisions until the maximum amount of information is available. Expert systems in domains with long lines of inference between evidence and conclusions, where following any line of reasoning is time consuming, would gain from the use of this strategy.

Most diagnosis strategies can be described in terms of three types of reasoning, namely abduction, deduction and induction. Abduction reasons from evidence to a hypothesis which has to be tested. By deduction the consequences of a hypothesis are obtained. The validity of the hypothesis is inferred by induction.

There are many more reasoning strategies which have been developed, like taxonomic reasoning and default reasoning [19] to name but two; only a few strategies and techniques have been presented above. But those described are fairly widespread, and often form the backbone of other strategies. In the following section the subject of dealing with uncertainty will be dealt with, for which many different strategies have been adopted.

2.4 Dealing with Uncertainty

2.4.1 Sources

The existence of uncertainty is a common aspect in many tasks for which expert systems are being applied. The creation of suitable techniques to deal with uncertainty is therefore important. Uncertainty techniques have had specific relevance to rule based expert systems, and techniques presented later will relate specifically to knowledge represented by production rules. The sources of this uncertainty are diverse but may be grouped under several headings.

- **Unreliable information:-** Unreliable knowledge can exist in the knowledge base, in the form of implications rather than concrete relationships, or in the data supplied to a working expert system.
- **Incomplete information:-** Information requested by a consultant expert system may neither be known nor ascertainable. Likewise, automatic means of data retrieval might not provide the desired information before the reasoning is commenced, or fail to provide it at all.
- **Languages lack precision:-** The lack of precision in languages may arise in the knowledge acquisition and knowledge representation processes, or in the dialogue between a user and an expert system.
- **Contradictory viewpoints:-** Although access to several experts is advantageous in that a greater breadth of expertise may be captured, a consensus may not be obtainable. Where such contradiction exists, the development of a coherent knowledge base is difficult; varying degrees of uncertainty must be applied to each expert's knowledge.
- **Lack of determinacy:-** Domains do not always exhibit a one to one correspondence between evidence/data and conclusions. Some ranking is

required to determine a conclusion.

2.4.2 Techniques

There are several techniques for dealing with uncertainty which are well known. Subjective probability theory, Dempster-Shafer theory and possibility theory had their theoretical bases established before the widespread interest in expert systems, but have since been incorporated into expert systems in order to manage uncertainty. The theory of certainty factors and the theory of likelihood ratios were developed specifically for expert systems (MYCIN and PROSPECTOR respectively). Whereas these methods relate in particular to uncertainty propagation through a rule base, the unverified evidence quantification method deals solely with a comparison of competing conclusions. The theory of endorsements was developed in order to deal with uncertainty qualitatively.

As this thesis reports on work in the field of diagnosis, the quantitative techniques will be briefly discussed with respect to the general production rule:

$$\text{if } E \text{ then } H$$

where E is evidence which can be caused by the unobservable hypothesis H . This is not restrictive, but is simply for illustrative purposes. For example, rather than evidence, the rule's antecedents could be intermediate hypotheses. Basically the techniques centre around trying to obtain the uncertainty that H is indeed the cause of E being observed. Subscripts will be applied to E and H to indicate different pieces of evidence or hypotheses respectively.

The following standard symbols shall be used.

\neg	negation
Σ	summation function
Π	product function
max	maximum function
min	minimum function
\in	member
\subset	subset
\cap	intersect
\emptyset	empty set

2.4.2.1 Subjective (or Bayesian) Probability Theory

Objective probability is measured using frequency ratio techniques, and is preferable being fully mathematically supported. However, as the required statistical information is commonly not available, subjective probability can often only be used. Subjective probability is a value between 0 and 1, which is set higher as the confidence which an individual (the expert) has that an event is true grows. The collection of all possible elementary events is called the sample space.

A further subdivision of probabilities can be shown with respect to the general production rule. The prior probability of hypothesis H , denoted $p(H)$, is the unconditional probability assigned to H prior to the observation of any evidence; the posterior probability of hypothesis H , denoted $p(H|E)$, is the conditional probability that H is true, given that E is true (ie. evidence E is observed).

For this same case of H causing E , and E being observable, Bayes' rule defines the probability of H being the unseen cause of observing E as:

$$p(H|E) = \frac{p(E|H) \times p(H)}{p(E)} \quad (2.4.1)$$

Bayes' rule involves the prior probability that evidence will occur and that a hypothesis will occur, and the conditional probability that evidence will occur given that the hypothesis has occurred—all three values relate to tangible events and should be quantifiable.

It can be shown that this can be expanded to the case where multiple hypotheses (H_1, H_2, \dots, H_m) exist. By assuming that the hypotheses are mutually exclusive and exhaustive, then

$$p(H_i|E) = \frac{p(E|H_i) \times P(H_i)}{\sum_{k=1}^m p(E|H_k) \times p(H_k)} \quad (2.4.2)$$

Where there is also multiple evidence, say E_1, E_2, \dots, E_n , then the above expression would contain $p(E_1 E_2 \dots E_n | H_l)$ for all l , ie. it would require the conditional probabilities of all possible combinations of evidence given all possible hypotheses. For most uses this is clearly impractical. However, if we assume condition independence among pieces of evidence given a hypothesis, then the product rule applies:

$$p(E_1 E_2 \dots E_n | H_l) = \prod_{j=1}^n p(E_j | H_l) \quad (2.4.3)$$

Substituting this in the above expression for $p(H_i|E)$, equation 2.4.2, gives a workable means of obtaining $p(H_i|E_1 E_2 \dots E_n)$.

Attempts have been made to remove the requirement for conditional independence. However, as well as the independence of evidence, this method has several other problems: exhaustive sets of mutually exclusive hypotheses are seldom possible to create; a large number of values is required of the expert (at least, the number of hypotheses multiplied by the number of pieces of evidence, plus the number of hypotheses); it requires a very large number of calculations. The method's main attribute is its well-formed methodology.

2.4.2.2 Dempster-Shafer Theory

Whereas probability theory relates to a solution space, Dempster-Shafer theory involves the frame of discernment. The frame of discernment is an exhaustive set of mutually exclusive events. Thus elements in the frame of discernment may contain one event or a combination of events. The basic quantity assigned to a particular element H_s is denoted $m(H_s)$, has a value between 0 and 1, and refers specifically to the set H_s (not to individual events). The (total) belief in H_s , denoted $Bel(H_s)$, is thus

$$Bel(H_s) = \sum m(H'_s) \quad \text{for all } H'_s \text{ such that } H'_s \subset H_s \quad (2.4.4)$$

That is, the belief in an element is the sum of the measures of those elements in the frame of discernment which are subsets of the element. For example, if $H_s = \{H_1, H_2\}$ then

$$Bel(\{H_1, H_2\}) = m(\{H_1\}) + m(\{H_2\}) + m(\{H_1, H_2\}) \quad (2.4.5)$$

Also, the plausibility of H_s , denoted $Pl(H_s)$, is defined as

$$Pl(H_s) = 1 - Bel(\neg H_s) \quad (2.4.6)$$

where $\neg H_s$ means everything but H_s . The belief and plausibility are viewed as upper and lower probabilities.

Where there are multiple beliefs in the same hypothesis, the two belief functions are combined by computing their orthogonal sums with Dempster's rule of combination. For the two quantities m_1 and m_2 supporting hypothesis set H_s ,

the combined probability assignment is given by

$$m_1 \oplus m_2(H_s) = \frac{\sum_{H_X \cap H_Y = H_s} m_1(H_X) \times m_2(H_Y)}{\sum_{H_X \cap H_Y \neq \emptyset} m_1(H_X) \times m_2(H_Y)} \quad (2.4.7)$$

where H_X and H_Y are elements in the frame of discernment. That is, the combined quantity for a hypothesis set is a ratio of sums of terms, each term being the product of values for two elements obtained from each of the two probability assignments, and each sum including all possible combinations of elements which satisfy a condition. The numerator condition is that the intersection of each of the two elements must make up the hypothesis set. The elements in the denominator term must have a non-zero intersection.

Although this technique has certain advantages, explicitly representing ignorance for example, it has serious shortcomings. Experts must provide beliefs in all possible combinations of hypotheses. Moreover no effective procedure exists for drawing inferences from belief functions as yet. Almost no expert systems have been built using this theory.

2.4.2.3 Possibility Theory

Possibility theory has its basis in fuzzy set theory. A fuzzy set is a class of objects with a continuous grade of membership in the set of all objects (say U). The degree of membership of an object u in a fuzzy set A is given by the membership function $\mu_A(u)$; the value of $\mu_A(u)$ is between 0 and 1, 1 being the highest grade of membership. The elementary set operators act as follows:

$$\mu_{\neg A}(u) = 1 - \mu_A(u) \quad (2.4.8)$$

$$\mu_{A \cup B}(u) = \max[\mu_A(u), \mu_B(u)] \quad (2.4.9)$$

$$\mu_{A \cap B}(u) = \min[\mu_A(u), \mu_B(u)] \quad (2.4.10)$$

Possibility theory relates possibility distributions with fuzzy set membership. Consider the case where evidence E in the general production rule is given by

$$X \text{ is } A$$

where A is a fuzzy set with membership function μ_A . For this the possibility distribution function π_X is defined by:

$$\pi_X(u) = \mu_A(u) \quad (2.4.11)$$

The possibility measure defining the possibility that a value x belongs to finite set A is denoted $Poss\{x \in A\}$, and is defined as:

$$Poss\{x \in A\} = \max_{u \in A} [\pi_X(u)] \quad (2.4.12)$$

Using the above form for H as well as E , so that the rule is,

$$\text{if } X \text{ is } A \text{ then } Y \text{ is } B$$

then,

$$\mu_{H|E} = \mu_{B|A}(u, v) = \min[1, (1 - \mu_A(u) + \mu_B(v))] \quad (2.4.13)$$

That is, the conditional possibility distribution is that of the antecedent subtracted from that of the consequent and plus one, except that a maximum of 1 exists. The possibility measure follows as above:

$$Poss\{y \in B | x \in A\} = \max_{u \in A} [\min[1, (1 - \mu_A(u) + \mu_B(v))]] \quad (2.4.14)$$

Where the evidence is made up of two AND-ed conditions, say,

X_1 is A_1 and X_2 is A_2

then,

$$\mu_{A_1 \text{ and } A_2}(u, v) = \min[\mu_{A_1}(u), \mu_{A_2}(v)] \quad (2.4.15)$$

Similarly

$$\mu_{A_1 \text{ or } A_2}(u, v) = \max[\mu_{A_1}(u), \mu_{A_2}(v)] \quad (2.4.16)$$

An additional technique which fuzzy logic provides is that, in pattern matching, the matching does not need to be exact. Modifiers, like “very”, “quite” and “extremely”, apply functions to the possibility distributions. For example,

X is mA

where m is the “very” modifier, could have, as suggested by Zadeh [20], the membership function

$$\mu_{mA}(u) = \mu_{\text{very}A}(u) = \mu_A^2(u) \quad (2.4.17)$$

(In addition, similar functions are used by the permitted structure

it is τ that X is A

where τ is a truth quantifier like “very true”.)

Therefore, given

X is mA

if X is A then Y is B

it is concluded that

$$Y \text{ is } mB \quad (2.4.18)$$

$$Poss\{y \in mB \mid x \in mA\} = \max_{u \in U} [\min[\mu_{mA}(u), \mu_{B|A}(u, v)]]$$

Although possibility theory permits the explicit representation of natural fuzziness, there are several problems associated with the theory[21]. For example, if two mutually exclusive propositions are AND-ed, equation 2.4.15 wrongly gives a non-zero result.

2.4.2.4 Certainty Factors

Certainty factors were developed specifically for MYCIN, but have been used in many expert systems. For each production rule there is an *a priori* certainty factor (CF) assigned:

$$\text{if } E \text{ then } (CF) H$$

This value is set by the expert as his subjective judgement of the degree to which the definite occurrence of evidence E is suggestive of the hypothesis H . It is denoted $CF[H, E]$ and has a value between -1 and 1: a positive value indicates that the hypothesis is supported by the evidence, while a negative value is used where the hypothesis is counter-indicated by the evidence. Aspects of evidence likewise have certainty factors ascribed to them. The uncertainty of aspects of evidence are combined in accordance with fuzzy set theory, to obtain one certainty factor for the evidence, $CF[E]$. For example, for the two aspects of evidence E_A and E_B ,

$$CF[E] = CF[E_A \text{ and } E_B] = \min[CF[E_A], CF[E_B]] \quad (2.4.19)$$

The certainty factor of the hypothesis is no longer its assigned *a priori* value.

Rather, it becomes

$$CF[H, E'] = CF[E] \times CF[H, E] \quad (2.4.20)$$

where E' is the actually observed evidence.

Where a hypothesis is multiply related to (supported or opposed), say, evidence E_1, E_2, \dots, E_n then evidence E_i is ignored if $CF[E_i]$ is less than a threshold value (0.2 is used in MYCIN). For the remaining evidence the measures MB and MD are set. MB is the measure of increased belief, while MD is the measure of increased disbelief. They are based on the modified certainty factor of the hypothesis as shown:

$$MB[H, E_i] = \begin{cases} CF[E_i] \times CF[H, E_i] & \text{if } CF[H, E_i] > 0 \\ 0 & \text{otherwise} \end{cases} \quad (2.4.21)$$

$$MD[H, E_i] = \begin{cases} CF[E_i] \times |CF[H, E_i]| & \text{if } CF[H, E_i] < 0 \\ 0 & \text{otherwise} \end{cases} \quad (2.4.22)$$

where E_i is the evidence which satisfies the above comparison with the threshold value.

Where two bodies of evidence, E_i and E_j , support the hypothesis (ie. *a priori* $CF[H, E_i]$ and $CF[H, E_j]$ both greater than 0) then

$$MB[H, E_i E_j] = \begin{cases} 0 & \text{if } MD[H, E_i E_j] = 1 \\ MB[H, E_i] + \\ \quad MB[H, E_j] \times (1 - MB[H, E_i]) & \text{otherwise} \end{cases} \quad (2.4.23)$$

Likewise, where two bodies of evidence, E_i and E_j , are against the hypothesis

(ie. *a priori* $CF[H, E_i]$ and $CF[H, E_j]$ both less than 0) then

$$MD[H, E_i E_j] = \begin{cases} 0 & \text{if } MB[H, E_i E_j] = 1 \\ MD[H, E_i] + \\ \quad MD[H, E_j] \times (1 - MD[H, E_i]) & \text{otherwise} \end{cases} \quad (2.4.24)$$

Having combined all relevant evidence, then $CF[H]_{supporting} = MB[H, E_{all}]$, $CF[H]_{against} = MD[H, E_{all}]$ and the certainty factor for the hypothesis is

$$CF[H] = CF[H]_{supporting} - CF[H]_{against} \quad (2.4.25)$$

The same results can be achieved by expressing the combination of evidence slightly differently. Having obtained the modified certainty factors for the hypothesis (by equation 2.4.20) based on the actually observed evidence E'_i ,

$$CF[H]_{supporting} = 1 - \prod_i (1 - CF[H, E'_i]) \quad \begin{array}{l} \text{for all } i, \\ \text{where } CF[H, E'_i] > 0 \end{array} \quad (2.4.26)$$

$$CF[H]_{against} = 1 - \prod_i (1 - |CF[H, E'_i]|) \quad \begin{array}{l} \text{for all } i, \\ \text{where } CF[H, E'_i] < 0 \end{array} \quad (2.4.27)$$

A modification of the original formulation for combining evidence E_i and E_j has been suggested [21]. Using equation 2.4.20 with the actual evidence E_i and

E_j ,

$$CF[H, E_i E_j] = \begin{cases} CF[H, E'_i] + CF[H, E'_j] \times (1 - CF[H, E'_i]) & \text{if } CF[H, E'_i] \text{ and} \\ & CF[H, E'_j] > 0 \\ CF[H, E'_i] + CF[H, E'_j] \times (1 + CF[H, E'_i]) & \text{if } CF[H, E'_i] \text{ and} \\ & CF[H, E'_j] < 0 \\ \frac{CF[H, E'_i] + CF[H, E'_j]}{1 - \min[|CF[H, E'_i]|, |CF[H, E'_j]|]} & \text{otherwise} \end{cases} \quad (2.4.28)$$

This certainty factor method has been used in several rule based expert systems. Mostly, though, these are based on MYCIN: it has not found wide acceptance. The certainty factors are not accurate statistical measures, but it is reported that it was found that the ranking of conclusions in the MYCIN expert system, resulting from certainty factors, was accurate. Despite this, some researchers have claimed to have found serious weaknesses in the method. One of these claimed is the limited validity of the implicit assumption of independence among hypotheses.

In addition to the original certainty factor method, some modified certainty factor schemes have been used. These can differ in their treatment of combined antecedents or aspects of evidence ($CF[E]$), in their treatment of the effect of evidence uncertainty ($CF[H, E']$), or in their treatment of combined multiple evidence ($CF[H]_{supporting}, CF[H]_{against}$). Marcus [22] describes a “standard” method, a fuzzy based method and a Bayesian method, while Winston [23] describes an “ad hoc” method and a Bayesian method.

Marcus' "standard" method redefines the combination of multiple evidence as;

$$CF[H]_{supporting} = \max_i [CF[H, E'_i]_{supporting}] \quad \begin{array}{l} \text{for all } i \\ \text{where } CF[H, E'_i] > 0 \end{array} \quad (2.4.29)$$

$$CF[H]_{against} = \max_i [|CF[H, E'_i]_{against}|] \quad \begin{array}{l} \text{for all } i \\ \text{where } CF[H, E'_i] < 0 \end{array} \quad (2.4.30)$$

The fuzzy based method maintains the *a priori* hypothesis certainty factor, ie. $CF[H, E'] = CF[H, E]$, and defines the hypothesis certainty factor where there is multiple evidence as;

$$|CF[H]| = \max_{\text{all } i} [|CF[H, E'_i]|] \quad (2.4.31)$$

Marcus' Bayesian method redefines the method of combining antecedent certainty factors:

$$CF[E] = CF[E_A \text{ and } E_B] = CF[A] \times CF[B] \quad (2.4.32)$$

Winston's methods only deal with positive certainty factors. His "ad hoc" method defines the hypothesis certainty factor where there is multiple evidence as,

$$CF[H] = \max_{\text{all } i} [CF[H, E'_i]] \quad (2.4.33)$$

Winston's probability based method modifies all three areas of the theory. Antecedent certainty factors are combined by multiplication as shown for evidence E_A and E_B :

$$CF[E] = CF[E_A \text{ and } E_B] = CF[E_A] \times CF[E_B] \quad (2.4.34)$$

The *a priori* hypothesis certainty factor is modified by the evidence certainty factor, through a single valued function mapping “input” certainty to “output” certainty. Multiply supported hypotheses’ certainty factors are obtained using a transformation to a certainty ratio, r . This general relationship is given by,

$$r = \frac{CF}{1 - CF} \quad (2.4.35)$$

Substituting an *a priori* certainty factor for the hypothesis in the latter equation gives an *a priori* certainty ratio, r_o . Likewise using the conditional certainty factors obtained for the actual evidence (E_i) gives

$$r_i = \frac{CF[H, E_i]}{1 - CF[H, E_i]} \quad (2.4.36)$$

Then the hypothesis certainty ratio is given by,

$$r[H] = r_o \times \prod_{\text{all } i} \frac{r_i}{r_o} \quad (2.4.37)$$

By rearranging equation 2.4.35, the hypothesis certainty factor is defined:

$$CF[H] = \frac{r[H]}{1 + r[H]} \quad (2.4.38)$$

These modifications to the original theory are intended to match the calculation characteristics with specific domain characteristics. For example, a domain problem which is more strongly influenced by volume of supportive evidence, rather than the strength of one in particular, would be more suited to a methodology which involved the effect of multiple evidence certainty factors. In such a case, the selection of those methods based on the maximum or minimum functions would not be suitable.

2.4.2.5 Theory of Likelihood Ratios

This method was used in PROSPECTOR, and is based on Bayesian decision theory and fuzzy logic. Uncertainty in evidence is represented by a number between 0 and 1. The strength characterising relationships is modelled by the pair of likelihood ratios (LS, LN):

if E then H (to degree (LS, LN))

LS is the sufficiency factor and measures the support favouring a hypothesis, while LN is the necessity factor and measures the support against a hypothesis. The likelihood ratios are defined as;

$$LS = \frac{p(E|H)}{p(E|\neg H)} \quad (2.4.39)$$

$$LN = \frac{p(\neg E|H)}{p(\neg E|\neg H)} \quad (2.4.40)$$

Given that *a priori*, and *a posteriori* odds are given by;

$$O(H) = \frac{p(H)}{p(\neg H)} \quad (2.4.41)$$

$$O(H|E) = \frac{p(H|E)}{p(\neg H|E)} \quad (2.4.42)$$

respectively, it can be shown, using Bayes' rule, that;

$$O(H|E) = LS \times O(H) \quad (2.4.43)$$

$$O(H|\neg E) = LN \times O(H) \quad (2.4.44)$$

The effect on the odds of H , of observing E is thus clear. As $p(\neg X) = 1 - p(X)$;

$$O(X) = \frac{p(X)}{1 - p(X)} \quad (2.4.45)$$

by which corresponding odds and probabilities can be obtained. Thus when E is observed or $\neg E$ is observed, the probability of H can be updated by using LS or LN .

If, however, the evidence is not established with certainty, because of uncertainty in the observation or because the evidence is the hypothesis of a previously applied rule, then the odds multiplier, λ , must be changed from LS or LN . The effect of uncertain supportive evidence E' on λ is given by a piecewise linear function. Such a function, adapted from Bratko [24], is shown in Figure 2.8.




Figure 2.8: The definition of the odds multiplier λ

In such a case;

$$O(H|E) = \lambda_{E'} \times O(H) \quad (2.4.46)$$

and the probability can thus be obtained from the resultant odds. Where a hypothesis is multiply supported by E_1, E_2, \dots, E_n then

$$O(H|E_1, E_2, \dots, E_n) = \prod_{i=1}^n \lambda_i \times O(H) \quad (2.4.47)$$

where λ_i is the odds multiplier corresponding to evidence E_i .

Where evidence consists of various pieces of evidence logically AND-ed, the minimum function is used as in fuzzy logic theory. That is,

$$p(E) = p(E_A \text{ and } E_B) = \min[p(E_A), p(E_B)] \quad (2.4.48)$$

Again, the maximum function is used when pieces of evidence are OR-ed.

Many expert systems have used this scheme. Questions have been raised over certain assumptions, and systems have performed poorly as a result of overlooking underlying constraints.

2.4.2.6 Unverified Evidence Quantification Method

This is the simplest method of dealing with uncertainty. It is not an attempt to manage inexact reasoning, but merely an attempt to rank hypotheses.

The measure of a hypothesis uncertainty is the number of pieces of evidence on which the hypothesis is based which are unverified. In a consultant expert system the measure would be more narrowly defined in terms of evidence which is unverifiable. Though without mathematical basis nor intrinsic appeal, this simple scheme has been widely used where the theories are not appropriate due, for example, to their computational demands and statistical requirements.

2.4.2.7 Theory of Endorsements

The theory of endorsements is aimed at dealing with reasons for believing or disbelieving a hypothesis, and influencing system inference thereby. These reasons for believing or disbelieving are represented in frame-like knowledge structures called endorsements. A backward chaining strategy is used to reason through the rule base which has endorsements. As endorsements are encountered (in consequents or antecedents) they are used to decide if a line of reasoning is certain enough or whether to reduce uncertainty. Various endorsements give reasons for believing or disbelieving rules, data and conclusions, and give potential reasons for doing a task. Also, resolution endorsements record the application of methods, in accordance with other endorsements, to resolve uncertainty. Since this theory is relatively new, it has seldom been used. There are clear advantages in having the

line of reasoning explicit, as through the theory of endorsements. However, during reasoning, the theory rapidly creates a large body of applicable endorsements.

2.4.3 Some General Points

The various theories and methods all have their problems, and work is continuing in this area. Some generally applicable problems will be touched on.

- Most of the quantitative techniques require the identification of point values for uncertainty measurements, something that can be a great overstatement of knowledge.
- Either in their derivation or for reasons of practicality, assumptions are introduced which often do not apply to real-world problems.
- For long inference chains the use of the uncertainty technique can be very costly in memory and computation time.
- In many domains a large number of possible pieces of evidence and hypotheses exist; vast quantities of data would have to be set by experts for some of the techniques.

Despite these problems, expert systems using such techniques have been found to operate satisfactorily. The selection of the technique most suited to a domain will be dependent on various characteristics of that domain, including number of possible conclusions, volume of data required, quality of statistics available, the effect of uncertain evidence, the mutual effect of multiply supported hypotheses, and the effect of opposing evidence.

2.5 Learning Facilities

Facilities for learning in an expert system would enable the system to update its own knowledge base on the basis of experience. Though desirable, the ideals of such facilities have generally remained unachieved.

Having developed an on-line expert system for diagnosing valve malfunctions, Schreurs and Bednar [18] continued with the development of a companion “refinement module”. The diagnosis expert system maintains a log of suspected malfunctions. After maintenance has been performed, the refinement module compares the actual malfunction disclosed by maintenance, with the log of diagnoses. When discrepancies are found, inductive learning techniques are used with the actual malfunctions and the corresponding log of sensor data, to extract new rules. Consistency between these new rules and the original rule base is checked before the rule base is updated; any conflicts found must be resolved by human experts.

One simple learning strategy is described by Abu-Hanna and Gold [25], and is based on deep/shallow knowledge base interaction. The knowledge in the “deep expert” (ie. the deep knowledge base) consists of a structural and behavioural model of the diagnosed system. The shallow knowledge base is initially empty. Following a successful diagnosis using the deep knowledge, the symptom(evidence)/diagnosis relationship is added to the shallow knowledge base. Such relationships are then available for future use; the “shallow expert” has “learned” the relationship.

By manipulating the numerical values used in the management of uncertainty, like certainty factors, the expert system reasoning can appear to be brought into line with experience. However, the onus of updating expert systems is very much with the system developer or user. Thus, rather than a learning issue, the provision of maintainability techniques is where the emphasis lies. Prerau et

al [26] describe some maintainability techniques realised from the development of a large multi-paradigm expert system, addressing such areas as modularity, style, knowledge and input/output maintainability, system management, data access and control flow.

2.6 Explanation Facilities

There are basically three types of explanation: interpretive, descriptive and reason giving [27]. Interpretive explanations give the meanings of terms, descriptive explanations relate to processes or structures, and reason giving explanations provide details of how conclusions were arrived at. Often all three types are supplied where explanation facilities are important; explanation facilities may be needed to help the user understand, or to justify the system's reasoning to the user in order to evoke confidence in the expert system.

In some systems a separate knowledge base is used to provide explanations. The knowledge explains the expert system's problem-solving knowledge, and is meta-knowledge—ie. knowledge about the knowledge. Causal models are often used to provide explanations. Frequently these are realised as semantic networks, with cause-and-effect, association and definition links incorporated.

For many consultant rule based expert systems, tracing the rules used is considered to give sufficient explanation. MYCIN makes this facility available to the user by providing “why (and why didn't you) and how” questioning. Research has been undertaken to manipulate explanations to suit the specific needs of the user—so called “user-tailored explanations” [28]. This uses the same trace of rules, but removes certain rules and concepts thought too simple. Rules and concepts considered too complex are also removed unless they are thought too important, for which a paragraph of explanatory text is provided. The bases for the manipulation are rule “complexity” and “importance” measures and user “expertise” and “detail” variables.

Some systems use explanation facilities for teaching. Such dedicated systems are called Intelligent Computer-Aided Instruction (ICAI) systems. These systems are competent expert systems since they contain knowledge which directly models techniques used by human experts.

2.7 Advantages of Expert Systems

There are many advantages which can be provided by expert systems technology. Some of these will be specific to the application, but some general points can be listed.

- **Extended computer use:-** Using expert systems provides the opportunity to solve problems which could not be tackled using computers before, or which were not satisfactorily solved using conventional computing techniques. An example of the latter is an expert system used to control cement-kilns, which, in comparison with model based control, saved £4 million per year in energy costs [11].
- **Human expert assistance:-** The expert system can provide a better environment for human experts which can enhance their productivity. For example, the integration of expert system techniques with Computer Aided Design systems has reduced design times by factors up to 90 [29].
- **Flexibility and extensibility:-** Expert systems can be prototyped rapidly and incrementally. Likewise, they can be straight-forwardly modified in line with new (ie. previously unexpected) situations. An increased life cycle can result.
- **Understanding:-** Implementation languages are close to natural language, and thus the represented knowledge can be more easily understood. Moreover, the explanations often provided aid the users' understanding.
- **Modularity:-** It is often easier to modify part of an expert system without affecting the rest than to do so with a conventional program. Modularity and understanding facilitate improved maintainability.

- **Portability:-** When knowledge is removed from a particular procedural context, it is free to be used in many contexts, within the existing application and in others. In addition, the problem-solving strategy may be applied to many different problem areas, as the inference engine is domain independent.
- **Availability:-** Expert systems provide the expertise whenever needed. The technology provides a means of distributing knowledge.
- **Consistency:-** Expert systems are less adversely affected by stressful environments. They can therefore provide more rapid reaction than humans in emergency conditions.
- **Capability:-** Expert systems can cope with uncertainty. Model based expert systems also have the potential for handling some unexpected situations.

2.8 Applications in Power Engineering

Although the power industry was relatively slow in first using expert system techniques, interest rapidly grew and is now very widespread. Most expert systems in power engineering are presently prototypes although some have been in use for some time. For example, the contingency selection expert system described by Wrubel in Schulte et al [30], has been in use since 1980.

Many expert systems have been developed for power engineering applications. A survey of expert systems in electric power systems is presented by Zhang et al [5]. Applications to power engineering in general include the following tasks:

- fault diagnosis of networks and of HVDC systems;
- reactor shutdown analyser;
- monitoring and diagnosis of valves, rotating machinery, large pumps and battery systems;
- alarm processing;
- system restoration and feeder reconfiguration;
- load forecasting;
- reactive power and voltage control;
- switching operations;
- security assessment;
- transient stability;
- unit commitment;
- maintenance scheduling;

- substation automation;
- reactor fuel routing;
- capacitor placement;
- power cable design, auxiliary system design, and substation grounding design and testing;
- protection setting and coordination evaluation;
- operator tutoring and training scenario creation;
- operation and loadflow planning;
- generation and distribution expansion planning;
- friendly interface

Chapter 3

Power Systems & Fault

Diagnosis

3.1 Structure of Power Systems

3.1.1 Network

The purpose of a power system is to produce and provide a supply of electric power of reasonable quality with reasonable security, to meet the load demanded by the customers, at minimum cost. To meet this task, the electricity supply network and industry consist of three divisions: generation, transmission and distribution.

Generation equipment is used in the production of electrical power and is presently centrally located in large power stations, which are dispersed geographically, usually away from load centres. The electrical equipment or plant includes rotating machinery, cabling, busbars, transformers and junctions; a vast amount of other equipment like turbines, boilers and pumps is also required. Various types of switchgear is installed between items of plant. As switching in electrical power networks can dissipate vast amounts of energy, switchgear is classified in terms of the current levels it can break and make (or reclose onto), and consequently the use it is put to. Under manual or automatic initiation, various types of isolator can disconnect dead plant (no current or voltage), energised plant (no current, only system voltage) or load carrying plant (system voltage and current). Circuit breakers can likewise reconfigure the network under manual control, but are designed to be able to break fault (ie. abnormally high) currents.

The transmission system connects the various power stations with each other and with the distribution system. The system can be viewed hierarchically. At the first level, the network consists of many geographically displaced substations connected by interconnector or feeder circuits, mainly overhead lines. Each substation may contain busbars, transformers, autotransformers, junctions, reactors, etc.; this plant connects circuits together and, via transformers, supplies

the distribution system. Circuit breakers and isolators are also located in the substations.

While some large industrial loads are supplied from a transformer directly connected to the transmission system, the majority of customers are supplied by the distribution system. Busbars, transformers, junctions, circuit breakers and isolators are housed in distribution substations, which are interconnected by lines and cables. Due to the dispersion of customers in a rural setting, transformers and switchgear are often also located out on the lines and are pole mounted; fuses also form an integral part of such distribution systems.

3.1.2 Protection Systems

Following the occurrence of a short circuit fault, abnormally high and possibly unbalanced currents flow which can cause fire damage at the fault location, cause damage to plant feeding the fault and can also cause serious problems with the stability of rotating machinery. To minimise these effects, protection equipment must detect abnormal system conditions, and cause circuit breakers to trip. The protection system must function to isolate the least part of the network. Sometimes autorecloser (or automatic circuit recloser) equipment is included with line protection so that lines may only be temporarily isolated and transient faults are cleared but the line is automatically restored.

Protection systems for generation, transmission and distribution systems follow similar principles, but with differing emphases. In generating stations there are extensive mechanical and thermal systems. Thus there is often a large number of protection devices monitoring physical and thermodynamic aspects of plant, like steam pressures. Moreover the process of isolation is often more complex as thermal, mechanical and nuclear processes may have to be safely shut-down in addition to the interruption of the flow of current. Therefore comprehensive

shutdown procedures are instigated by operated protection.

Protection equipment for the transmission and distribution systems is generally located in substations. As the transmission system carries bulk power, the power available to feed a fault (and hence cause damage) is very high and the effect on the power system as a whole is significant. Moreover, generator stability is at risk due to the changes in system conditions. Consequently investment in protection systems is very high. Conversely, nearer the customer in the distribution system, simpler protection equipment is employed. In rural systems and on low voltage distribution systems the simplest form of protection, the fuse, is extensively used.

Protection relays detecting electrical conditions are supplied with a proportion of system quantities by current transformers (CTs) and voltage transformers (VTs). Relay output contacts change status according to the relay operating characteristic, which is set up in order to detect fault conditions. The operation of such relays (primary activity) controls secondary activity. Secondary activity may be circuit breaker tripping or trip relay operation (which will cause further secondary activity). Relays are connected to each other by local wiring, perhaps with auxiliary switches; where they are in different locations, communication is by pilot wire, optical fibre, or signaling superimposed on the power line or cable (power line carrier or mains carrier). In this way the functional requirements of a protection system are achieved.

There are two main philosophies used in protection systems: unit and non-unit protection. Unit protection schemes protect a particular well-defined zone, the zone being defined by CT location. For internal faults (ie. between the CTs) the relay operates, for external faults the relay does not operate. Non-unit protection, however, acquires system conditions at a point (the relaying point) to recognise fault conditions. In order to have the desired level of reliability, protection systems

include multiple main protections and backup protection. Backup relays will cause a break in the flow of fault current in the event of the main protection not doing so. Being sited “upstream” of the backed up main protection, the operation of backup relays will normally result in the isolation of a larger part of the network; coordination, that is setting the backup protection to give the main protection enough time to isolate the fault before operating, is therefore very important.

The means of detecting fault conditions are varied. The differential type of unit protection has CTs, on all phases, surrounding the protected plant and connected in parallel by phases and with the relay detection circuit such that, by Kirchhoff’s current law, there is virtually no current flow in the detection circuit unless another path for current is provided by the occurrence of a fault. The principle is illustrated in Figure 3.1(a). Restricted earth fault protection acts by




Figure 3.1: A simplistic illustration of the principles of (a) unit protection and (b) non-unit protection

the same principle to protect transformer windings. Non-unit protection types (generically as in Figure 3.1(b)) generally either measure a current magnitude or determine the effective impedance seen from the relaying point. Overcurrent protection follows the former principle: each phase’s CT has its own detection circuit which causes the relay to operate when a level of current is exceeded. An inverse current-time characteristic is common. Some relays are directional, operating for faults on one side of the relaying point only. When supplied by

a CT on a transformer neutral or by the sum of the currents in the CTs on each of the three phases, the same relays form non-unit earth fault protections. The various types of distance relay operate on the principle of determining the effective impedance “looking into” the network. Such devices use CT and VT outputs, and various voltage-current characteristics to recognise the occurrence of a fault.

3.1.3 Control and Monitoring Systems

3.1.3.1 Function

The purpose of the control and monitoring system is to automatically change, and permit a human operator to change, the network in line with changing conditions. The changing conditions can be classified as load changes and unexpected events. While future load levels cannot be guaranteed, it is known that demand will change and loads can be forecasted with reasonable accuracy. Unexpected events include system faults and the consequential isolation of plant. Such occurrences can affect three aims of system operation: the active and reactive power generation should match demand to within a reasonable accuracy (load fluctuations make absolute matching impossible); the voltage and current levels on items of plant should not exceed their rating; the delivery of power should be economical.

The first two technical factors are presently the most critical, and their level of compliance indicates the operating state of a power system.

- The normal state is that which satisfies both criteria.
- In the alert state, although both criteria are met, a reasonable contingency could cause the system to enter the emergency state.
- A system in the emergency state is violating operating limits on some plant, while meeting demand.

- When plant has been isolated such that the demand can no longer be met, the system is in the restorative state.

Clearly the aim of system control is to keep or restore the system back to, as near the normal state as possible. Christie and Talukdar [31] state that most power systems [in the USA] usually operate in the alert state.

Automatic control is provided in the generation system by automatic voltage regulators, governors and stabilisers. In the transmission and distribution systems automatic tap changing equipment is used to control secondary transformer voltage levels, and automatic load shedding is performed in response to frequency drops. Static VAR compensators and synchronous compensators are also used to control transmission voltages. The automatic control provided by autoreclosing equipment helps to restore the system to normal conditions.

On a regular basis, system control involves unit commitment, economic dispatch and security assessment. Unit commitment orders generating units on or off the system, in anticipation of the future demand, and works hours to days in advance. Economic dispatch orders the loading of connected generating units to meet the aforementioned three criteria, and has a time-scale of minutes to hours. Every few minutes the security assessment simulates the effect of possible unexpected events to determine if the system is secure or not (ie. whether it is in an unacceptably alert state).

Following a fault on the power system, the operator's duties escalate. He must obtain information from the data generated in order to (at least partially) understand the situation. Decisions have then to be made and perhaps action taken to prevent the situation worsening and to restore things to normal. In all these functions a reliable knowledge of the system configuration and status is required; hence the significance of the monitoring system. The provision of valuable data is especially critical when the rareness of major disturbances is

recognised. As Boutacoff [32] reports: “According to a working group from the IEEE, approximately 90% of the problems which have caused emergency, near-blackout, or blackout conditions were problems that the operators had never encountered before”.

3.1.3.2 Remote System

The above system-wide control functions are exercised from a few control centres. Telecontrol facilities provide the control engineers working in the control centres with the ability to reconfigure the network. Not all control actions can be remotely performed; some control requests must be telephoned to field engineers. The Supervisory Control And Data Acquisition (SCADA) system updates the database containing system topology in line with switchgear activity, and also retrieves other status changes, and voltages and currents from hundreds of measuring points throughout the network. The data is operated on by loadflow programs, load forecast programs and other power application software. Man-machine interfaces (MMIs) are provided for the control engineers to access the software and data.

Alarms or messages are transmitted by the SCADA system following a change of state of protection or switchgear equipment. These are time-tagged, displayed to the operator and printed. Different alarm formats are used by different utilities. Moreover, utilities convey differing levels of detail in their SCADA alarms. The time-tagging is to varying degrees of accuracy, since some utilities tag times near the source, while others only tag times on receipt at the control centre. Some examples of alarms from different utilities are shown (a ScottishPower SCADA alarm is shown first, then one from Wollenberg [33], Arès and Girouard [34], Ashmole [35], Cockburn et al [36] and McDonald et al [37]):

13:23:30:28	COCK2	W20	OPEN				COCK2
ALARM	09/20	14:01:10	S	JENKIN	JES-HO	OP	
890415	161507	A	B300-2	-Open-			
2013	MANN4	AXMI MAIN PROT	OPERATED				
I	04:55:37	HEADTOWN	132KV	OCB	2T0	OPENED	
0924	MILTON	33 BSECT	OPEN				

The system operator can set alarming limits on analogue parameters. Then, if a measured analogue falls outwith the set limits, an appropriate alarm is displayed to the operator. The analogue parameters are retrieved at a much slower rate than the status change points.

3.1.3.3 Local System

Most control and monitoring equipment is located in generating stations and substations. Outstations provide the link to the remote system: to retrieve data, the master station at the control centre polls the outstations. More detailed data of pre-fault, during fault and post-fault conditions is recorded in some transmission substations by disturbance recorders. Most recorder data is retrieved by a site visit, however some digital recorders have been connected to modems which can be remotely dialed up for data transfer.

3.2 Some Developments in Power Systems

There are many developments progressing in power systems throughout the world which will influence the task of controlling power systems. In particular, those developments which influence the fault diagnosis of systems will be briefly presented.

3.2.1 Network

Throughout the world, power systems are growing. Networks have been developed to a high level of interconnectivity. Moreover, and more significantly, utilities have been establishing high capacity interconnectors between formerly isolated networks, and networks which had low capacity connections. The use of HVDC technology has been significant in this. Links within countries and across borders have been developed, the trans-national interconnectors forming new trading routes. As a result, transmission network conditions are more critical to more people, and system conditions are more variable as supply and demand patterns differ from the traditional trends established when only utility customers were supplied.

While transmission network development has proceeded, expansion of generation systems has been less significant. Primarily it has been environmental arguments and legislation which have stayed development. Consequently margins of capacity over demand have narrowed.

According to Roman and Marian [38], the major emphasis of “year 2000 [generation] power plant” will be operation and maintenance, with availability and capacity factors becoming more important. They predict extensive use of computerised instrumentation, expert system analysis packages and robotic/remotely operated systems.

A major development in the United Kingdom electricity supply industry has

been the privatisation of the utilities, and the progressive removal of restrictions on private generators' access to the network. It is expected that this will lead to the development of many small private generating stations (probably factory waste heat units, or gas turbine units which can be rapidly installed). The addition of non-utility generation to a power system will likely complicate the task of coherently operating the system.

SF_6 switchgear, and gas insulated switchgear equipment generally, have become popular at all levels of power systems. Such equipment is very efficient and effective, and more reliable as it is unaffected by humidity fluctuations and does not have a fire risk. The units are sealed for life, and require reduced maintenance.

3.2.2 Protection

There has been much activity in the development of digital relays in recent years. Whereas the original electronic relays lacked reliability due to the number of electronic components, microprocessor based devices overcame this problem. Such devices have been used to realise the same operational characteristics as analogue relaying equipment. With the advent of 32-bit microprocessors and of digital signal processing chips, the theoretical performance levels are attainable [39]. Having few moving parts, these devices require less regular maintenance. Digital relays, unlike their analogue counterparts, use the same components/circuits under fault and normal system conditions. This feature permits the self-diagnosis of relays. This ability is a major advantage, providing data on device performance problems before the device fails to isolate a fault. The operator will thus have more data available to him concerning the operational state of protection. As such equipment is becoming widely used, the cost is expected to decrease. Such equipment shall then become competitive at distribution levels as well as at transmission levels.

In addition to transposing traditional relaying algorithms from analogue devices, more sophisticated algorithms have been developed. In addition, multi-function relaying units have been developed incorporating many protection schemes for particular applications. Some digital relays also have the facility for the limited storage of phasors. Some relaying algorithms involve the identification of the location of the fault on a line or cable. Microprocessor relays could also store this and make it available to other electronic equipment in the substation. Significant amounts of fault details would thus be provided without additional hardware and without interfering with the protection's main task.

Advances in integrated protection and control systems will permit more advanced backup schemes (facilitated by fast communications), the use of adaptive protection responsive to system conditions, and the simultaneous use of any manufacturer's equipment due to the standardisation of protocols. The interpretation of protection system behaviour could thus become more complex.

3.2.3 Local Control and Monitoring

The aim in local control and monitoring system developments has been integrated protection and control systems. In conventional substations, protective relaying, disturbance recording, metering, automatic control and SCADA functions are performed by stand-alone separate systems. Due to the critical importance of protection integrity, the separation of protection functions has been strictly observed. Integrated protection and control systems maintain functional independence with a degree of integration. Local area networks (LANs) and high quality optical fibre communications links, integrate digital protection, digital recorders and enhanced remote terminal units (RTUs, the interface to the SCADA system). As well as the aforementioned protection advances, such an arrangement will allow: shared access to local databases which will increase data volume and

facilitate more local diagnosis; shared access to local transducer/data acquisition which will improve management of redundancy.

Once these developments have been more fully proved, it is expected that the use of non-power-industry-specific electronics will reduce the presently restrictive costs.

Some transmission substations have had modems installed which are connected to digital disturbance recorders. These provide the means for down loading the recorder data, thus retrieving valuable information much faster.

Modern RTUs have the flexibility to incorporate a modem card which provides for the remote collection of sequence of event logs via the public telephone network using a PC and independent of real-time computer operations [40, 41]. Modern RTUs are also multiported. Thus, without the delays of master to master communications and without requiring duplicate RTUs, substation data can be simultaneously provided to utility national and regional control centres and to adjacent utilities. Both of these developments in RTU technology will provide more data than was previously possible.

These RTUs also have the flexibility to be used as satellite masters. Micro-processor based control modules have much lower power demands, and can thus be used as pole mounted devices in distribution systems. As a result advances have been made in the development and use of distribution automation. Radio has been used as the communication medium from the pole mounted device in SEEBBoard's system [42]. With a receiver at the substation, the enhanced RTU could be used to transmit both substation and feeder data to the SCADA master.

The application of proved technology is ongoing. By the nature of power systems, being required continuously, the upgrading of equipment has to fit in with limited planned outages. Consequently the complete implementation of new equipment in the transmission systems, as well as being costly, is a relatively slow

process. Knowledge of both old and new systems is needed for power system control. In addition, the data availability and the control possibilities in distribution systems are growing.

3.2.4 Remote Control and Monitoring

Much development activity is being carried out in power system control centres. Utilities have set about concentrating control activity in fewer centres. The aforementioned extensions of distribution telecontrol have resulted in greater remote control capabilities. The expanded SCADA systems provide more monitored data at the control centres.

With these changes in functionality, a major change in computer system architecture has been adopted. Distributed open system architectures are being used in upgraded or new control centres. Computing functions (man-machine interface, power application software, SCADA, etc.) are distributed among processors with shared facilities and communicating via high performance channels or LANs. Planned developments involve the use of workstations, which are connected via LANs.

The developed control centres include much improved man-machine interfaces (MMIs), with high quality graphics screens, standard window systems and standard keyboard and mouse driven applications. This provides a powerful means of getting information to operators and the means to control the system efficiently.

These MMIs display network diagrams to varying degrees of detail which are automatically updated in line with received SCADA activity. Moreover these diagrams are interactive and provide the means of issuing telecontrol commands. In addition, Geographical Information Systems (GISs) are being included, by which operators can see geographical network diagrams. Such functions clearly require extensive system databases, and these are held in shared facilities on

the LAN. Comprehensive relational databases provide the data for the various network manipulation and analysis programs. A language oriented database [43] and a knowledge based database (implemented in Turbo Prolog) [44] have been developed and tested. Moreover the integration of data sources within utilities is being developed: corporate databases would hold the source data for on-line and off-line, operational, planning and design uses. Thus system data is now much more readily available.

The computer systems are being developed to incorporate electronic documentation. Permanent documentation such as codes of practice and regulations will be available for inspection on the operator's display. The means for helping the preparation of documentation generated on-line will be incorporated. Aids for the production of fault reports, safety documentation and switching schedules will be provided.

The developments in remote control and monitoring equipment will permit the receipt of more data and the issuing of more commands, and will improve the display of system configuration and activity. However the task of interpreting the system behaviour and condition from the received data, and identifying required operator intervention remains a difficult task at times. A tool to help the operator with system fault diagnosis would be a valuable asset.

3.3 Fault Diagnosis of Power Systems

3.3.1 Events

The fault diagnosis of power systems is primarily concerned with identifying faults on power system equipment. Due to the scale and geographical dispersion of power system plant, the development of a dedicated system for monitoring every piece of plant and reporting back within a reasonable time period is not feasible. However, the protection system fulfills such a function, and the SCADA system monitors protection and switchgear activity. Thus, the monitoring of “secondary effects” is used to diagnose power system faults. This, however, complicates fault diagnosis as the evidence is influenced by the behaviour of power plant, protection and switchgear, and SCADA system equipment. Fault diagnosis thus includes the identification of system faults, protection maloperations, switchgear maloperations and telemetry malfunctions. Protection and switchgear maloperations include failure to operate (ie. false non-operation), false operation, and slow operation. Telemetry malfunctioning relates to lost alarms. These aspects are unpredictable, widespread, very fast, and can lead to complex changes in the network. Fault diagnosis is by no means a trivial task.

3.3.2 Reasons for Fault Diagnosis

With the task of fault diagnosis sometimes being very involved, it must be clearly justified. The reasons for performing fault diagnoses may be looked at in three time frames: immediate, short term and long term.

Immediately following a disturbance and its isolation, a fault diagnosis would identify the plant in which faults occurred. Only then can appropriate restoration procedures be adopted. Different items of plant are not always connected through circuit breakers, with the result that more than the faulted plant may

be isolated. Failures of protection or switchgear equipment can also result in excessive areas of the network being isolated. This loss of plant can necessitate or cause the loss of supply to customers, and hence result in lost revenue and even financial penalties for loss of supply. Fault diagnosis can enable operators to identify the minimum necessary isolated area; appropriate telecontrol actions can then be generated to restore plant and load. Together with the help of a Geographic Information System (GIS), investigation teams can then be more effectively guided to the location and nature of the fault. The fault diagnosis, in keeping the operator informed, will also help in the reassessment of plans (eg. planned outages, dispatch).

With regard to short-term applications, the fault diagnosis will identify protection and switchgear maloperations. Such findings will indicate that detailed investigation of the protection system and maybe switchgear is required; alternatively unnecessary investigations can be avoided. The performance of protection (and switchgear) is critical to the operation of power systems. In addition, on performing a fault diagnosis and identifying a fault, the remote retrieval of relevant disturbance recorder data or microprocessor based relay data could be initiated and details such as the fault location along a line sent to the investigation team. The post-fault investigation could thus be speeded up.

If properly maintained, records of fault diagnoses would provide a valuable source of data for assessing the long term performance of plant, protection types and switchgear. This is a long term investment and must be done conscientiously as faults and maloperations are not very common.

3.3.3 Basis for Diagnosis

As was mentioned before, by monitoring the behaviour of protection and switchgear activity, inference can be made concerning the behaviour of power plant. As

switchgear and some protection equipment is wired into utility SCADA systems, their behaviour is known at control centres; protection and switchgear activity is therefore the observable evidence which is caused by and which indicates the unobservable system faults.

The sample point measurements of voltage and current supplied by the SCADA system are not of direct relevance to fault diagnosis. This is because the significant polling period is unlikely to capture the system parameters during the short period during which the fault current is flowing. However these measurements can provide a means of verifying breaker status following isolation of the fault.

Another source of data is disturbance recorders. Valuable information can be obtained from fault recorders [45]. Digital recorders can generate fault summaries which can be transmitted by the SCADA system; the full complement of data is too voluminous for the SCADA system, but can be gathered via a modem. The complete recorder data can thus aid the diagnosis, but not until the period of time required to transmit the high volume (especially analogue) data has elapsed. In the meantime the summary data provides useful, though limited, information. In addition, microprocessor based relays can provide current magnitudes measured at set time instants before they operated and sometimes the fault location may also have been identified.

3.3.4 Reasons for a Fault Diagnosis System

Traditionally, fault diagnosis has been performed by the operator in the control room, and by protection engineers much after the event. However, utilities are seeking to develop tools to help their operators diagnose faults because of the changing nature of the task and the advantages of such aids. Even back in 1969, Dy Liacco and Kraynak recognised some of the advantages of having an operator aid for analysing system trouble [2].

Utilities have found it difficult to strike a balance between supplying more SCADA information to increase knowledge of system conditions especially following disturbances, and restricting the generation of alarms in order to avoid swamping operators with unnecessary information. The dissatisfaction of operators came out in a recent survey of utilities [46]. Whereas 65 out of 87 respondents declared the operators to be satisfied with alarms presented to them during normal conditions, this figure fell to 21 out of 87 when related to emergency conditions. Moreover among the alarms viewed as a nuisance rather than being helpful were relay alarms. A fault diagnosis system could increase knowledge of system conditions, but not require the operator to be presented with excessive alarms. A closely related, but shallower, analysis which is primarily concerned with converting raw alarms into a more digestible form, is called alarm processing or alarm handling.

More alarms are now generated as a result of the developments described in section 3.2, and the computer systems performing SCADA functions have been upgraded. The need for operator aids has become more critical. Due to the interconnection of different power systems, conditions of systems other than the operator's own network influence system operation: adjoining system data must also be monitored. More equipments are being wired into SCADA systems. Possible message retrieval is now at a rate above that which operators are able to interpret the messages. In many utilities, and now (with privatisation) in Britain also, the bulk power network operators are obliged to balance their utilities generation with private generation levels. All these factors add to the problem of more real-time data and hence more variables in decision making. Together with the extension of telecontrol, this means that operators have more possible courses of corrective action. A cognitive barrier may be reached during disturbances, at which the operator's interpretation of system behaviour becomes arrested. This is

also a result of the increasing complexity of substation configurations and highly interconnected networks. A system with the required diagnosis ability, and with consistent behaviour, is therefore required to give the advantages described in section 3.3.2.

Due to the general slow-down in large scale generation system development, the operational policy in many utilities has been to have reduced reserves. As a direct result of this, there is greater demand for rapid fault location and restoration of unnecessarily isolated plant. The greater utilisation of existing plant is stressed due to the environmental constraints many utilities have had imposed. A fault diagnosis system would help the operator better manage the utility's available resources by rapidly identifying faults and thus the minimum part of the network which need remain isolated.

Major disturbances are often the result of relatively minor faults following on from other faults which reduced system security. Automatic control and protection fulfills a role in maintaining system voltage and frequency, eg. by load shedding. But these actions are often initiated too late to arrest development into a major disturbance. A fault diagnosis system could help in the rapid recognition of the vulnerability of the system, and hence permit pre-emptive control. (Consider the New York blackout of 1977 which illustrates the need for such a system [47]: a line fault, followed by another line fault later, caused the frequency to drop and load shedding to commence, but by then it was too late as generators were then lost to overspeeding, following which more load shedding was initiated—ultimately the system was blacked out.)

A fault diagnosis system could also give added advantages to the function of a control centre. The system could be used to automatically generate fault reports of consistent quality; one of the problems in existing control centres is that, following major disturbances for which accurate reports are especially de-

sirable, the operator is too busy to generate good reports. Comments could be added by the operator, but the majority of the report material would be automatically generated. More time would then be available for the development and implementation of restoration procedures.

A fault diagnosis system could play an important role in statistics gathering. Having identified a particular system fault, correct operation, wrong operation or wrong non-operation, the appropriate entry in a statistics database could be upgraded. The integrity of such a statistics gathering function would produce more accurate statistics which would be viewed with less skepticism than are the statistics presently published. These statistics could then prove valuable in the identification of type faults, for example. Utilities are now requesting automatic statistics gathering in their energy management systems (EMSs). Fault and maloperation statistics would provide a useful addition to straightforward operation counts.

3.3.5 An Expert System

The provision of expert systems is one way of improving the operation of a power system. Computer programs for numerical computation have been widely used to great advantage in comprehensive EMSs (eg. loadflow, load forecasting, state estimation, etc.). However there is much judgemental activity which could gain from the use of expert systems both in control centres and in utilities' other functions. And not only as an operational aid. In a paper describing problems encountered in power system restoration [13], Adibi and Kafka say: "The knowledge-based, expert systems approach can be considered as a goal not only for system restoration, but for eventual automatic operation of the power system." This is indicative of the maturity of the techniques, and the continuing interest in expert systems for power system operation.

The general benefits of expert systems, described in chapter 2, are available for exploitation. Naser [48] points out some of the benefits for the electricity supply industry such as: increased productivity; increased efficiency; improved quality; enhanced safety; improved consistency; improved objectivity; reduced costs; improved methods for capturing, packaging and distributing corporate expertise.

3.3.5.1 Implementation of Expert Systems in Energy Management Systems

Because of the difference in nature of expert systems from other application software, the implementation of expert systems in the electricity supply industry is of significance. As has been already described, expert systems are usually developed and realised in symbolic languages, shells and development tools. EMSs do not have these in-built. To implement expert systems there are therefore three main alternatives [49]: installation in an EMS computer with the symbolic language/tool/shell system, translated to C/Pascal/FORTRAN and installed in an EMS computer, or installed on a separate processor.

An expert system installed on an EMS computer has the advantage of direct access to databases, man-machine interface, other application software and communications software. However, direct installation is disadvantageous since symbolic reasoning systems can use significant memory, and the EMS computer would be excessively loaded by running, development and maintenance [31]. Use of a language which is already supported, like C, Pascal or FORTRAN, will often involve translation as these are seldom used for development.

Distributed implementation involves installation of the expert system on a separate processor. In a distributed arrangement, only the processor need support the development tools/languages/shells, the expert system is not as dependent on

EMS hardware nor as demanding on EMS processing time, and there is potential for including many expert systems. A distributed implementation can be accomplished by point-to-point or network communications. A network is beneficial in that more expert systems can be added on separate processors without further changes to EMS hardware or software. Moreover, network communications are supported by most workstations which support expert systems.

The developments in utility computer systems have made the industry much more suitable for expert system implementation. Communication networks are being used by utilities and several new control centre computer systems have distributed architectures (some based on LANs). Moreover, the establishment of widely accessible corporate databases mean that the necessary data can be more easily obtained.

3.3.5.2 Fault Diagnosis

Expert system techniques could provide the means of developing a system for the fault diagnosis of power system networks. Due to the knowledge intensive and symbolic nature of the fault diagnosis task, the development of a knowledge based system or expert system should be a suitable means of achieving the benefits described previously.

This thesis reports work undertaken to seek to produce an expert system for the fault diagnosis of power networks. The expert system was intended to operate in the background, requiring minimal operator input, and providing helpful information relating to faults when they appear.

Work was originally concentrated on a generation system. The distinguishing complications in generation system fault diagnosis are the addition of many mechanical as well as electrical trips, and the comprehensive tripping logic which makes identification of maloperations following a shutdown difficult.

However attention was shifted to transmission networks. Transmission networks have complex topology, are very extensive and the time for responsive control is normally very short, making rapid fault diagnosis very important.

Distribution networks often cover inaccessible areas (cables in built up areas, lines in rural communities) making identification of faulted plant more important for the investigation team; as distribution SCADA systems generally provide less specific information there is less scope for detailed diagnosis.

There is therefore scope for the development and use of expert systems for fault diagnosis in generation, transmission and distribution networks.

3.4 Review of Other Expert Systems for Network Fault Diagnosis

The fault diagnosis of power system networks has received a significant amount of interest in recent years, the aim being to develop expert systems which would be useful in the control centre environment or in normal post fault utility activities. Some published work will be looked at.

3.4.1 Automatic System Trouble Analysis

Automatic System Trouble Analysis (ASTA) was presented by Dy Liacco and Kraynak [2] as a procedure for identifying the general state of a network and diagnosing “trouble”. Areas where faults occurred, protection and switchgear failures and false operations are identified as the SCADA alarms are received. Although the authors make no mention of artificial intelligence (their work was published before any significant interest in expert systems), the expression of the analysis as discrete logical “equations”, qualifies this system for being classed as the fore-runner of expert systems for power system network fault diagnosis.

The network is represented as a collection of circuits, each circuit being the smallest element which can be automatically isolated from the rest of the network. The breakers that should trip for a fault on each circuit are defined together with their associated primary and backup relays. The analysis is based on the received switchgear and protection alarms, with voltage indications being used to establish the status of generating stations and certain customer installations.

Operating in real-time, ASTA is triggered by a primary or backup relay operation or breaker trip alarm, but does not commence analysis until the system is no longer changing. Subsequently, faulted circuits are identified as those circuits which are disconnected from all live sources by tripped breakers and have at least

one related main protection relay which has operated. Related protection relays are relays which control circuit breakers in or bordering the disconnected circuits, and which can detect faults on the isolated circuits. Protection malfunctions are identified on the basis of the faulted circuits and relay alarms. The above analyses are performed using at first the final network configuration, and then the history of alarms. Then false circuit breaker and protection relay activity is identified. Basically, diagnosis is performed by a pattern matching approach. The authors also recognised the need for special logic to deal with specific problems; they include the difference between ground and phase fault protection for some circuits. The described analysis equations could have produced a valuable operator aid.

Although there is no evidence of ASTA actually being realised, the establishment of these analysis principles and the identification of the potential advantages, given the date of publication, deserves recognition. The authors themselves recognised some of the short comings of the proposed system. Some problems include:

- The system does not identify the faulted plant; only areas of faulted circuits are identified. These areas may contain several items of plant.
- The system requires the perfect receipt of all breaker trip alarms. For example if, following a fault, a breaker trip alarm was lost (or corrupted), then false relay operations would be wrongly inferred. The integrity of the diagnosis is therefore dependent on the integrity of the SCADA system.
- For a large network, the definition of the breakers that should trip for a fault on every circuit requires the manual setting up of a very large look-up table.

3.4.2 An Expert System Presented By Fukui and Kawakami

An expert system was developed in Prolog for fault diagnosis in electrical power networks [50]. Most likely faulted plant, and switchgear and protection failures and false operation are identified in an off-line role. Network topology is stored as a connection of nodes, a node being equivalent to ASTA's "circuit". Relaying scheme knowledge contains sufficient details of system relays to describe their characteristics in terms of one of five "reduced relay" types. A reduced relay is a simplified view of a relaying scheme. Non-unit relays' protected areas are automatically derived given the controlled circuit breaker position; the location of faults protected by unit relays is included in the definition of the relays.

With the above knowledge the rules for diagnosis use the post fault network configuration: a list of tripped circuit breakers and a list of operated relays is entered by the user of the expert system. Diagnosis of a single fault is first attempted, and, if that fails, multiple faults are considered. The preliminary identification of fault locations is according to the principles of ASTA, namely that the fault must lie in an isolated (blackout) area of the network. Possible fault positions are then identified using related relays' protected areas: faults are located in the intersection of related relays' protected areas. Where more than one blackout area exists in the network, multiple faults are assumed and each one diagnosed on the above basis. In addition, where two blackout areas are separated by only one open breaker, the possibility of one blackout area not containing a fault—being isolated falsely, or as a result of backup protection operating before main protection—is also included. For each possible fault, the unoperated primary relays which should have detected the fault are classified as failed relays; untripped circuit breakers whose relays have operated are classed as failures also (à la ASTA). The "most possible" fault locations are identified from

the alternatives and are those which include the least number of failed and false operations.

As a post fault diagnosis aid, this system is quite useful, and it also proved that a system could be realised along the lines given by ASTA using expert system techniques. It also demonstrates the suitability of Prolog. There are, however, certain shortcomings. In common with ASTA, this expert system is dependent on receiving perfect switchgear information. Also:

- The diagnosis inference has no capability for explanation of the development of the disturbance. The authors describe how, by the estimation of the sequence of activity using basic rules, an explanation facility could be added.
- The single level nodal representation of the network means that, for a large system, very ineffective searches through every relay in the whole network may result. The nodal representation also appears to overlook the fact that, for example, busbars may be reconfigured by isolators. Indeed the only switchgear included in the network representation was circuit breakers.
- Although there is a degree of relay modelling, unit protections still require explicit definition of their protected items of plant.
- Finally, the output of the expert system, as illustrated in the paper cited, is not in a suitable form. Although this might have been a decision aimed at minimising diagnosis time—the diagnosis is very fast—the importance of the user interface for ultimate acceptance of the expert system has often been described [51].

3.4.3 An Expert System Described by Wollenberg and Sakaguchi

Wollenberg and Sakaguchi describe a fault diagnosis system to illustrate some of the general points they make in their paper [49]. While the resources for realising this diagnosis system in FORTRAN were evaluated, the diagnosis system was developed as an expert system using OPS5 in, it is estimated, half the time. Although little detail is given of the implemented expert system, the diagnosis philosophy adopted is the same as that presented by Fukui and Kawakami [50]. Therefore, it is likewise dependent on knowing the correct status of every circuit breaker. In addition, there is no mention of ranking possible solutions.

3.4.4 Distributed Diagnosis

Distributed Diagnosis (D2) was developed by Cardozo and Talukdar [52]. Its predecessor was the Operator's Assistant [53, 54] which consisted of a black-board connecting various problem-solvers. The expert system uses on-line circuit breaker alarms in order to identify faulted lines and busbars, temporary faults, breaker failures, failures of timing relays and blocking signals, telemetry malfunctions, and line overloads and out-of-tolerance voltages. In order to diagnose such occurrences in a real-time setting, a distributed problem solving approach was adopted. A hybrid of non-algorithmic and algorithmic programs was developed, OPS5, Lisp and C being used to realise various programs. A distributed problem solving kernel supports D2, and allows programs on various processors to run in parallel; communication is mostly by frame-like structures using distributed shared memory. A "manager", "proposal crew", "construction crew" and "demolition crew" make up the team of cooperating programs.

A nodal representation of the network is again used. The Operator's Assistant had a significant amount of protection knowledge, and had access to a discrete

event simulator which could be used when unsure of a hypothesis. A deliberate attempt was made to localise the diagnosis by use of a proposal crew.

The diagnosis attempts to identify hypotheses which could account for the observed activity (some telemetry malfunctions can be tolerated). The Operator's Assistant used a pattern matching approach to identify valid hypotheses from those it generated. In order to cope with lost telemetry, D2 uses a synthesis approach instead. The distributed programs effectively grow hypotheses by incrementally adding faults and maloperations in order to cover the received alarms. The created hypotheses are ranked according to descending number of equipment maloperations and telemetry malfunctions. Some shortcomings can be recognised:

- The primary disadvantage of D2 is its costly usage of resources. In response to the discussor's query, the authors revealed that, for a line fault plus false operation example, running time per hypothesis was: 75 seconds with 3 Microvaxes, 110 seconds with 2 Microvaxes, while one Microvax failed to perform the diagnosis.
- The ranking method used in the Operator's Assistant is also used in D2, and is the same as adopted by Fukui and Kawakami. However, with the added ability to diagnose telemetry malfunctions, D2 can generate more possible conclusions, conclusions which could have the same total number of maloperations plus telemetry malfunctions, but in different combinations. The ranking system could thus fail to discriminate between diagnoses which are known to differ greatly in likelihood. One result of this is that, in order to be reasonably confident of including the correct hypothesis, the operator would need to have displayed many unnecessary hypotheses.
- While the Operator's Assistant used the pre-fault and post-fault network configuration, D2 is intended to read in alarms from the SCADA system.

However the authors do not indicate what means they use to trigger the diagnosis. The control of the initiation of the expert system is an important aspect for an on-line diagnosis system. The expert system's role must be as a tool for utility staff without placing an addition burden on the engineers—it must not require input from an operator to start its diagnosis. Conversely, to attempt to invoke, say, the construction crew rules after every primary alarm would be wasteful of resources. The “manager” could probably fulfill such a triggering role, but to omit to mention this is to fail to recognise its importance.

3.4.5 Fault Analysis Using Simulation and Testing

Fault Analysis Using Simulation and Testing (FAUST) is an expert system developed using the Poplog language environment (incorporating POP11, Prolog and an object oriented programming package called Flavours)[55, 36]. Its purpose is to diagnose line and transformer faults, temporary faults and stuck breakers on 132kV and 33kV distribution networks, and uses mainly switchgear alarms supplied via a filtering program developed in C. (FAUST is also designed to cope with planned outages.) An object-oriented representation of the network includes information slots of switchgear state, voltage level and normal direction of current flow in the appropriate objects. Limited protection details are included. A software blackboard system (not separate processors) includes a new hypothesis generator, a simulator, a monitor and a reviewer.

Alarms are compared with any existing hypotheses' expected alarms, which are then updated accordingly. Also, on receiving a breaker trip alarm, the adjacent circuit breakers are identified. If a sufficient number of these are also open, then FAUST generates hypotheses based on faults on the lines and transformers in the area bounded by these breakers. These faults are simulated by propagating

along the path between hypothesised fault location and the above breakers.

A unique answer, a “best guess” or a “joint best guess” is displayed for each isolated area. When no more alarms are expected, or when a few minutes have passed, FAUST identifies the hypotheses with the fewest number of expected alarms not arrived. Hypotheses accounting for more alarms “take preference” over others. (Cockburn, Brailsford and Beaton [36] do not make clear whether the identification of the lowest number of expected alarms not arrived or the highest number of alarms accounted for, is applied first.)

For what FAUST was designed to do, it performs fairly well. Many of the issues for on-line application have been addressed. Yet, some potential problems can be identified. These include:

- Limiting the diagnosis of failures to breaker failure has permitted the use of a simple method of selecting the favourite hypotheses. Although this method might appear to be an overstatement of confidence, it will work given these limitations. For example, despite failures not directly influencing the selection method, the extra alarms generated by backup activity affect the relative confidence in the hypothesis. This works for breaker failures as there is almost always extra breaker trips. But it could not be relied upon for protection failure (even if protection relays were also wired into the SCADA system) as non-unit protection relays can operate as main and as backup protection.
- The division of the selection of the “best guess” into a two tier process can result in errors of judgement. If the first check to be performed is that of the alarms accounted for, then a hypothesis with many still expected alarms may wrongly appear as the best guess because, for example, it incorrectly accounts for a false alarm. Conversely, there may be problems if the check for the minimum expected alarms is performed first: the correct hypothesis

may not be selected as best guess because one of its expected alarms failed to arrive due to a telemetry malfunction, while an unlikely hypothesis has no expected alarms. These can be illustrated by minor changes to the example given by Bramer et al [55].

- The simulation only checks the topology remote from the potential blackout area under certain circumstances; generally plant is assumed live. This is not always the case; wrongly expected activity may thus appear in some hypotheses.

3.4.6 An Expert System Developed by Girgis and Johns

A hybrid expert system was developed by Girgis and Johns [56] in order to locate a faulted busbar, line or common area between a bus and a line, identify breaker failures, classify power network faults, and, when the faulted plant is identified as a line, determine the physical location of the fault by automatically selecting and using an appropriate fault location algorithm. The expert system is mainly written in Prolog, with the algorithms in FORTRAN. A nodal representation of the network is used.

In order to determine the fault location the tripped breakers are used. Like ASTA, the basis is the identification of isolated portions of the network. Breaker (or protection) failures are diagnosed if the breaker bounding the suspected fault location is untripped, but adjacent breakers have tripped. A “depth-first” first answer approach is taken: the expert system reasons with one attempted answer at a time, and stops when one solution is found.

In anticipation of the transmission of voltage and current phasors from microprocessor based relays to control centres, the expert system was designed to consult a database of phasors in order to classify the suspected fault. The technique adopted is also described by Watanabe and Ooba [57], who describe it as a

scalar comparison method. Firstly, the three phase and neutral current phasors which relate to the network around the suspected faulted plant are retrieved. The change from pre-fault to fault magnitudes are determined and normalised by division by the biggest change. By comparison with threshold values, each normalised current change has a binary state ascribed. A pattern matching approach relates these four states to a fault type (eg. single line to ground fault on phase A, etc.). The fault type classified, together with the type and number of phasors available, are the bases on which line fault location algorithms are selected.

Transmission line fault location algorithms have been realised in FORTRAN. Two algorithms exist which are based on the phasors at one end. Fundamentally, they both use selected voltage and current pairs to compute an apparent impedance, but one of them requires knowledge of the source impedance. Algorithms are also provided for the cases where the phasors at the line ends are all available; some require the phasors to be synchronised.

This expert system illustrates how faults may be classified and located accurately by using current and voltage phasors which would become available in control centres in the not-too-distant future. However, the identification of faulted plant is dependent on perfect breaker status information, like ASTA. Similar to the approach described by Wollenberg and Sakaguchi, Girgis and Johns' expert system makes no attempt to discriminate between competing diagnoses. Other inadequacies can be recognised:

- Only a single fault and one maloperation per diagnosis are considered. These two strictures curtail the diagnosis ability of the expert system.
- Relay information is underused. The rules presume that a relay and its controlled breaker are in the same state. Despite the expert system having access to phasor quantities obtained from relays, protection appears to play a very minor role in the identification of faulted plant, its presence being

apparently only assumed.

These shortcomings are especially unfortunate for this expert system since it has access to current and voltage phasors. The phasors, and the numerical algorithms provided, could be used to resolve (to some degree) the aforementioned shortcomings and make much deeper inference than would be possible with only alarms.

3.4.7 Diagnosis of System Faults

Kirschen et al describe an expert system for the diagnosis of system faults (DSF) [58]. DSF uses circuit breaker and relay activity alarms on-line to identify faulted plant, and protection and switchgear maloperations. But, in recognition of the fact that not all system relays will be wired into the SCADA system, DSF is not dependent on relay alarms. A hypothesis and test strategy was adopted.

DSF starts operating when a relay operation alarm is received, but does not commence diagnosis until alarms cease (à la ASTA). Subsequently, blackout areas are determined; a multiple fault is supposed if there is more than one blackout area. Hypotheses of faults on every item of plant in a blackout area are generated. Impossible fault locations (eg. due to maintenance outages) have their hypotheses removed. The remaining hypotheses are tested using the telemetered relay operations, and have implied relay and breaker maloperations added. The results are used to rank hypotheses on the basis of certainty factors.

Although DSF does not require all system relays to be wired into the SCADA system, its diagnosis, like ASTA's, is dependent on the perfect receipt of all breaker alarms. One thing which is problematic for developers of on-line fault diagnosis expert systems, is deciding when to commence diagnosis. This is another area where DSF could be lacking:

- The DSF designers have used the ASTA technique of making the expert

system wait for operations to cease, the reason being that a fault normally causes many operations to occur within a short period. During this period, the messages would be logged in. However more major disturbances generate messages for quite some time. Meanwhile, DSF would be stalled, awaiting the cessation of activity while logging the messages. This may be desirable in traditional control centres for which ASTA was developed, where as much mainframe processor time would be required for SCADA functions. However in a modern control centre, the fault diagnosis would be more likely to be performed on a separate workstation, or at least not on the processor used for SCADA functions. The stalled DSF would offer absolutely no help to the operator, who, in such circumstances, is most in need of help.

3.4.8 An Expert System Developed by Jeyasurya et al

The advantages of implementing an alarm processing expert system in an RTU came out in the discussion between Dr. Mahjouri and B.F. Wollenberg [33]. The advantages of analysis at the substation include: faster reasoning achieved by parallelisation, complicated communications requirements are removed as only the reasoning results need be sent back to the operator, and more data from relays is readily available. Wollenberg pointed out that the main retardation was the insufficient level of RTU hardware. This is now changing [40].

Jeyasurya et al authored a paper [59] describing an expert system for implementation in the substation computer of an integrated control and protection system. Information on relay operations and breaker trips is used to identify single fault location and failed protection and switchgear. The expert system was developed in OPS83, and uses the forward chaining strategy.

The diagnosis determines the protected plant of each of the operated relays,

and finds the items of plant common to all the relays. These are the possible fault locations. Where local information is inadequate for a diagnosis, the authors suggest that the station computer could obtain other substation data via the inter-substation communications provided by the integrated control and protection system.

Diagnosis at the substation is simplified since the reliability of the telemetry system does not influence the diagnosis. Still, there are some inadequacies in the described system which should be noted. Clearly, this expert system is for on-line application, and yet the authors, as those describing D2, do not mention a triggering rule. Also, like the system described by Wollenberg and Sakaguchi, there has not been any means of ranking competing possibilities incorporated. Although it is less of a problem without telemetry malfunctions, the authors did recognise this omission and suggested special rules or counting failures. This expert system, like that described by Girgis and Johns, deals only with single faults. At the local level, the likelihood of multiple simultaneous faults is less than when considering the whole network from the control centre. However, it is feasible for two lines connected to the same substation to fault (due to lightning or cascade overloading), or for one fault to cause instability in nearby protection. A local fault diagnosis system should therefore be able to diagnose these events. One further point should be made, relating to this system uniquely:

- It would appear that the inter-substation communications used to retrieve remote information to help the local diagnosis, is the same communications link which, for example, line differential protection would use. If this is the case, special rules would need to be incorporated to deal with communications failure; as the monitored and monitoring system are both dependent on the same link, care must be taken to make inference from communications failure and not simply to fail.

3.4.9 An Expert System Described by Hotta et al

Hotta et al describe an expert system for a restoration guide in a control centre [60]. One of the functions of the restoration guide is fault diagnosis based on protection and switchgear activity. The diagnosis involves faulty equipment and failed breaker identification. OPS83 and C were used to realise the restoration guide.

The identification of faulty equipment is based primarily on ASTA: finding a blackout area using breaker activity and voltage levels. If main protection relays have operated, the plant in the area which the relays protect is identified as possibly having faulted. If, however, only backup relays have been active, then the possible fault location is that plant which, although protected by the backup relays, is not protected by (unoperated) main protection.

Although the authors present an interesting expert system for guiding restoration, there are certain necessary aspects which the authors make no mention of. As Wollenberg and Sakaguchi, Hotta et al make no mention of ranking competing candidates for fault location. In addition, as in Girgis and Johns' paper, there is no mention of knowledge being incorporated for multiple faults; it appears that false operation due to protection instability is not considered. This would be a valuable aid to restoration, in that it could identify unnecessarily isolated areas. Further points may be made:

- Firstly, despite breaker failures being included in the illustration in the paper, there is no mention of relay failures.
- Although the location of a fault is a means to developing a restoration procedure and not the ultimate goal in itself, an explanation should still be provided. This would gain the operators confidence in the restoration procedure, by justifying the proposed action more clearly.

3.4.10 GESTAL

Arès and Girouard [34] describe a tool for building expert systems for alarm processing and fault diagnosis of power systems. Lisp and ART were used to develop GESTAL, but it has also been implemented using C and OPS83. GESTAL has a frame-based model for each substation. Switchgear and protection alarms are used to identify possible fault locations, the type of fault (ie. the operated protection relay type), failed equipment, relationships between multiple faults and the sequence of operations. Also, the network database is consulted to determine the resulting off-line/on-line state of affected plant. Substation models essentially contain knowledge describing the characteristics and behaviour of the protection and recovery systems. This knowledge is defined in a simple structured language.

A GESTAL-based expert system uses a focussed data-driven approach. By dynamically adjusting the focus of attention, the expert system's inference strategy controls which portions of the knowledge base are invoked. The data-driven inference basically grows directed graphs in accordance with received alarms with causal/associative arcs linking alarms. While obsolete alarms and inconsistent diagnostic graphs are removed, accepted and completed diagnostic graphs are translated to a natural language form for presentation. Diagnoses have been produced by this technique in less than one second.

Arès and Girouard have demonstrated the validity of a GESTAL-based expert system for real-time fault diagnosis. There are however certain deficiencies evident. Although GESTAL can display a set of possible faults where there is more than one potential answer, Arès and Girouard, like Wollenberg and Sakaguchi, have not recognised the need to discriminate between alternatives by a ranking technique. One further unnecessary limitation has been noted:

- The fault diagnosis appears to require the correct receipt of all alarms; there is no mention of diagnoses possibly including missing alarms. However,

unlike ASTA, this stricture is not a feature of the diagnosis philosophy adopted, and is therefore unnecessary.

3.4.11 Computer Assisted Intelligent Storm Outage Evaluator

Balakrishnan and Pahwa [61] propose a Computer Assisted Intelligent Storm Outage Evaluator (CAISOE) for distribution systems. That part of the distribution system which uses fuses is also included in the diagnosis. It was implemented using a Lisp-based tool. CAISOE can locate multiple faults and classify them as high impedance or short circuit and broken or unbroken conductor faults. Semantic networks are used to represent the network (in terms of primary feeders, distribution transformers, and the laterals which connect them) and to generate abstract fault hypotheses. Production rules are used to represent knowledge which refines these.

CAISOE requires the placement of voltage sensors at the location of breakers, fuses, lateral tap-off points, and after small groups of distribution transformers. These need to be connected to the control centre by a dedicated radio channel. When voltage sensors change state, they generate “symptoms” or messages. An intelligent symptom processor suppresses secondary symptoms according to the network topology. The fault hypotheses generator uses the primary symptoms to generate abstract fault locations. The hypotheses validation system determines actual faults, using symptom-fault rules specific to the abstract fault locations.

This expert system can locate faults on distribution systems more accurately than could a system based solely on SCADA system switchgear alarms. CAISOE can also locate faults which are only cleared by fuses. There are certain shortcomings in the proposed system:

- The authors recognise that hardware and communications issues need to

be addressed. For example, dedicated sensors need implementation. The vast number of sensors required by CAISOE is a significant deterrent to its implementation.

- Exclusive symptom patterns are required for each fault (which eliminates the appearance of alternative diagnosis conclusions).
- CAISOE does not cater for sensor or communication malfunctioning.

3.4.12 An Expert System Presented by Okamoto et al

Okamoto et al [62] describe an expert system which identifies fault sections (single and multiple faults), protection and switchgear maloperations and telemetry malfunctions. An object-oriented representation of the power system (in a Lisp environment) is involved, but is not used directly in the fault diagnosis. It is used to automatically derive a semantic (“cause-effect”) network every time the power system is reconfigured. The semantic network links fault section, relay and breaker nodes by the following arcs: protected by (main protection), causes tripping of, and (breaker) failure causes operation of (backup protection). Protection and switchgear alarms and the semantic network are used in the diagnosis process which is suited to parallel processing.

The diagnosis process commences with the identification of protected fault sections from the operated protection relays, and the association of each tripped breaker with its controlling protection relay. For each fault section hypothesis, protection and switchgear maloperations and telemetry malfunctions are identified by recognising patterns in linked nodes. Multiple fault hypotheses are created if false relay activity has been identified, and are dealt with as above. The most probable hypotheses are those with the lowest combined total of system faults, maloperations and malfunctions.

This fault diagnosis system is an interesting piece of work and would appear to hold the promise of effective implementation in a parallel processing environment. There are, however, certain shortcomings. One of these is that the authors like those describing D2, make no reference to a real-time strategy of when to commence the diagnosis. In addition:

- The method of selecting the most probable hypothesis, like the ranking method of D2, can fail to discriminate between diagnoses of significantly differing likelihood. This is demonstrated by the example presented by the authors.
- While there are clear processing advantages in using the described technique, the automatic derivation of the semantic network is time consuming. It is unclear if modifications to the power network can be restricted to a limited updating of the semantic network. This would be required for a practical system since the creation of the semantic network for the very small example network is reported to have taken 6 seconds.

3.4.13 An Expert System Developed by Partanen et al

An expert system is being developed for the fault diagnosis of 10-20kV distribution feeders and also to help reconfigure the network [63]. The toolkit GoldWorks was used, involving frames, rules, object oriented programming and Lisp. The diagnosis is based on data available from the SCADA system and includes switchgear and relay alarms from the primary substations, status of fault current detectors on the distribution lines, and weather data. It also requires the distance to fault data from the operated relay.

The fault diagnosis is initiated following a breaker trip alarm from a primary substation. The breaker's feeder is then considered to be faulty. Using the status of the fault detectors, the possibly faulted line sections are identified. The most

likely fault location is then that section whose electrical distance from the primary feeder matches the relay's calculated distance to fault. The nearest upstream isolator is then advised to be closed.

There is still a lot of work to be done by the authors. However, from the work which has been presented, a few points should be made:

- There is no means for coping with failure of equipment. For example if the fault detector failed, the expert system should be able to use the distance to fault information to at least identify the alternative possible fault locations.
- The authors make mention of heuristic knowledge concerning such things as weather conditions. How this is used in what would appear to be a largely determinate diagnosis is unclear. Moreover, use of such knowledge would then require the inclusion of a ranking technique.
- Few distribution SCADA systems provide access to relay's distance to fault calculation results.

3.4.14 An Expert System Developed by Hsu et al

Hsu et al [64] describe an expert system for locating faults on a distribution system. It is specifically that part of the distribution system between primary feeders and customers on which this expert system performs fault diagnosis. Faults on this plant are cleared by fuses; no knowledge of the isolated plant is automatically available at the control centre. Neither is there an accurate schematic of the network in the control centre for which this expert system was developed. This expert system uses trouble calls to diagnose faults. Prolog was used to develop this system. Network data includes a nodal feeder circuit configuration representation, plant location, location and telephone numbers of key customers (those with technicians responsible for the electricity supply), and plant with a higher

risk of failure.

Following trouble calls (calls from customers whose supply has failed), the expert system determines those branching points in the vicinity of each customer, and hence the possible supply routes from primary feeders. The intersection of customers' possible supply routes is taken as containing the likely fault location. The telephone numbers of key customers supplied from branching points on these routes are displayed to the operator so that he may contact the technician on these premises to ascertain if the supply has been lost there too. Easy access to and presentation of such data is provided by Geographic Information Systems (GISs), whose potential is beginning to be recognised by the electrical power utilities. The resulting information, together with the knowledge of higher risk plant, is used to narrow down the possible fault locations.

This expert system, though less automatic, can be commended as it is designed to cope with the same limitations faced by present operators; it does not require extensive modifications to the existing distribution system. Nor does it require the addition of information which is not already available in the distribution control centre, though in a different format. As presented, it is however restricted to only single faults.

3.4.15 Customer Restoration And Fault Testing

Customer Restoration And Fault Testing (CRAFT) is an expert system whose purpose is to identify the faulted section of multiply tapped 115kV transmission lines and propose restoration actions [65]. It follows from work by Tomsovic et al [66], to which has been added knowledge bases specifically for the validation of isolator status. The expert system has been operational in a control centre since mid-1988. CRAFT is production rule based and is realised in OPS83. The basis of its diagnosis is telemetered switchgear alarms and knowledge of the function

and time settings of automatic switchgear.

When a non-control alarm is received, CRAFT considers that switchgear's line to be faulty. The diagnosis is delayed for a period of 1 minute to allow for the completion of the system's response to the fault, during which other alarms may be processed. Following this, the need to apply the isolator status validation knowledge is checked. Its application corrects the isolator status data before the main CRAFT knowledge base is used. The identification of the faulted section is achieved by the application of empirical rules which use the post-fault configuration of the switchgear on the line. Empirical rules also advise some control actions.

Although a valuable tool for appropriate utilities, CRAFT has a very limited scope in that it relates specifically to multiply tapped transmission lines which have automatic sectionalising switchgear. It has, however, been successfully implemented in a control centre. The major shortcomings include:

- Although automatic isolator alarm malfunctions are acknowledged, CRAFT assumes that all breaker alarms are correctly received. This would appear to be unreasonably inconsistent.
- All automatic switchgear activity is assumed to be correct. However, switchgear does maloperate, though not often.

The authors recorded these assumptions but do not appear to recognise the full limitations they place on the applicability of CRAFT.

3.4.16 AHFA

AHFA is an expert system for fault diagnosis described by Esp and Cahill [67]. It is developed in Smalltalk-80, and uses both protection and switchgear alarms to identify faulted busbars, lines and transformers, and recognise temporary faults. The network topology is extensively represented in an object-oriented framework.

The diagnosis is based primarily on switching patterns, with protection activity fulfilling a supportive role. AHFA commences diagnosis on receipt of a breaker trip alarm. Then, in a similar way to FAUST, the plant on either side of the tripped breaker and up to adjacent breakers, are identified as possible fault locations. As adjoining breakers trip, the possible fault cases which expected their activity increase in confidence. Fault cases whose expected breaker activity is not received are eventually removed. Diagnosis also uses autoreclosing (breaker and isolator) activity in order to identify if the fault is permanent or temporary.

This extensive use of autoreclose data is quite significant, and use of the logic helps AHFA to locate faults more accurately. However, the diagnosis power is severely curtailed by certain omissions:

- The authors do not indicate any ability of AHFA to cope with multiple faults, nor to diagnose breaker or relay failures. If these capabilities are indeed not included, then the number of competing diagnoses ever produced will be smaller; no means of ranking may therefore be necessary.

3.4.17 An Expert System Described by Claus et al

Claus et al [68] describe an expert system for the fault diagnosis of power systems. The expert system has been developed in Prolog. It has been modelled on the Fukui and Kawakami system with only a few changes. One of these is the addition of knowledge to cope with intertripping in a second stage of diagnosis.

The authors recognised a potential for bad diagnoses in the Fukui and Kawakami expert system and modified the fault diagnosis strategy accordingly. Instead of limiting the possible fault locations to that plant in the intersection of the related relays' protected areas, this expert system considers every item of plant in the blackout area as a possible fault location. Consequently, false operations as well as failures may be related to possible faults. This change results in a correct

diagnosis being made for the example cited by the authors. However, the expert system still suffers from the need to receive every switchgear tripped alarm, the lack of an explanation capability, and the requirement to define the protected areas of some relays as described in section 3.4.2. In addition:

- The “wave approach” for identifying blackout areas which is described by the authors is not suitable for an on-line diagnosis system. Details of switchgear activity would be made available to the expert system in a piecemeal fashion. Therefore, to have a search propagate through the network from every source, following a change in switchgear position would be extremely inefficient.

3.4.18 SEPT

Fauquembergue and Esquevin [69] describe SEPT, an expert system for diagnosing incidents on a power system. SEPT is a relatively mature expert system: its prototype was developed in 1987, and has been operational in an Electricité de France (EDF) area control centre since September 1990. The most likely faulted plant and protection maloperations are identified in an on-line role. The expert system uses event recorder data from one substation as the basis of its diagnosis. Data is automatically retrieved via a PC and modem following a protection operation has been identified in the control centre RTU information. The expert system has been realised using a development tool; knowledge is represented by production rules, objects and methods. A deep protection system model is included.

The diagnosis consists of four parts. The “verification of the protection systems” uses a pattern matching approach to verify the dynamic activity of the protection or identify anomalies. Heuristic rules and interval filtering are used in the “search for elementary incidents” which seeks to divide the event recording

into intervals of time containing activity relating to individual incidents. The “fault section estimation” considers each high voltage item of plant to be possibly faulty, for which each verified protection relay (no anomalies) can vote for or against on the basis of the protection system model; the section with the highest vote is defined to be faulty. The “synthesis” uses heuristic rules to perform a final check, to use the faulted section conclusion to identify further protection maloperations, and to distinguish between temporary and permanent faults. Summary reports are automatically printed out following completion of the diagnosis. In addition, SEPT sometimes consults the operator where one of the diagnosis parts fails, for example, due to a tie in the voting process. Static explanatory text, voting results, and the rules used are also available on request.

SEPT has been proved to be a practical tool for a utility, and demonstrates how substation sequence of event recorders may be consulted by expert systems in order to perform automatic fault diagnoses. Despite this, there is a major shortcoming of the strategy adopted:

- The means of diagnosis, especially in the “verification of the protection systems”, is inflexible. The knowledge is very much specific for a particular substation arrangement and standardised protection system. As a result, modifications to the knowledge base necessitated by changes to network or protection equipment are complex and lengthy. The authors recognise this and indicate that it was due to the original design aims. The system is therefore not particularly suited to system-wide application.

3.4.19 An Expert System Developed by Amantegui et al

Amantegui et al [70] describe an expert system for fault diagnosis in power systems. An initial prototype was developed using a shell (M1), C and assembler, and was operational in real time for one and a half years. The final system

was however developed using ART and C, and involves rule-based parts, object-oriented parts and procedures. The expert system uses SCADA alarms.

The fault diagnosis reasoning is performed by five separate modules. The “problem solving kernel” controls the performance of the expert system. The shallow knowledge which forms the “selector of the possible fault locations” generates hypotheses of the fault location with associated certainty factors. The “hypothesis validation and fault location” module identifies the alarms related to each hypotheses and adds deviations where necessitated by the received alarms. The hypothesis which includes the fewest maloperations is selected as the conclusion. If a conclusion is not reached, then the “simultaneous disturbance analyser” attempts to generate multiple fault hypotheses by trying to recognise geographically distinguishable clusters of alarms, each relating to one fault. Among the roles of the “time pursuit of a disturbance” module, is the validation of hypotheses in line with automatic resetting/reclosing and operator action. Further developments have been engaged in to make use of sequence of event recordings, and to provide restoration knowledge.

The expert system provides a very good fault diagnosis role. Moreover, the performance results cited by the authors appear strong. It is, however, unclear how the two means of uncertainty measurement (certainty factor and maloperation count) work together. One serious shortcoming is apparent from the authors’ paper:

- The diagnosis knowledge appears to make use of the network arrangement, but not of the present configuration. It is unclear how the expert system would perform a correct diagnosis if certain circuit breakers had remained in the open state. Recent events should be able to be diagnosed because of the role of the “time pursuit” module. Distant incidents or operator action which have left the network in an abnormal configuration, may (based on

the cited paper) cause some problems for the expert system.

3.4.20 Other Systems Related to Fault Diagnosis of Power System Networks

There has been much development work in the related field of alarm processors. A number of intelligent alarm processors have been developed. In addition, a large number of systems have been developed for the fault diagnosis of specific items of, especially nuclear, plant. For example, research in the Centre for Electrical Power Engineering, University of Strathclyde has been active in both these areas: a practical alarm processor, APEX, has been developed [37]; a turbo-alternator fault diagnosis expert system, called MODEES, is being developed [71]. EKA and SUBES are more closely related to the work reported in this thesis.

3.4.20.1 EKA

Keronen [72, 73] is developing an expert system for operation planning and event analysis, called EKA (a Finnish acronym for an expert system for power system operations). It has been developed using KEE, Lisp and FORTRAN, and contains a comprehensive hierarchical frame-based representation of the power system. Although the event analysis module has still to be developed, Keronen intends to include causal, eventual and temporal relations together with pattern matching and simulation based reasoning. Insufficient details have been established to merit a fair evaluation.

3.4.20.2 SUBES

Kaminaris et al [74] developed an expert system for the related area of substation troubleshooting. It is called SUBES and is described as a rule-based expert system approach for (manned) distribution substations troubleshooting and maintenance

scheduling. Protopapas et al [75] described a similar expert system, which could have been the forerunner of SUBES. Developed in Turbo Prolog, the expert system uses a tree structure with various classes of node, and involves production rules to represent the diagnosis knowledge. The expert system is implemented on a PC, and can interface with a substation SCADA system. Much development effort has been invested in the user interface, which provides menus and also displays portions of the substation one-line diagram.

The fault diagnosis is performed by traversing the tree. As the inference progresses down the tree, the expert system goes from identifying parts of the substation, through the protection level, to the cause and ultimately to remedial measures and explanation. At certain nodes the system requests information from the user. Tests, like gas chromatography, are proposed at one of the tree levels. SUBES also generates substation maintenance schedules.

Insufficient details are supplied to support a detailed evaluation. However, one apparent omission is the ability to generate multiple possible causes—the reasoning stops when all the activity is accounted for. Also, when travelling down the tree, the “left hand child” of a node is always attempted first, and the “right hand child” (a child is a connected node on the next lowest level of a tree) is only ever attempted last. Thus the technique requires the arrangement of next steps (or node children) in a static order of best first. System knowledge would not always fit in with this arrangement.

3.4.21 General Comments

There are certain functions which have been provided together with the core fault diagnosis role of the above systems. These include alarm processing/filtering (D2 and FAUST) and live interactive graphics (AHFA and EKA).

On the other hand, it has been demonstrated how many of these expert sys-

tems have certain deficiencies. These include unrealistic assumptions, vast computer resources requirements, dependency on having complete information (ie. their diagnosis is not robust), and indecisiveness (ie. inability to discriminate between several possible diagnoses).

In addition, with the exception of Dy Liacco and Kraynak and of Girgis and Johns, researchers appear to neglect earth faults. Often a general short circuit fault is, instead, only considered. Dy Liacco and Kraynak only mention the need for special logic for earth faults; Girgis and Johns' expert system only considers the fault type after the location of the fault has been established. Utilities employ some dedicated earth fault protection. Therefore, the behaviour of the network following the occurrence of a fault, is influenced by the nature of the fault. Power network fault diagnosis systems could therefore include the identification of the fault type from the protection and circuit breaker response. In fact, to ignore this, will limit a system's ability to identify protection maloperations correctly.

The classification of the fault type would help the operator assess damage to system plant. For faults on overhead lines without autoreclose equipment, the classification of the fault type may influence the decision whether or not to reclose the line. For example, a phase fault might indicate conductor clash, following which reclose procedures may be initiated; a ground fault may indicate a more permanent fault such as insulator string failure.

This, together with the identification of the faulted phase from digital recorder data, could become more important as utilities seek to increase the utilisation of existing plant. Engineers at one utility are known to be considering phase segregated operation.

Chapter 4

Diagnostic Knowledge

4.1 Preliminary Discussion of Implementation Issues

The developments in power systems have indicated that expert system techniques could provide systems that are both suitable and necessary for on-line network fault diagnosis. However, fault diagnosis of power networks differs in several respects from those problem-solving tasks for which general (or “first generation”) expert systems were developed. This has influenced the design of the fault diagnosis system reported in this thesis. These differences include:

- automatically supplied evidence:- The SCADA system is the main source of the information on which the diagnosis is based. Data retrieval is thus automatic, fast and a relatively random process. Consequently, the provision of a user interface dialogue is less significant than in the many expert systems in which the human user plays an important role in the provision of evidence. Moreover, some of the explanation techniques developed for other systems are irrelevant since extensive two-way dialogue will not be engaged in.
- size of the object of diagnosis:- The power system network is very large, and the number of possible conditions and diagnosis conclusions is vast. (For example, the ScottishPower transmission network contains 120 substations, with 2,000 controls and 15,000 alarms in total.) Semantic networks or scripts representing the knowledge would be very large. Such abstract networks were therefore considered unsuitable. Their unsuitability is corroborated by the observation that knowledge is equally appropriate in various parts of the system. Also, unlike many systems, power networks have a regularly changing topology as well as changing conditions.

- depth of reasoning:- Combining the latter point with the desired level of robustness, results in a significant depth of reasoning being required. A degree of system modelling is necessary.
- legitimacy of many answers:- The process of fault diagnosis can reasonably generate several valid answers. The reasons for this shall be detailed in section 4.7.2. The means of producing several alternatives rather than one possibility must be provided.
- interfacing to numerical algorithms:- The ability to interface with standard power application software would be a valuable attribute. Although not an immediate necessity, this facility would extend the system's potential benefits.

The work reported in this thesis relates to the production of such an expert system. In the light of the above, Prolog ("PROgramming in LOGic") was chosen as the means of implementation. Some of the advantages of Prolog may be listed as:

- it is a general purpose language with simple semantics
- symbolic reasoning is provided
- it is based on logic, which suits the task
- a backward chained (or goal directed) inference mechanism with pattern matching and backtracking are provided (Curtice [76] claims that Prolog may be more suitable than OPS5 or OPS83 (Official Production System), which are data driven, since many tasks in power system operation are goal directed.)
- rapid prototyping is possible

- it has the flexibility to establish grammatical structures
- it supports the development of shells, other reasoning techniques and knowledge representation methods, and aspects of other paradigms
- its established record of use has proved its validity
- it is widely supported on different levels of hardware (described by Subrahmanyam [12])
- interfaces to conventional languages are provided by some Prolog systems
- reasoning can be transparent
- a degree of parallelism is natural (described by Bobrow [16])

Many of these attributes were utilised while developing the fault diagnosis expert system. As a result, the expert system's development has been more flexible. The Prolog system used was Quintus Prolog [77]. This provided a foreign language interface by which the expert system used a loadflow program written in the C Programming Language. Appendix A gives a brief introduction to some of the fundamentals of Prolog.

4.2 Structure of Diagnosis Results

4.2.1 Observations

In a generation system, the plant is more localised and strongly interconnected than in a transmission or distribution system. There is seldom more than one simultaneous (ie. close in time) disturbance. However, in a transmission or distribution system, the volume, dispersion and vulnerability of plant is such that there may be several on-going independent disturbances at any point in time. There could be several alternative explanations attempting to describe any disturbance. Moreover, since the fault diagnosis is performed on-line, the diagnosis conclusions may be subject to change. As a result, the diagnosis conclusions cannot only be expressed as textual output, but must be held in the knowledge base as knowledge.

4.2.2 Initial Structure

In the early prototypes, alternative diagnosis conclusions were expressed in scenario facts. Each scenario accounts for every received alarm, and is therefore in competition with every other scenario. Consequently, there is one ranking table which contains all existing scenarios in order of decreasing chance of fully describing the genuine system behaviour.

Each scenario may contain several causal events. Causal events include system faults, failed equipment, false activity and telemetry malfunctions. A system fault includes the faulted item of plant, the generic type of fault (ground or phase fault), and the fault permanence (permanent or temporary fault). Failed equipment is particularly the protection and/or switchgear which has wrongly not operated, while the phrase false activity is applied specifically to devices which have wrongly been active when they should have remained inactive. Failed

equipment includes both the failure to detect a fault and the failure to operate following a control signal from a relay. Missing alarms result from telemetry malfunctions, and relate to action which has taken place but has not been alarmed. The telemetry which supports the causal events is also expressed in the scenario fact. Each scenario fact has eight arguments in total: ranking position, measure of uncertainty, system faults, device failures, supportive alarms, false activity, predicted/missing alarms, and an argument for extra knowledge. This arrangement has the legibility advantages of the frame structure. The example which follows describes the disturbance shown diagrammatically in Figure 4.1.

```
scenario(
  rank: 1,
  uncertainty_measure: 0.001,
  faults: [phase fault_on ss1 to ss2::l1 which_is
    (permanent or temporary)],
  failures: [],
  supportive: [operated(r1 at ss1), operated(tr1 at ss1),
    operated(r1 at ss2), operated(tr1 at ss2), tripped(x1 at ss1),
    tripped(x2 at ss2)],
  false: [operated(r2 at ss10)],
  predicted: [operated(tr2 at ss10), tripped(x2 at ss10)],
  extra: [events_alarms(details(1990,1,11,15,41,37.0),
    phase fault_on ss1 to ss2::l1 which_is (permanent or temporary),
    [operated(r1 at ss1), operated(tr1 at ss1),
    operated(r1 at ss2), operated(tr1 at ss2), tripped(x1 at ss1),
    tripped(x2 at ss2)], [], isolated),
    events_alarms(details(1990,1,11,15,41,39.0),
    operated(r2 at ss10) is false,[operated(t2 at ss10),
    tripped(x2 at ss10)], [], _)]
).
```

This example illustrates how the extra knowledge relates alarms and failures to specific causal events.

While this structure proved satisfactory in many respects, it has a few shortcomings:

- unrelated events are indistinguishable: Each scenario groups together

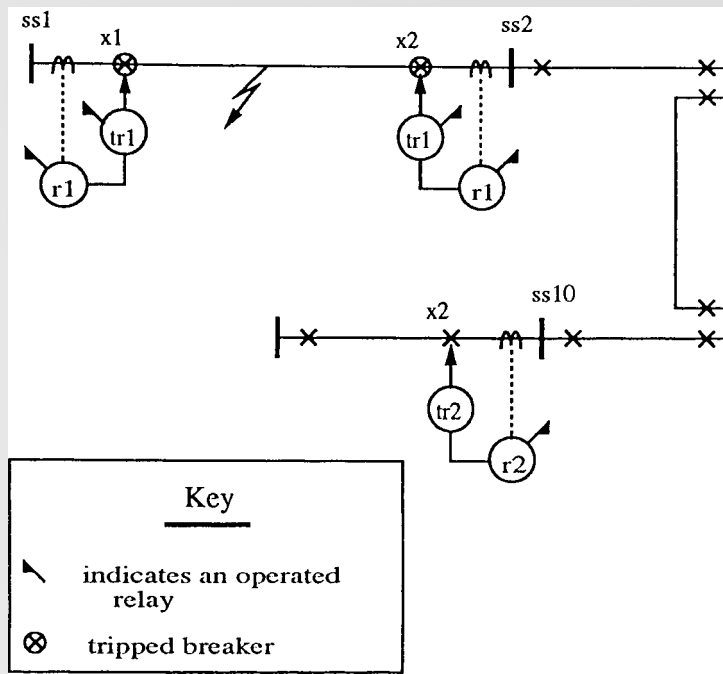


Figure 4.1: The diagrammatic representation of the example scenario

both related and unrelated causal events. There is no distinction made between events which may relate to the same disturbance and those which are clearly independent. This limits the inference which can be drawn.

- on-line management of conclusions:- After a certain period of time, the details relating to a particular disturbance will not change. However, changes due to other on-going disturbances will continue to change the scenarios. Extra work in recalculating combined uncertainty values is therefore involved. The recognition of a final conclusion on the completed disturbance and its removal from continued consideration is unnecessarily difficult.
- indication from ranking:- By combining all causal events together and evaluating a combined uncertainty measure, wrong indications can be given. For example, suppose two independent disturbances have explanations $C_{11}, C_{12}, C_{13}, \dots, C_{1x}$ and $C_{21}, C_{22}, C_{23}, \dots, C_{2y}$ respectively. $C_{11}, C_{12}, C_{13}, C_{14}$ are all very similar in content and hence in levels of uncer-

tainty, whereas C_{21} and C_{22} are slightly more different. The ranked list of scenarios may then have the order $C_{11}C_{21}, C_{12}C_{21}, C_{13}C_{21}, C_{14}C_{21}, C_{11}C_{22}$, etc. That could give the false impression that only C_{21} should be considered as a reasonable explanation of the second disturbance.

- combinatorial growth:- By not separating unrelated disturbances, the diagnosis can appear more complex than it is. For example, two independent disturbances, each with three alternatives, would form nine scenarios. A list of nine scenarios will always appear more complex than two lists of three.

As a result, a third layer was added to the diagnosis conclusion structure.

4.2.3 Three-Tier Hierarchical Structure

Diagnosis conclusions are formed according to a three tier hierarchy. This provides for the separation of unrelated disturbances, the comparison of competing conclusions and the combination of related occurrences. The three levels which each diagnosis conclusion has are the disturbance, the competing scenario and the causal event.

Potentially related behavioural activity belongs to the same disturbance. Each disturbance has an area of involvement which identifies that part of the network which could be involved in the same disturbance. A ranked table of competing scenarios is maintained for each disturbance.

Competing scenarios are alternative explanations of the same disturbance, rather than of the whole system condition. A disturbance's diagnosis conclusions are expressed in the same format as that described previously, except that the scenario facts have an extra argument describing the disturbance to which the scenario relates. The following example illustrates this.

```

scenario(
  disturbance_id: -19999,
  rank: 1,
  uncertainty_measure: 0.001,
  faults: [phase fault_on ss1 to ss2::l1 which_is
    (permanent or temporary)],
  failures: [],
  supportive: [operated(r1 at ss1), operated(tr1 at ss1),
    operated(r1 at ss2), operated(tr1 at ss2), tripped(x1 at ss1),
    tripped(x2 at ss2)],
  false: [],
  predicted: [],
  extra: [events_alarms(details(1990,1,11,15,41,37.0),
    phase fault_on ss1 to ss2::l1 which_is (permanent or temporary),
    [operated(r1 at ss1), operated(tr1 at ss1),
    operated(r1 at ss2), operated(tr1 at ss2), tripped(x1 at ss1),
    tripped(x2 at ss2)]), [], isolated)]
).

```

The above again relates to the disturbance illustrated in Figure 4.1, although this time the 'ss10' activity is not included since it belongs to a different disturbance.

Although some unrelated events will still appear in the same "disturbance", this will occur less frequently. The problems with the initial structure are largely overcome.

4.2.3.1 Disturbance Area

The disturbance area must include that part of the network which could be involved in the disturbance due to its connective proximity. That is, it contains the protection which could reasonably be involved in the system's response to that disturbance.

Initially it was thought that a planar grid could be used to divide the network into disturbance areas. Along the boundaries of these areas, protection on one side of the boundary would not be affected by system events on the other. However, for such boundaries to be created, then there must be parts of the

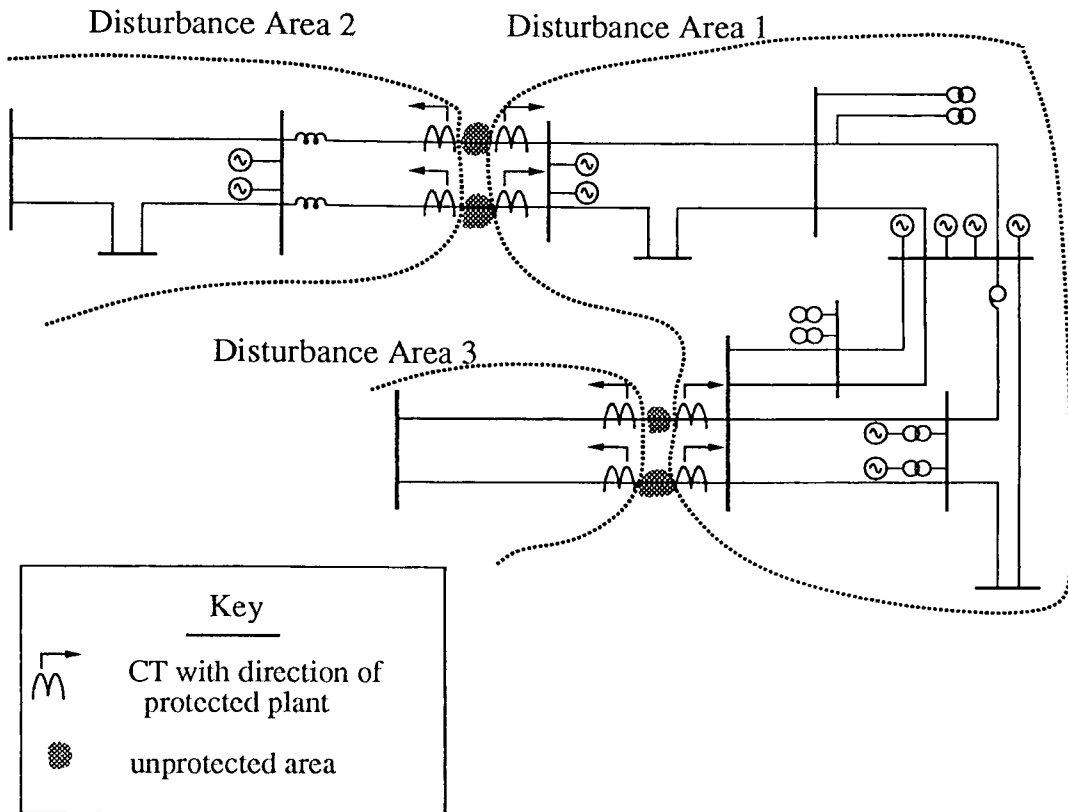


Figure 4.2: Illustration of the effect of a planar disturbance area

network which would be unprotected. This is shown in Figure 4.2. Protections usually overlap for this very reason of avoiding unprotected areas in the network, as shown in Figure 4.3. Settings of non-unit protections are calculated so that the relays provide remote backup. The area of overlap can therefore be significant. For example, the standard 80%/150%/225% settings of 3-zone distance protection have a reach greater than two substations. A planar grid is therefore inappropriate.

An array of overlapping areas could be established. A high degree of overlapping would be the consequence of the degree of overlapping protection zones and of network interconnectivity. A disturbance area would be dynamically formed from those areas overlapping a particular point in the network. Due to the level of overlapping, disturbance areas would consist of several defined areas.

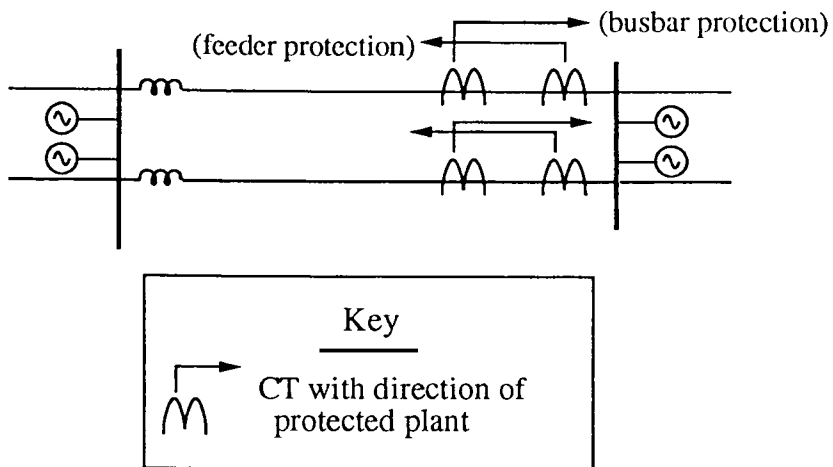


Figure 4.3: Overlapping of protections

It was felt that using several defined areas gave little advantage over forming the disturbance area directly from the topology. Therefore, the solution adopted is to create the disturbance area during the diagnosis based on the network topology. The actual reasoning used to form the area is dependent on the diagnosis philosophy used.

4.3 Structure of the Means of Diagnosis

4.3.1 Observations

In the fault diagnosis of transmission system networks, there are two competing factors, namely scale and detail. The scale factor requires the expert system to be able to diagnose faults throughout the network; the detail factor dictates that the diagnosis should focus down to as much detail as permitted by the available data. Having the diagnosis knowledge at different levels and using a “focusing” approach resolves this conflict. With further development, the hierarchical diagnosis could give further flexibility: a shallower diagnosis could be performed when speed is critical during heavy alarm streams; a deeper diagnosis could be performed once detailed knowledge from disturbance recorders or microprocessor based relays is received at the expert system site.

4.3.2 Hierarchical Structure

4.3.2.1 Gut Reaction

The top level diagnosis establishes a “gut reaction” or first impression based on the interpreted alarm. The diagnosis knowledge can identify activity which can only be as a result of manual control. Additionally, automatic activity alarms may be placed in the context of the state of the network by establishing the alarm’s relationship to disturbances. The alarm type and nature are determined, with some simple abnormalities being checked for.

The knowledge is represented by production rules for clarity. These are stored in rule/2 facts as shown by the following example:

```

rule(hyp_gut,
if judgement(isnt control) and
  judgement(present_alarm is Alarm-Type) and
  ( Alarm = tripped(_) or
    Alarm = reclosed(_) or
    Alarm = reset(_) or
    ( Alarm = operated(_) and
      Type = trip))
then assert(judgement(secondary_alarm))
).
```

The first argument of the fact identifies the group of rules to which the rule belongs—the hypothesiser “gut reaction” (‘hyp_gut’) in this case. The second argument contains the rule logic expressed in the classic IF-THEN format. Satisfaction of the rule antecedent requires satisfaction of all ANDed antecedents. In the above rule, satisfaction identifies the alarm to be a secondary alarm. For satisfaction, previously asserted judgements are used. These identify the alarm not to relate to operator control, but to indicate a (breaker) trip or reclose, or a (relay) reset, or a trip relay operation.

A supplementary inference engine was developed in Prolog in order to use knowledge in this form. The reasoning commences with finding a fact containing the appropriate production rule. The rule’s conditions are checked, and actions are taken according to the logic of the rule. This is then repeated until all the appropriate production rules have been attempted. Mostly, the rule conclusions add (ie. “assert”) facts to the knowledge base which, as well as being subsequently used in further production rules, builds up a view of the alarm. These details are held in the argument of judgement/1 facts, like the judgement(secondary_alarm) created by the above example. The opposite of the conclusion of some rules are required. Rather than count on the absence of a judgement/1 fact which could be a result of the appropriate rule not having been used, a facility has been provided for storing an opposite conclusion. These rule conclusions include a make_judgement/1 predicate. If the rule conditions are

satisfied then the argument is stored in a judgement/1 fact and the rest of the conclusion is performed. Otherwise, the operator 'isnt' is appended. Thereby facts like judgement(isnt wrong_msge) may be created.

The picture of the alarm which is built up, is then used to guide the knowledge contained in the central diagnosis. Therefore, the production rules are somewhat different for the different diagnostic philosophies.

4.3.2.2 Central Diagnosis

The central diagnosis performs most of the diagnostic work, and creates and modifies scenarios. System faults, protection and switchgear maloperations and telemetry malfunctions are inferred at this level. Differing philosophies have been developed to establish this diagnostic ability, and these are described in the following sections.

The standard Prolog logical rule structure is used to realise the central diagnosis knowledge. The detailed inference required by the depth of reasoning is realised in these logical rules. Though less understandable than the gut reaction production rules, these logical rules have realised knowledge which could not have been effectively represented in such production rules. Despite this, the knowledge remains reasonably understandable due to Prolog's semantics.

4.3.2.3 Deeper Diagnosis

The deeper diagnosis makes use of the ranked scenarios produced by the central diagnosis. As research work has been concentrated on the central diagnosis, the deeper diagnosis has not yet been fully implemented. Some ideas have been put together as a feasibility test.

The idea is to use deeper relationships in system behaviour, use non-electrical knowledge, use deeper plant models, and have the knowledge to validate and add

detail to previous diagnoses using data from disturbance recorders and microprocessor based relays. It is envisaged that logical rules will be the best suited for the diverse aspects of knowledge.

4.3.3 Performance Monitor

As well as performing a fault diagnosis, the expert system monitors the performance of the protection and switchgear. The performance monitor records and indicates to the user the failed protection or switchgear equipment which is included in the top few scenarios. As this is influenced by incoming alarms, the survey of a disturbance's scenarios is delayed until the disturbance is considered complete. Correct operations are also recorded as these, in contrast to the above, indicate good performance.

4.3.4 Management

Since the fault diagnosis is a continuous process, the reasoning will not reach a conclusion and stop. Therefore, the expert system needs the knowledge to recognise when a system event or a disturbance is complete and can be taken out of continued consideration. The deeper diagnosis can then be applied to the central diagnosis' final ranked table of conclusions. Indications of bad or good performance will then be identified by the performance monitor.

4.4 Philosophies of Diagnosis

4.4.1 Historian Philosophy

Initially, the “historian” philosophy was used. In this context, a historian is regarded as one who seeks to explain the past while looking at present evidence. Therefore, this philosophy did not realise an on-line diagnosis. The developed post-fault diagnosis knowledge was based on the work of Fukui and Kawakami [50] (see section 3.4.2). The expert system requests the operated relays and tripped circuit breakers. Isolated (blackout) areas are determined using the post-fault network configuration. Possible fault locations within the blackout area are narrowed down by only considering plant which is in the intersection of the protected areas of related operated protection relays. Multiple faults are resolved as described by Fukui and Kawakami.

As an on-line aid was desired, the “wait-and-see” and “hypothesiser” philosophies were developed.

4.4.2 Wait-and-See Philosophy

Faced with the appearance of protection and switchgear alarms, a power system operator may wait until he has established that an area of the network has been isolated, before attempting to identify possible faults. The appearance of a blackout area is regarded as being indicative of the occurrence of a genuine system fault. Having located a blackout area, the operator identifies related protection activity: from the record of operated relays, he selects those whose operation caused the tripping of those breakers which isolate the blackout area. He then resolves the system faults inside the isolated area which account for the related protection activity. This is the philosophy used for the wait-and-see fault diagnosis.

The on-line SCADA system alarms describing protection and switchgear activity are used. It is the switchgear activity which is of primary importance; protection activity has a secondary role. Following every switchgear trip alarm, the electrical status (ie. dead/live status) of the switchgear terminals is checked. Protection activity is used to identify possible faults and to detect the occurrence of maloperations, only when a new blackout area has been detected. Additionally, if a blackout area is created not long after an adjacent blackout area, multiple faults and false operations are inferred using the ideas suggested by Fukui and Kawakami [50]. A remote blackout area forms a disturbance area, as do related adjacent blackout areas.

4.4.3 Hypothesiser Philosophy

In contrast with the aforementioned example, another operator may consider it unreasonable to wait for the recognition of an isolated area of the network, especially when he is aware of the possibility of alarms being lost. He determines which system faults a relay can detect following the receipt of an alarm describing that protection's operation. For each possibility (or "hypothesis"), he resolves the alarms which he would expect to see if the suspected fault was genuine. By checking whether these expected alarms have been received, hypotheses are discarded or kept. Furthermore, he may add device failures to the hypotheses. Subsequent activity would influence his confidence in these hypotheses. The hypothesiser follows this philosophy, which is based on the generate-and-test paradigm.

Again, on-line switchgear and protection alarms are used. However, the hypothesiser philosophy is primarily protection activity driven. It involves the prediction of activity, and hence has the potential to cope with missing (or corrupt) data. Data can be missing due to lost alarms or due to equipment not being wired into the SCADA system. Just as faults are hypothesised, so are malop-

erations. However, there are many fault and maloperation combinations which can cause relay operations, and consequently there are many fault plus maloperation hypotheses which could be generated following such an operation. Consider, for example, the simplest case of a unit protection relay controlling two trip relays which control two breakers: following the protection's operation, and only considering one possible fault, there is one no-failure-hypothesis, four one-failure-hypotheses, six two-failure hypotheses, etc. The hypothesiser philosophy has had several changes in order to contain this combinatorial avalanche. Multiple fault conclusions are formed by merging new single fault scenarios with previous scenarios for the same disturbance area. A disturbance area, according to the hypothesiser strategy, is the network in the vicinity of the substation which generated the original protection alarm. This region is derived on the basis of electrical topology, not geographical proximity.

4.5 Wait-and-See

4.5.1 Gut Reaction

Firstly, the “gut reaction” performs general checks on the interpreted alarm. The validity of the alarm is tested by verifying that the alarm has not already been effected in the knowledge base. The primary/secondary nature of the alarm is ascertained and, if secondary, the alarm is checked if it was controlled by an operated device and hence expected. Indeed, further rules could easily be added.

The gut reaction also has rules which relate specifically to the wait-and-see strategy. The wait-and-see is primarily circuit breaker activity driven, and so the reasoning following such is greater. Following a breaker trip alarm, the status of the breaker terminals is established. It is inferred whether each terminal is electrically dead or electrically live by topologically searching for connected power infeeds. A power infeed is plant which could feed a system fault, and therefore includes machines and supply points (plant connected to undefined parts of the network which can supply power).

The knowledge to determine electrical state is realised in logical rules due to the reasoning involved. The topological search commences at the plant connected to the breaker, and proceeds by identifying the plant connected to it. This is shown in Figure 4.4 by part 1 of the broken line representing the search which starts from the dead terminal of breaker '195 at chap'. (In this figure, the topology has been shown in accordance with isolator status, but without the isolators being shown.) Plant connected to the found plant is progressively identified until either an open breaker or a power infeed is reached. When an open breaker is reached in the search, the traversed path is effectively retraced in order to find an alternative path. This is illustrated by part 2 of the dashed line in Figure 4.4. If a connected item of plant is a power infeed, then there is no blackout area—the




Figure 4.4: The search to determine the electrical status of breaker terminals plant is live. Otherwise, as shown in Figure 4.4, if the rule fails to find a power infeed, the traversed plant forms a blackout area.

The path formed by parts 1 to 4 in the figure indicates that basically a depth-first strategy is adopted. The decision not to use a width-first search was based on the high degree of network interconnectivity: a width-first search could quickly divide into a large number of on-going search paths which would use significant memory resources. In order to improve the efficiency, the idea of the heuristic search has been incorporated into the search (see section 5.5): knowledge of normal load flow is used in order to check initially in the direction most likely to have a power infeed.

If the status of the two breaker terminals is 'live' on one side and 'dead' on the other, then a new blackout area has been created. A gut reaction rule compares the breakers isolating the new blackout area with those of existing blackout areas in order to identify adjacent blackouts. Thereby, it is inferred whether it is a new disturbance area or belongs to an old one.

An example of the knowledge derived by the gut reaction follows:

```

judgement(present_alarm is tripped(x1 at ss1)-breaker).
judgement(isnt wrong_msge).
judgement(secondary_alarm).
judgement(msge_expected).
judgement(msge_not_false).
judgement(terminal_status is dead/live).
judgement(blackout is [ss1 to ss4::l1\\e1, ss1 to ss4::l1\\e2,
    ss4::b1\\s1, ss4::b1\\s2])).
judgement(new_disturbance).

```

4.5.2 Central Diagnosis

The gut reaction indicates if the alarm is reasonable or not. If it is, then the action described by the alarm is effected on the representation of the network held in the knowledge base. The rest of the response of the central diagnosis depends on the type of the alarm.

When an alarm is received which relates to the operation of a protection relay, then the expert system checks that at least one of the relay's current transformers (CTs) is on live plant. If all its CTs are on dead plant, then the relay's operation is added to scenarios as slow or false: slow if the blackout area was created not long (an arbitrary time period) before the relay operation and the relay was classed in the scenario as a failure. The operation is presumed reasonable at this stage when a CT is found on live plant. For these relays, the possible faults which could have caused the relay activity are identified and asserted in the knowledge base. The protected areas of some protections are influenced by switchgear activity, and therefore the possible faults need to be determined before subsequent changes in switchgear status is carried out in line with alarms. All reasonable protection and trip relays are initially recorded as "unaccounted for".

Following circuit breaker activity, the degree of diagnostic reasoning used in the central diagnosis is guided by the status of the breaker terminals as determined by the gut reaction. There are three alternatives:

- live/live:- The breaker activity is logged as unaccounted for, and no more reasoning is performed.
- dead/dead:- If a controlling relay is known to have been slow then no more reasoning is performed; otherwise, the breaker is classed as slow. In either case the scenarios are updated accordingly.
- live/dead:- A system fault is presumed to have occurred and been cleared, and the diagnosis continues into more detail.

The purpose of the more detailed reasoning is to determine the possible system faults. Scenarios are updated or generated accordingly.

4.5.2.1 Determination of Possible System Faults and Maloperations

The possible faults are located within the blackout area which was determined in the topological search of the gut reaction. Related operated relays are then determined from the recorded unaccounted protection and the tripped breakers defining the blackout area: from the log of unaccounted operated protection, the reasoning identifies those protection relays which controlled at least one of these breakers and which could have seen a fault in the blackout area, and the trip relays through which the breaker was controlled. Unaccounted operated protection relays, whose breakers have not tripped, but which could have seen a fault in the blackout area are also identified as related relays, as are the operated trip relays they controlled—this covers those relays whose breakers have thus far failed to trip. The possible faults are then those faults which could have been detected by all the related primary relays and which are located in the blackout area. Each possible fault will be the basis for a temporary single fault scenario.

From the operated relay list, and using relay auxiliary switch logic, trip relays which failed to operate are identified. These failures, together with the failed breakers previously mentioned, are independent of the fault, and so are common

to all new temporary scenarios. In contrast, primary relay failures are dependent on the fault location and type. For each possible fault, unoperated protection relays which can still “see” the fault are identified as failures. To these are added every unoperated relay which could have “seen” the fault—these are identified as those which can “see” the fault if all breakers are assumed closed and for which a relay with the same CT positions has operated. These possible faults and maloperations are recorded in temporary single fault scenario facts.

4.5.2.2 Formation of Complete Scenarios

If the gut reaction resolves that the fault was a new disturbance, then the temporary single fault scenarios are copied and recorded as complete single fault scenarios. However all the temporary scenarios are combined with all the disturbance’s existing scenarios if the fault was related to an existing disturbance.

The first way of combining a new scenario with an old one, is on the basis that the events in both scenarios genuinely occurred. The faults and maloperations in the two scenarios (and the corresponding alarms) are added to form a new scenario. On the other hand, due to the vicinity of the isolated regions, it is feasible for a blackout area to have been falsely created because of the vicinity of the genuine fault. That is, a nearby fault destabilised protection which caused the breakers to trip, creating a blackout area. The falsely created blackout area can be the latest region to be isolated, or one on which the old scenario was based. In changing a fault-caused blackout area to one created by false operations, the supportive secondary alarms remain as such, but the previously supportive primary alarms are changed to false. Similarly, the failed secondary devices (trip relays and switchgear) remain as failed, but the protection relays previously identified as failures are removed since that classification was based on the fault.

4.6 Hypothesiser

4.6.1 Gut Reaction

The hypothesiser “gut reaction” checks for the same general judgements as does the wait-and-see gut reaction. These are described in section 4.5.1. There are also rules specific to the hypothesising philosophy.

Unlike the wait-and-see philosophy, the hypothesiser suspects a fault following one alarm. Therefore the disturbance area is created, not after the isolation of an area of the network, but following alarmed activity which could have been caused by the occurrence of a system fault. The disturbance area is defined by a list of substations. If an automatic activity alarm is from a substation inside a disturbance area, then a rule records the judgement that the activity belongs to that old disturbance. Otherwise, a new disturbance area is created.

A new disturbance area of involvement is generated by forming a list of substations from the alarm “source” substation and those found a prescribed number of substations away which are not already in an existing disturbance area. The disturbance area should contain that plant potentially involved in the disturbance. Therefore, if protection and switchgear was 100% reliable, then the “radius” around the source substation would be one substation. For example, consider the network shown (without switchgear and with substations represented by busbars) in Figure 4.5. The protection relay which has operated may protect substation equipment, or one of the substation’s lines (assuming no device failures). In either case, only protection at that substation, and protection at the adjacent substation will be involved. Hence the disturbance area is as shown, with an effective radius of one substation.

Unfortunately, however, the basis for this restriction of the disturbance area does not hold; protection and switchgear equipment does fail. Therefore, pro-

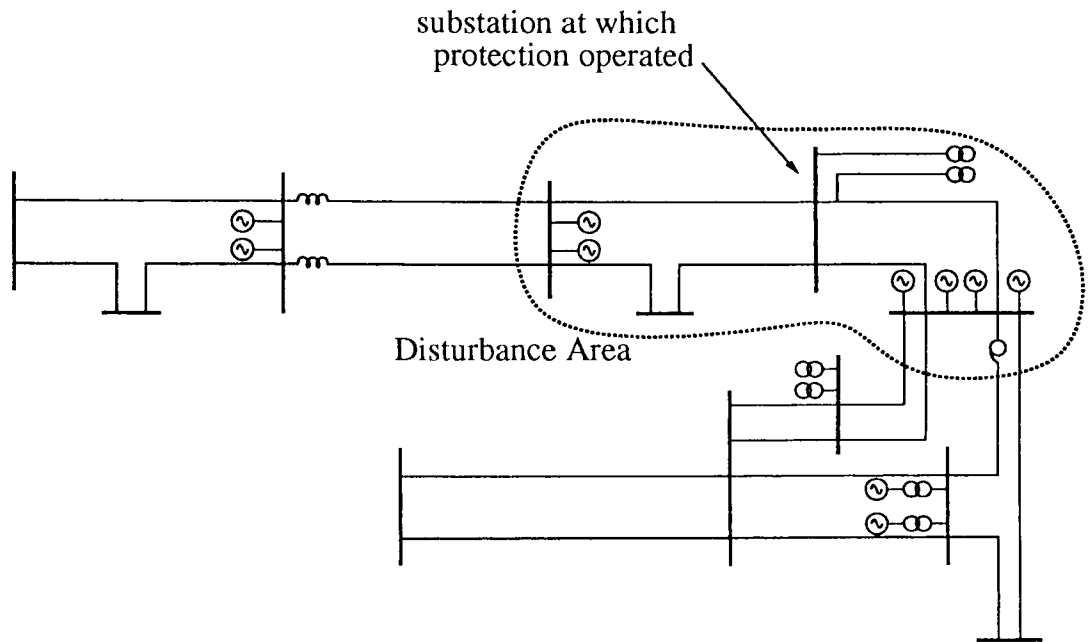


Figure 4.5: A 1 substation radius disturbance area

tection schemes are designed with remote backup capability. As a result, non-unit relays are set to detect faults further than one substation away from the relaying point. The disturbance area radius must cater for the distance from the protection's substation to the fault, and the distance on the other side to another non-unit protection since the fault can be fed from both sides in an interconnected network. It would be possible to use protection models, however a quick and general derivation is desired. Figure 4.6 illustrates how using standard 80%/150%/225% reach settings in a three zone distance scheme requires a radius of four substations. Overcurrent protections could likewise have a similar reach. It should be noted that such occurrences are unlikely since they necessitate several simultaneous failures. Also, in an interconnected network, such long reaches would be uncommon due to the effect of tie lines on fault currents, and the problem of coordination.

The disturbance area radius is set in the gut reaction production rules and so can be easily adjusted to suit. If unit protection operates, then the fault will

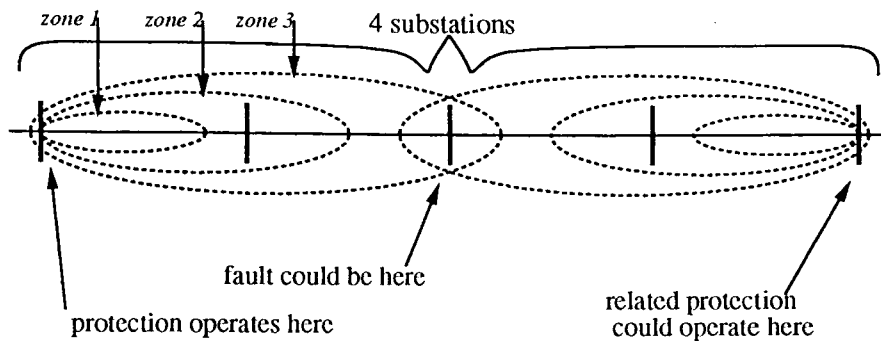


Figure 4.6: The disturbance area of a three zone distance protection scheme on a simplistic series of lines

be within one substation away from the relaying point. Therefore the radius is less. This discrimination is used by the gut reaction. Two production rules set the radius at one of two levels based on whether or not the alarm relates to a unit protection operation. Existing disturbance areas also limit the scope of a new disturbance area.

As an example, the gut reaction may resolve the following knowledge

```

judgement(isnt control).
judgement(present_alarm is operated(p1 at ss1)-protection).
judgement(present_time is details(1990,1,11,15,41,37.0)).
judgement(present_ss is ss1).
judgement(primary_alarm).
judgement(isnt old_disturbance).
judgement(new_disturbance).
judgement(present_disturbance is-199999).
judgement(disturbance_area_radius is 3).
judgement(isnt wrong_msge).
judgement(isnt definitely_expected).
judgement(isnt maybe_related).

```

4.6.2 Evolution of Hypothesising Strategies

There are largely two roles fulfilled by the central diagnosis. Originally, hypotheses are generated and their validity immediately tested in order to create scenario facts. Secondly, the central diagnosis performs subsequent testing on existing sce-




Figure 4.7: A schematic of the first hypothesising strategy

becomes active.) Similarly two scenarios are disproved on the receipt of a_3 and, of the remaining scenarios, two are invalidated when a_4 is received. Two scenarios predicted alarm a_5 , while three scenarios did not involve this alarm. This strategy then continues to generate and simulate hypotheses which only account for the latest alarm—this is indicated on the figure by the new origin O' . The twenty six resultant partial scenarios are then combined with the three old scenarios to form seventy eight complete scenarios (the propagation of all of these scenarios is not shown in the figure for clarity). Figure 4.7 illustrates how every one of the eighty scenarios existing after alarm a_5 predicted alarms a_6 and a_7 .

According to this strategy, many hypotheses are produced from the start, and many competing hypotheses are maintained.

4.6.2.2 Incremental Maximum Hypothesised Failures

The maximum number of hypothesised failures which may be included in hypotheses is increased by one when an unexpected alarm is received. Only at this point are new disturbance hypotheses generated. This strategy can be shown schematically as in Figure 4.8, which relates to the same case as Figure 4.7. The




Figure 4.8: A schematic of the second hypothesising strategy

figure illustrates how the strategy suppresses the generation of as many scenarios from the beginning. For example, if the protection, switchgear and telemetry acted perfectly, then the diagnosis would progress with only no-failure scenarios. The maximum failures limit of the disturbance in the illustrated case is maintained at zero until alarm a_5 is received. Being unexpected, alarm a_5 results

in the disturbance's maximum failure limit being incremented. All possible new hypotheses are then generated and simulated, resulting in twenty eight partial scenarios. Twenty six of these are then combined with the three old scenarios to form seventy eight complete scenarios. (Some of the lines showing the propagation of these scenarios have been left out of Figure 4.8 for clarity.) In addition, two of the created partial scenarios account for alarms $a_1 \dots a_5$. These therefore form complete scenarios in themselves, as illustrated in Figure 4.8.

Since the number of failures is only incremented by one on the receipt of the fifth alarm, the strategy considers fewer hypotheses in total than did the previous strategy. However, many hypotheses are again generated and many competing hypotheses maintained.

4.6.2.3 Derived Hypotheses

When an unexpected alarm is received, the maximum number of hypothesised failures which may be included in hypotheses is again incremented. However, this strategy seeks first to develop existing hypotheses before seeking new hypotheses. The existing hypotheses have a failure added to them. The candidates for failures are those devices with predicted activity in the existing scenarios. Figure 4.9 demonstrates the action of this strategy on the same example as Figure 4.7. Again the simultaneous failures limit is zero for the first four alarms, with the result that only three scenarios, two fault-based and one based on a false alarm, are in existence. On receipt of alarm a_5 , the two existing fault-based scenarios have a failure added to them in an attempt to form hypotheses which also account for a_5 . Two scenarios are successfully created by this means. In addition, a partial scenario detailing alarm a_5 as false is generated, and combined with the three old scenarios to form three complete scenarios.

This evolved strategy can significantly reduce the number of hypotheses gen-




Figure 4.9: A schematic of the third hypothesising strategy

erated and maintained. In addition to the above description, however, the system needs a way to go back to previous alarms in order to generate different hypotheses. This facility would be needed to correctly diagnose the above alarm stream if a_5 actually related to a different fault. The beneficial improvements in this last strategy make it more suited to real use. The complete central diagnosis using this strategy will now be described.

4.6.3 Central Diagnosis

The central diagnosis reasoning is used following the receipt of an alarm describing automatic activity.

4.6.3.1 Original Hypothesis and Scenario Generation

When the gut reaction has judged an alarm to belong to a new disturbance, then the knowledge in the hypothesis and scenario generation rules is used. As has been previously mentioned, the hypothesiser is primarily concerned with protection relay operation. The alarm detailing this primary activity could, however, be lost by a telemetry malfunction. Therefore Prolog rule clauses cater for the case where a secondary activity alarm has caused the creation of a new disturbance area: the rule body determines a possible controlling protection relay and then recursively tries the hypothesis and scenario generation rules with the protection relay operation in the argument list, in place of the actually received secondary alarm.

The clause of the hypothesis and scenario generation rule for primary protection relay operation alarms, controls the main hypothesis and scenario generation inference. The rule body first uses a protection modelling rule to determine the possible faults which could have caused the relay to operate. The modelling rule ignores all other protection relays; it simply determines the location and generic type (phase or ground) of detectable faults. No new failures are added since there is no reason at this stage to believe that maloperations have occurred.

A detectable system fault is taken from the list. Firstly, a fault duration term is added to complete the description of the fault. Delayed autoreclose (DAR) schemes should be able to discriminate between temporary and permanent faults. Therefore, if the protection relay on which the hypothesising is based belongs to a DAR scheme, a 'permanent' fault and then a 'temporary' fault are hypothesised; a 'permanent' fault is taken first, while the 'temporary' fault is selected on backtracking. Otherwise, only one '(permanent or temporary)' fault is hypothesised.

The fully formed (including location, generic type and duration) hypothesised

system fault can be qualitatively simulated. The effect of the latest alarm on the knowledge base is first hidden in order to take the system back to the state it was in when the fault occurred. The removal of the effect of the alarms is achieved in the same way as simulated changes: facts are created which are treated as an overlay above the system data, rather than changing the representation of the actual system status. The qualitative simulation then obtains the system response to the situation where the hypothesis is a genuine fault. The simulated activity includes the expected and the possible expected activity. Also, any activity which has an alternative, and was chosen for no other reason than because it was the default, is recorded as a default with its alternative. The actual reasoning of the developed qualitative simulation knowledge is described in section 5.4. The possible consequences of the simulation are threefold. This is the immediate (as distinct from the subsequent) testing of hypotheses.

- fault not isolated:- Faults should never remain unisolated in any real power and protection system. However, the qualitative simulation can sometimes reach this conclusion because of insufficient system details. Either, certain protection equipment has been omitted from the knowledge base, or the faulted plant is so near the edge of the defined network that isolation would involve switchgear in the undefined part of the network. The fact that this occurred is simply stored in the knowledge base with the fault description.
- genuine activity not included:- In some cases the simulated activity will not include the actual activity (on which the hypothesis is based) due to the operation of protection nearer the fault location or of higher priority protection. This arises because it is detectable, and not protected, faults which are hypothesised and simulated. The genuine activity would only be observed following the hypothesised fault if some equipment failed. The

qualitative simulation generates “blocking” facts describing activity which it suspects may have been simulated instead of the actual activity. Then if the hypothesised fault does not explain the genuine protection operation, the devices in the blocking facts are asserted as candidates for profitable hypothesised failures. These may be used for derived hypothesis generation.

- genuine activity included:- The previous conclusions cause the inference engine to backtrack for another fault hypothesis. However, if the genuine activity is included and the hypothesised fault is isolated in the simulation, then the reasoning proceeds. The activity involved in the successful simulation is classified as supportive or predicted depending on whether a corresponding alarm has been received or not. A partial scenario is then formed from the classified simulated activity, the hypothesised fault and any extra information obtained from the simulation.

All scenarios must cover the whole disturbance area. Since the case presently being considered is that of a new disturbance, there are no existing scenarios, and so the partial scenario details are copied into a complete scenario fact which is stored in the knowledge base. Backtracking is then started in order to find another hypothesised fault.

Once all such hypotheses have been simulated, a false activity hypothesis is created. The hypothesis that the device operated wrongly is simulated to obtain the secondary activity which should follow from the observed activity. A partial and full scenario are generated in the same way as for a hypothesised fault: the partial (and hence full) scenario includes the classified simulated activity and the false activity.

The same hypothesis generation rule is sometimes used following an alarm in an old disturbance. Following the subsequent testing and the hypothesis derivation, there may still remain old scenarios which have not been updated. If this

is so then the original generation knowledge is used to generate new original hypotheses specifically for the new alarm. The rules for obtaining complete scenarios from partial scenarios have different clauses for the new disturbance case and the old disturbance case. Thereby the partial scenarios generated based solely on the latest alarm, have the old scenario's details included before being asserted as a complete scenario: multiple events are added. For each partial scenario, a new scenario is generated for every remaining old scenario. Again a false alarm hypothesis is simulated, and the resulting partial scenario likewise combined with all old scenarios to generate new complete scenarios.

Hypotheses for new disturbances have no failures since the maximum failures limit starts at zero. However, an old disturbance may have a non-zero failures limit due to previously received alarms. A successful partial scenario is used to generate sub-hypotheses: devices whose activity is predicted are candidates for failures. Successful sub-hypotheses' partial scenarios are likewise used to generate further sub-hypotheses by the addition of another failure. This is repeated until the limit on the number of simultaneous failures is reached.

4.6.3.2 Subsequent Testing

As the operation of the power system is continuous, the process of fault diagnosis is an on-going one. Consequently, diagnoses may be performed before all the relevant alarms are received. The hypothesising philosophy involves most of its reasoning effort before the majority of related protection and switchgear activity is known to have occurred. Therefore, as well as the immediate testing of hypotheses when they are generated (as per the generate-and-test paradigm), subsequent testing plays an important role. Subsequent testing updates scenarios following alarms subsequent to the alarm which invoked the generation of the scenario.

When an alarm has been judged by the gut reaction to belong to an old

disturbance, then all that disturbance's existing scenarios are tested. There are four tests which may be performed on a scenario. They are attempted in turn until one succeeds or all fail, for each scenario.

- activity predicted:- The scenario was generated on the assumption that the newly alarmed activity would occur. It was classed as predicted activity within the scenario. The scenario is replaced with an equivalent new one, except for the reclassification of the newly observed activity as supportive.
- active device was supposed failed:- The scenario includes a device as a failure, and that device has now been active. The developed hypothesiser strategy suppresses the inclusion of failures until there is evidence which is thought to support it. Therefore there must have been activity that indicated that the device did not act correctly. Therefore, the scenario is not removed as invalidated, but is replaced with a new scenario. This new scenario is basically the old scenario with any secondary activity which may be caused by the latest activity as predicted, and with knowledge that the device activity was slow.
- activity possible:- The scenario has some activity classified as possible, rather than definite (predicted), which includes the observed activity. An example of the use of this is a 'temporary' fault scenario. Delayed autoreclose equipment may clear the fault after the first attempt, or after subsequent trip/reclose cycles. The corresponding scenario therefore includes more than the first trip/reclose cycle activity as a possibility until a member of the list is received, when the rest of the activity becomes predicted. In general then, such scenarios are replaced with similar scenarios with the group of possible activity removed, the received activity classed as supportive and the other activity in the group included as predicted.

- activity was alternative to default:- The scenario has the activity listed as an alternative to a default (ie. an assumed) operation. The requirement for this feature stems from the uncertainty which sometimes marks the behaviour of some schemes. For example, in an intertrip feeder protection scheme, when protection at one end operates, the remote end protection and the intertrip signal “race” one another. That is, the remote end circuit breaker may trip because of the operation of the remote end protection relay or that of the intertrip receive relay. In the absence of additional details, the generated scenario therefore may include the protection operation as the default, and the intertrip received as a possible alternative. When alternative activity in a scenario is received, the scenario is recreated in the light of the most recent activity. The old scenario is removed and regenerated, including the newly received activity rather than the default.

4.6.3.3 Derived Hypothesis and Scenario Generation

Hypotheses are not only generated originally from the relay protection model. As has been mentioned, a scenario may need to be recreated if the alternative to default activity is observed. Secondly, a hypothesis may be derived on the basis that the latest alarm is evidence of a device failure. Thirdly, a hypothesis may be derived based on the failure of a previous qualitative simulation.

Having identified a scenario whose alternative activity has occurred, the actual fault and failure hypothesis to which this activity pertains is identified—the scenario may be a multiple fault scenario. That single fault is resimulated including any failures previously included.

Before repeating the qualitative simulation, the knowledge base representation of the network has to be taken back to the state the network was in when the fault occurred (ie. the time of the alarm which first caused the scenario to be

created). This is the same requirement as for original hypotheses. However, it is more important for derived hypotheses and scenarios as several operations might have occurred since the first related alarm. The expert system hides the effect of the alarms received since the suspected occurrence of the fault. If this were neglected, the resulting partial scenario could be wrong: an incomplete qualitative simulation would result if switchgear has isolated a path for fault current which initially fed the fault.

Once this preparation is done, the hypothesised fault and failures are simulated as before—the simulation automatically uses the knowledge of the latest alarm to select the alternative-to-the-default activity, and therefore not regenerate exactly the same (wrong) alarms. The partial scenario is created from the hypothesis, the classified simulated alarms and any extra information generated by the qualitative simulation. The newly generated partial scenario is the new scenario if the original scenario contains only one hypothesis. If, however, the old scenario also includes other (fault or false) hypotheses, then the unaffected hypothesis details must be preserved. All but the affected hypothesis' old details in such a scenario are added to the new partial scenario details. The resulting new complete scenario replaces the old scenario. Old scenarios which contain exactly the same affected hypothesis, but have different unaffected hypotheses, are likewise replaced by a scenario containing their unaffected hypothesis details and the same partial scenario details.

Following the subsequent testing of an alarm, including dealing with alternatives to default activity, a further attempt may be made to generate a derived hypothesis and scenario. A derived hypothesis is generated on the basis that the most recent relay operation alarm is evidence that another device has failed. Again, if the backup operation alarm is lost, a possible primary relay would be identified from the unaccounted secondary alarm. The derived hypothesis is

based on an old scenario which remains after the subsequent testing, and which contains a hypothesised fault which can be detected by the newly operated protection relay. The new hypothesis is derived by adding another failure to the old hypothesis. The disturbance's maximum number of simultaneous failures is increased by one, and this applies to all subsequent scenarios in the disturbance. Devices with predicted activity in the old scenario are candidates for the new failure. Each possible hypothesis derived in this way is simulated, following the reversal of the effect of alarms received since the suspected occurrence of the fault. There are three possible results from the simulation, in exactly the same way as for the no-failure simulation, and they are treated in the same manner. If a partial scenario is successfully generated, then it is added to the details of the old scenarios' unaffected hypotheses.

Finally, it was mentioned in the description of the original hypothesis generation, that not all hypotheses' simulations will account for the genuine activity. (This applies equally well to the new disturbance and old disturbance situations.) In certain cases the reasoning may fail to generate any partial, and hence complete, scenarios. This arises when a backup operation alarm has been received and the necessary number of failures has not been included. Those qualitative simulations which did not account for the genuine activity, will have generated facts detailing the attempted hypotheses and the possible blocking devices. Blocking devices are those relays and circuit breakers whose correct behaviour could have been the reason that the genuine activity was not simulated. Such devices, which have not yet been operational (ie. would have been classified as predicted if a partial scenario had been formed) are potentially profitable failures. If the absolute limit on the number of failures would not be exceeded, these failures are added to the unsuccessful hypotheses, and the resulting hypotheses simulated. The disturbance's number of simultaneous failures is increased by one. Partial

and complete scenarios are attempted as before.

.

4.7 Dealing with Uncertainty

4.7.1 Importance in Diagnosis

It was seen in chapter 2 how dealing with uncertainty plays a significant role in many problem-solving areas. This is so for the on-line fault diagnosis of power system networks. Since it is performed on-line, the diagnosis may well be lacking evidential data. The diagnosis may fail to generate only one conclusion; some other expert systems generate one conclusion, but do not “know” if there are any other possibilities. What is desired is a diagnosis which is robust, but also discriminating: a robust diagnosis should not just fail, and a discriminating diagnosis should not just replace a list of alarms with an unhelpful long list of possible conclusions.

4.7.2 Sources of Uncertainty

One thing that should be noted concerning the fault diagnosis of power system networks is that it is not an inadequacy of the diagnosis which results in there being more than one conclusion. There are several reasons why there can be legitimately more than one possible answer.

- alarms lost:- Telemetry system malfunctioning can cause an alarm to be lost. As a result, when some activity is expected but has not been alarmed, it is uncertain whether the activity has failed or the alarm has indeed been lost. Performing the diagnosis before the receipt of all relevant data results in the same effect.
- limits of SCADA system:- Most, if not all, utilities have not wired all relays into their SCADA systems. The resulting incomplete evidence has the same influence on the diagnosis as does telemetry malfunctioning.

- protection relay characteristics:- Relays can correctly detect a range of faults. An operated relay may have detected one of various fault types and locations. Indeed, the protection scheme may be designed such that one relay protects more than one item of plant.
- protection relay failure:- Relays can fail. Most of the relays installed in power systems are not microprocessor based and so do not have self-checking facilities. Consequently, it is often difficult to distinguish between correct and incorrect non-operation: for both, no alarm is received.
- remote backup:- Although remote backup operation fulfills an important role in the protection of power systems, it can result in large areas of the network being isolated. Isolation by remote backup activity, can indicate that a fault occurred on one of many possible items of plant. The simultaneous operation of more precise protections, which would enable the narrowing down of the possible fault locations, does not always occur.

4.7.3 Problem Factors

One of the major problematic factors is the size of the system to be diagnosed. The amount of power network, protection and switchgear equipment involved in power systems means that the volume of data required by any technique for dealing with uncertainty is an important factor. Whatever technique is adopted, a significant number of values will need to be set.

The nature of the system to be diagnosed causes problems for some techniques. The diversity of equipment involved in power systems and the diverse ways in which equipment is configured, restrict the use of simplifying generalities which could reduce the volume of necessary data. The diversity of behaviour involved is also a source of increased data requirements. Together with the possibility of multiple faults, the total number of possible diagnosis conclusions is colossal.

This would invalidate those techniques which require a complete definition of all possible conclusions (like the Dempster-Shafer theory).

Furthermore, sound statistics are not available. The power system operator, who is responsible for producing fault reports, has many other duties to perform following a fault. Consequently, the quality of statistics produced on the basis of operator logs is questionable. Moreover the rarity of fault events, and the small number of large power systems means that the expertise to accurately set objective values is not common.

4.7.4 The Developed Measure - Relative Likelihood

4.7.4.1 What it is

The measure developed is relative. The absolute value is not the critical feature, rather it is the value in comparison with values for competing aspects. The measure has been called likelihood because the values are subjective, are not calculated from statistics, and the confidence in their validity is not sufficient to merit being called probability. However, like probability, a value range of 0 to 1 is used. A relative likelihood value cannot be set for every possible scenario because of the number of possibilities which exist. Relative likelihoods are instead set for possible aspects of scenarios, and are combined to form values for specific scenarios. Values must be set for possible system faults, maloperations and telemetry malfunctions.

4.7.4.2 Why Use it

There are mathematically based techniques, like those described in section 2.4, which could have been used instead of this relative likelihood measure. Although these techniques are statistically sound, it must be remembered that only poor statistics are available for use. The assumptions on which these techniques are

based, when compared with the real behaviour of power and protection systems, would indicate that the selection of the techniques on the merit of their mathematical soundness would be unjustifiable. Moreover, these techniques would appear to be unmanageable due to the volume of data required, the computational complexity involved, and the processing time which would be required. As the uncertainty measure is to be used to judge conclusions rather than in the derivation of conclusions, it is questionable whether some of these techniques are needed. Although it was decided to develop a simplified strategy, the validity of these more established theories has not been rigorously disproved.

Another alternative is to use a primitive integral count of certain aspects. The use of a technique like the unverified evidence quantification method (see section 2.4.2.6) would be reasonable in terms of data requirements and on-line time constraints. The review of FAUST (section 3.4.5) indicated one of the problems associated with this method: how to make the distinction between different types of unverified data. In addition, there is more knowledge available than would be used in such a technique.

Given these shortcomings in the alternative methods, the question may still be asked why relative likelihood should be used. Using relative likelihood is a compromise between adopting either of the above two sets of alternatives. It does not require sound statistics, and uses knowledge whose availability may be reasonably expected. It is not mathematically sound, but does give a reasonable comparative indicator. By permitting the identification of unlikely alternative conclusions, relative likelihood provides the means of employing a display cut-off in order to restrict the number of possible answers which need be shown to the engineer.

4.7.4.3 How to Use it

Initially it was thought that a value would need to be set for a potential aspect of a scenario in relation to all other aspects. That is, complete interdependence was assumed. However, it was realised that comparison was only required between co-existing scenarios, and hence between co-existing scenario aspects.

Figure 4.10 illustrates the relationship of scenario aspects in the context of relative likelihoods. The “1ary” and “2ndary” items referred to in the figure

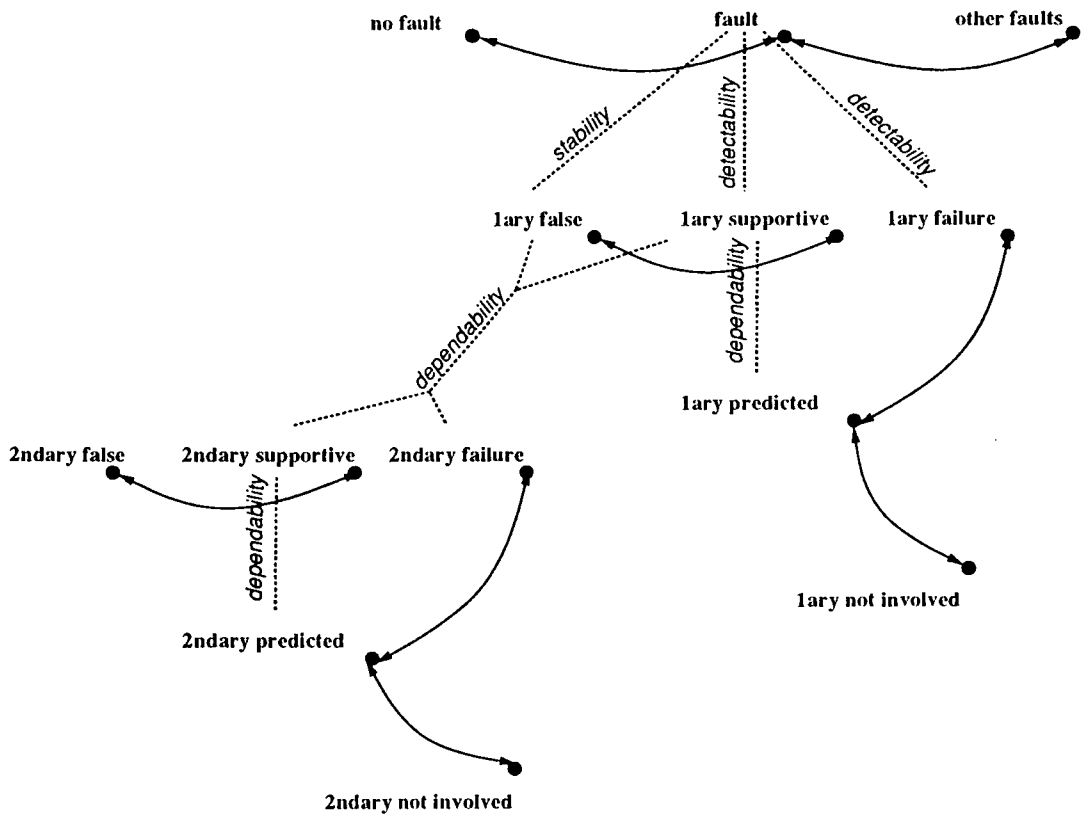


Figure 4.10: Schematic of the relative likelihood interdependence

indicate aspects which relate to primary and secondary activity respectively. The broken lines set the aspects of scenarios in the context of their corresponding behaviour. This interdependence is not, however, used to set relative likelihood values. It would only be used if the uncertainty value was required to indicate the progressive change of scenario uncertainty over time as more alarms are received. This is outwith the purpose of the present uncertainty measure. The full line arcs

indicate the aspects of a scenario which can co-exist. These, therefore, indicate the scenario aspects whose relative likelihood values have to be set in relation to one another. The “no fault” aspect and the two “not involved” aspects are set to absolute values which do not influence any combined value, ie. they perform an identity function. Without these aspects, other factors would influence scenarios. For example, since primary or secondary failure usually leads to backup operation, the resulting additional predicted activity would further influence the scenario’s relative likelihood.

In order to make practical use of these values, they must be combined to form a measure by which complete competing scenarios may be compared. The method of combination should be such that the effect of combined relative likelihoods mirrors the effect of individual aspects of a scenario on the whole. This effect can be described by the following observations:

- Scenarios may include independent faults, all of which influence the likelihood of the scenario.
- Scenarios may include related causal events which influence the likelihood of the scenario.
- Competing scenarios may have varying proportions of predicted and supported alarms.

These aspects of system behaviour show how every aspect of a scenario influences its total likelihood. This indicates that the selection of the least likely aspect to represent the whole is unreasonable. The use of a minimum function, as in fuzzy logic (equation 2.4.15), is therefore unsuitable. Reasonable results were given by combining relative likelihoods by the product rule, as are independent probabilities. (The “no fault” and “not involved” aspects are therefore set to 1.0 in order to act as an identity function.)

For each aspect of a scenario, a typical value may be combined with (ie. multiplied by) a particular value. If no particular value is supplied, only the typical value is used. The inclusion of a particular value is a means of including knowledge of plant or equipment which is, for example, less reliable than other similar plant or equipment. A learning influence can thus be incorporated. The product of values for each aspect in a scenario, gives a comparative measure. The actual data required will be described in a section of the next chapter.

4.8 Ideas in the Deeper Diagnosis

The knowledge contained in the “deeper diagnosis” would allow further reasoning to be made of the disturbance. It would make use of the conclusions of the “central diagnosis”, that is, system faults and maloperations. To this would be applied deeper modelling knowledge of plant and protection, knowledge of related disturbances and additional knowledge to deal with uncertainty.

Whereas the two tier plant item-subunit representation is sufficient for the likes of busbars, the complexities of transformers and generators deserve better representation. Better identification of the disturbance and of any maloperations could then be made. Further details of protection behaviour could be included. For example, the operation and alarming of phase specific protection should add the faulted phase to the diagnosis. Likewise the behaviour of gas or oil related protection can be used to locate the fault more accurately, such as distinguishing between an internal transformer fault and an external bushing fault. The characteristics of less definable protections, like generator negative phase sequence protection would allow the identification of more complex aspects of network behaviour.

The “deeper diagnosis” would include knowledge of complex behaviour, such as the effect of surges following line faults in causing busbar protection to operate, the occurrence of flashovers in substations following line faults, and the occurrence of faults in the common area between adjacent protections. The “central diagnosis” would continue to generate multiple system fault or system fault plus false operation scenarios for complex disturbances. The “deeper diagnosis” knowledge would be used to group together apparently discrete system faults in a cohesive disturbance explanation.

Some aspects of network behavioural uncertainty cannot be satisfactorily characterised numerically. The “deeper diagnosis” could include knowledge which

would recognise certain disturbance patterns which are in fact more likely than they might appear. External influences would generally be the cause of this effect. For example, in a lightning storm the occurrence of multiple faults is not uncommon, and therefore multiple temporary line fault scenarios should be ranked more highly than their combined relative likelihood might suggest. In a similar way, the nearness of genuine system faults could be used to adjust the likelihood of protection/switchgear equipment maloperations.

Chapter 5

Network Knowledge

5.1 Significance of Network Knowledge

A significant part of the work reported in this thesis is the level of network knowledge representation and use. This expert system makes a clear distinction between network knowledge and diagnosis knowledge whereas some fault diagnosis systems have been developed in which the network knowledge and diagnosis knowledge have been almost synonymous. The separation of such knowledge affords modularity and increases the system's portability to other power systems. By the same means, the expert system can be more easily maintained in line with modifications and additions to the power system being diagnosed.

The scope of the network knowledge incorporates details of the power system (ie. data) and behavioural knowledge of systems in general. Both of these have been kept distinct since the behavioural modelling is applicable to any power system. In order to make use of the behavioural knowledge, details are required of the system to be diagnosed. The network data can be broadly categorised as dynamic and static data. The dynamic data is accessed through the SCADA system and is the basis of the diagnosis. Details of the power network and the protection and switchgear equipment are largely static, and are related to the dynamic data by telemetry data. Relative likelihood data is included in order to manage uncertainty. Through the application of a loadflow algorithm to static network data, further static data is derived; the derived data relates to the normal direction of power flow. The use of the loadflow algorithm in cooperation with the expert system is a demonstration of the important area of interfacing knowledge based and algorithm based modules. As the hypothesiser needs to be able to predict the alarms which would follow a system fault, the behavioural knowledge includes qualitative simulation knowledge. The qualitative simulation can determine the activity caused by a system fault. Other Prolog logical rules model the behaviour of classes of network and protection equipment.

5.2 Dynamic Data Requirements

Clearly, an on-line fault diagnosis expert system requires access to dynamically changing data which describes, in some way, the behaviour of the system to be diagnosed. This expert system uses Supervisory Control And Data Acquisition (SCADA) system alarms as the basis of its diagnosis. The alarms are treated as symptoms of the fault to be diagnosed, and are also used to keep the network representation up-to-date.

The fault diagnosis is based on the activity of protection and switchgear equipment which is alarmed and conveyed to the control centre by the SCADA system. Use is made of alarms concerning protection relay operate and reset, trip relay operate and reset, circuit breaker trip and reclose, and isolator open and close. At the control centre, the alarms can be accessed directly, such that the expert system has minimum dependence on the power system operator.

The development of the expert system was based on knowledge of the ScottishPower transmission system. A typical alarm in this system might be:

13:45:13:00 DUMB2	COAT2	FIRST MAIN PROT OPTD	ON	DUMB2
-------------------	-------	----------------------	----	-------

The fields in this alarm relate respectively to the alarm time tag, the source substation (dumb2), the circuit or feeder (identified by remote end substation, coat2), a legend identifying the operation of that circuit's first main relay, the new status of the relay (on), and finally a repeat of the source substation.

ScottishPower's SCADA messages, unlike those of many utilities, are time-tagged at the outstation. The times tagged to alarms are consequently much closer to those of the original activity. However, the reliability of the alarm time tags is not depended upon. The times are only used to relate or dissociate an alarm from a disturbance: if the difference between the alarm time and a causal event time (ie. the time tagged to the first alarm judged to relate to the event) is above a limit then the two events are adjudged to be unrelated. The diagnosis

does not concern itself with the exact order within narrow time scales in which alarms are received. It can therefore be applied to less developed systems.

Most of the diagnostic reasoning is related to relay operations and circuit breaker tripping. With the exception of straightforward automatic reclosing schemes, automatic recovery systems have not yet been included. Isolator activity and reclose and reset alarms can be included in scenarios and are used to update the internal system database, but are not as yet used to their full potential in the diagnosis.

5.3 Static Data Requirements and Structure

5.3.1 Scope

The static data requirements can be very extensive for typically large transmission systems. The data incorporated covers a description of power system plant and of protection and switchgear equipment. This must be sufficient to maintain an accurate representation of the network topology and to draw inference from protection activity. Relative likelihood values are required in order that competing scenarios can be ranked. Some details pertaining to the SCADA system must be provided in order that alarms can be read and related to the activity of protection and switchgear equipment.

5.3.2 Structure of Power System Data

In the early stages of development, no structure was used to organise power system data. The network topology was represented by a nodal network, and relays, switchgear and network plant were simply identified by a unique name. While this was satisfactory for the small test systems being used at the time, its limitations arose when applied to realistic power systems.

Power systems were recognised to have an inherent hierarchy. Connections to neighbouring power systems relate to the whole system. The local system can be described in terms of substations in the system. Each substation consists of various items of plant, protection equipment and switchgear equipment. Finally, certain specific details describe each item of plant, and each protection and switchgear device.

As a result, network, protection and switchgear equipment is identified by a substation name, and a local name which is unique in that one specific substation. (Many utilities use such a labeling system, having many substations with breakers

labeled, say, 'x120'.) For example, the 'x120' breaker and the 'r1' relay in substation 'ss1' would be identified by the Prolog terms 'x120 at ss1' and 'r1 at ss1' respectively. Terms such as 'x120 at ss1' are not simply text strings: 'at' was defined as a Prolog infix operator, while 'x120' and 'ss1' are constants (text strings). Therefore, the substation and local names can be easily obtained from complete device names.

A second observation was made which influenced the structure of network data. It was noticed that, without focusing down to the level of physical fault location, some protection relays can discriminate between faults within items of plant. Transformers are the most obvious example of this: often, different earth fault protection relays protect the high voltage (HV) and low voltage (LV) windings.

This discrimination between parts of items of plant motivated the development of a lower "subunit" level for the plant data structure. The complete hierarchy for power system plant consists of a substation identifier, a plant item identifier and a subunit name. For example, the HV winding of a transformer 't1' in a substation 'ss1' would be identified by the Prolog term 'ss1 :: t1 \\ hv_wdg'. ('::' and '\\\'' have been identified as Prolog infix operators.)

Lines and cables between substations also fit into this structure despite not being restricted to one substation. The substation part of the data structure describes the interconnected substations. A Prolog infix operator 'to' has been defined to relate interconnected substations. The adopted structure would describe a line, say, 'l1', between substations 'ss1' and 'ss2', with the term 'ss1 to ss2 :: l1'.

This developed hierarchical structure has several advantages over the nodal representation of power system equipment. Firstly, the top part of the structure (ie. location in substations) is natural to power systems, and most utilities will

use similar structuring in their labeling scheme. Also, SCADA messages generally involve a source substation identifier in addition to legends, device names, etc. The use of the lower level of the hierarchical structure (ie. subunits) allows more knowledge to be incorporated into the expert system, as shall be seen in section 5.6. An example of the knowledge included is that which limits the reach of an earth fault relay to one winding of certain transformers. As well as allowing more detail to be gone into, the data structure provides a graduated focusing means: generally, a search for devices which satisfy some condition can be restricted to certain substations.

5.3.3 Plant Data

The details required to define an item of plant are dependent on the plant type. The required data shall be described in the following sections. In addition, details of the relational data between items of plant will be given.

5.3.3.1 Plant Types

The behavioural knowledge contained in the expert system has details for various types of electrical plant. The detailed plant types are transformer, three-power-winding transformer, autotransformer, busbar, junction, line, three terminal line, cable, reactor and machine. Additional data defines busbars which are supply points from an undefined part of the network or from an adjacent network, and loads on busbars. These are illustrated in Figure 5.1 by the busbars at substations *ss1* and *ss2* respectively. In addition a slack busbar is defined. Any particular network can be defined by these details.

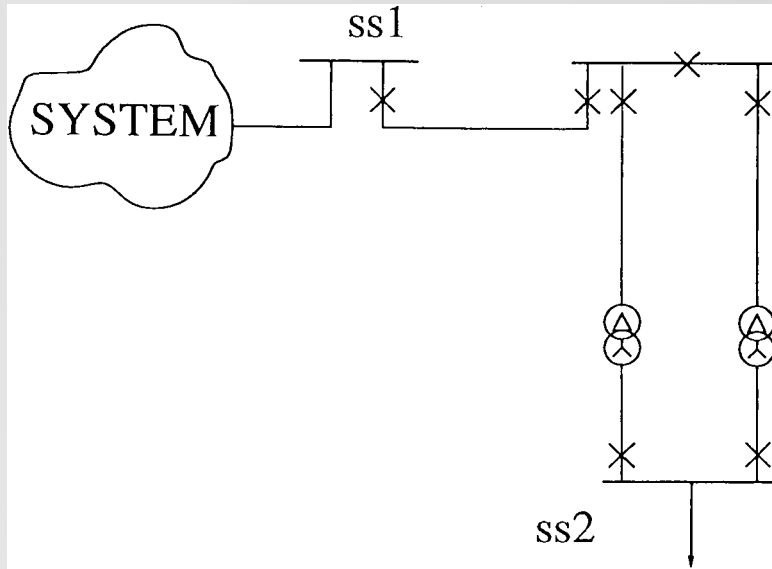


Figure 5.1: A portion of network showing a supply point busbar (at ss1) and a loaded busbar (at ss2)

5.3.3.2 Plant Template

All definitions of plant types conform to the same structure. This may be represented graphically as shown in Figure 5.2. The two tier representation of items




Figure 5.2: The general template for a type of plant item

of plant which is evident consists of a device level and a subunit level. At the device level there are details which relate to the whole plant item. In addition, the subunits, represented in Figure 5.2 by the small squares, may have specific details. Both the device and the subunit level data is specific for the type to which the item of plant belongs. (The interconnection of subunits is defined by the plant type modelling.)

5.3.3.3 Transformer Details

The device level details for a transformer relate to the number of phases (1 or 3) and whether or not there is a tertiary winding. The other details are specific to a particular subunit. A transformer may have between two and four subunits: two windings and no, one or two winding neutral connections. These subunits are identified as: 'hv_wdg', 'lv_wdg', 'hv_ntrl' and 'lv_ntrl'. The subunit level details cover three phase winding arrangement (star or delta), winding neutral connection, winding voltage and winding impedance. The following Prolog fact describes transformer 't1' at substation 'ss1', a typical three phase star/delta transformer with no tertiary and with per unit winding impedances of $0.003 + j0.11$ (only positive sequence impedance values have been included in the expert system):

```
transformer(  
  ss1,  
  t1,  
  3,  
  hv(star, ntrl_con is earth, 400),  
  lv(delta, no, 23.5),  
  no,  
  imped(z(0.003, 0.11), z(0.003, 0.11))  
).
```

5.3.3.4 3-Power-Winding Transformer Details

The definition of a 3-power-winding (double LV) transformer is very similar to that of a normal transformer (with two power windings). This description in terms of the number of power windings does not include the tertiary winding. The details are the same as for normal transformers except that there are more windings. The possible subunits are identified by 'hv_wdg', 'lv1_wdg', 'lv2_wdg', 'hv_ntrl', 'lv1_ntrl' and 'lv2_ntrl'. A Prolog fact which describes a three

phase star/delta/delta 3-power-winding transformer which does not have a tertiary winding is shown:

```
pw3transformer(  
    ss2,  
    pw3t1,  
    3,  
    hv(star, ntrl_con is earth, 275),  
    lv1(delta, no, 11),  
    lv2(delta, no, 11),  
    no,  
    imped(z(0.003, 0.11), z(0.003, 0.11), z(0.003, 0.11))  
).
```

5.3.3.5 Autotransformer Details

All details, except HV and LV voltages, relate to the whole device. These describe the number of phases, the existence of a neutral connection and of a tertiary winding, and the impedance between the HV and LV winding terminals. The possible subunits are 'hv_wdg', 'lv_wdg' and 'ntrl'. The following is a Prolog fact describing a three phase autotransformer which has no tertiary winding:

```
autotransformer(  
    ss3,  
    sgt1,  
    3,  
    ntrl_con is earth,  
    no,  
    400,  
    275,  
    imped(z(0.001, 0.12))  
).
```

5.3.3.6 Busbar Details

The device level details simply list the subunits which form the busbar. The subunits are the busbar sections. The distinction between different busbars and

different sections of the same busbar is made on the basis of how they may be isolated from one another. The convention used is that two busbars have only automatic switchgear (not manually operated isolators) between them. A fact describing a busbar is shown:

```
busbar(  
    ss4,  
    b2,  
    [s1,s2,s3]  
).
```

Unlike other types of plant, a busbar may have some subunits which do not have a connection outside the plant item. Also, the subunit names are not fixed, being specified in the busbar fact.

5.3.3.7 Junction Details

There are parts of power systems which are not busbars, but have similar functional characteristics. That is, they are multi-terminal nodes, connections or junctions—the term “junction” is used in the expert system. Junctions are defined in a similar way to busbars as shown:

```
junction(  
    ss5,  
    j2a,  
    [t1,t2,t3]  
).
```

5.3.3.8 Line Details

A line simply has its impedance detailed. The two halves of a line are its subunits, and are labeled 'e1' and 'e2'. The following fact defines a line, and illustrates the use of 'to' in relating interconnected substations:

```

line(
    ss1 to ss2,
    l1,
    imped(z(0.004, 0.016))
).

```

This describes a line between substations 'ss1' and 'ss2', labeled (locally) 'l1'. For reasons of consistency, a convention has been adopted for ordering interconnected substations: higher voltage substations come first, while the ordering is alphabetical between substations with the same voltage level.

5.3.3.9 3 Terminal Line Details

The subunits of a 3 terminal line are called 'e1', 'e2' and 'e3'. These subunits represent each of the three limbs of the teed line. Impedances are detailed for each limb. This is shown in the following fact which describes a 3 terminal line interconnecting substations 'ss1', 'ss4' and 'ss5':

```

line3t(
    ss1 to ss4 to ss5,
    l1,
    imped(z(0.003, 0.015), z(0.006, 0.03), z(0.006,0.03))
).

```

5.3.3.10 Cable Details

The definition of a cable is directly analogous with that of a line. For example:

```

cable(
    ss2 to ss3,
    c1,
    imped(z(0.001,0.01))
).

```

5.3.3.11 Reactor Details

The subunits of a series reactor are its two terminals, 't1' and 't2'. A reactor has its impedance defined, as shown in the following fact:

```
reactor(  
    ss6,  
    r1,  
    imped(z(0.001,0.2))  
).
```

5.3.3.12 Machine Details

A machine has its effective/synchronous impedance detailed. The machine terminal ('terminal') is the only subunit used in the description. The following fact defines a machine:

```
machine(  
    ss7,  
    gen2,  
    imped(z(0.0,0.2))  
).
```

In addition, some more machine details are required solely in order to perform a loadflow. The details are expressed in a separate fact and describe the type of loadflow node, and define the “normal” constant parameters appropriate to the node type ie. per unit real power and voltage are defined for a pv node, whereas per unit real and reactive power are defined for a pq node. Thus far, reactive power limits have not been included. These details are expressed in the form shown by the following example:

```

machine_lf(
    ss1,
    gen2,
    pv,
    0.7,
    1.01
).

```

5.3.3.13 Supply Point Details

Transmission systems are seldom “island systems”. They generally form part of large interconnected systems spanning vast areas. There are, therefore, busbars in a network which are supply points from undefined neighbouring networks (refer to Figure 5.1). Likewise, a complete transmission system, especially during feasibility testing of the expert system, will often not be defined; supply points from undefined parts of the same network will often exist. These supply points could feed a fault on the system and therefore influence the diagnosis. Such infeeds are therefore defined as shown for section 's1' of busbar 'b1' in substation 'ss1':

```
infeed(ss1::b1\\s1).
```

Infeeds of this sort clearly influence the load flow in the system. Therefore, a fact with the same arguments as the corresponding machine fact is supplied for each infeed. An example is shown:

```

infeed_lf(
    ss1,
    b1,
    pv,
    0.4,
    1.01
).

```

5.3.3.14 Loaded Busbar Details

Busbars may be loaded. While knowledge of this is not directly needed for diagnosis, it is necessary for the loadflow analysis which is performed off-line. Each loaded busbar has its “normal” real and reactive load declared in a fact. Such a fact follows:

```
load(  
    ss2,  
    b1,  
    0.02,  
    0.005  
).
```

5.3.3.15 Slack Node Definition

The loadflow also requires the identification of an infeeding busbar or a generator as the slack node. The following fact identifies 'gen4' at substation 'ss1' as such:

```
slack(  
    ss1,  
    gen4  
).
```

5.3.3.16 Single Phase Plant Identification

All plant is assumed to be three phase unless it is specified otherwise. Single phase plant is identified by 'single_phase/1' facts. Using Prolog's “don't care” underscore variable, general facts can be specified which cover many parts of the network. For example, the following five facts define all transformer neutral subunits to be single phase:

```

single_phase(_::_\hv_ntrl).
single_phase(_::_\lv_ntrl).
single_phase(_::_\lv1_ntrl).
single_phase(_::_\lv2_ntrl).
single_phase(_::_\ntrl).

```

5.3.3.17 Plant Relational Data

The aforementioned plant details define the existence of busbars, transformers, etc. To be fully defined, a network must also have the plant's interconnection (ie. relationship between plant, or network topology) detailed. This topological data is basically static since it always exists. However, the status often changes on-line because many of the connections between network equipment are made through switchgear.

Topological details describe the connections between items of plant. The subunits of the respective plant items are related, with any switchgear between them also identified. Such a connection is defined by a fact like the following:

```

network(
  x420 at ss1,
  ss1::b4\s1,
  ss1::b5\s2,
  on
).

```

This example describes a network connection formed by closed (ie. 'on') switchgear 'x420 at ss1', between plant 'b4\s1' and 'b5\s2' in substation 'ss1'. If the connection is not formed by switchgear, then a 'network/4' fact is used with the switchgear argument filled with a dummy name and the status argument permanently containing 'on'.

In addition, the topology of some individual items of plant needs to be similarly defined. The internal connectivity of busbars and junctions, unlike other plant types, changes on-line. Busbar and junction subunits may be connected

through manually operated isolators. These connections are defined in a similar way to network topology. This is shown by the following fact which relates to the connection between subunits 's1' and 's3' of 'b1' in substation 'ss4':

```
sectn_net(  
    x528 at ss4,  
    ss4::b1,  
    s1,  
    s3,  
    on  
).
```

Again, the isolator name can be a dummy name if there is no genuine switchgear. Genuine switchgear is detailed in separate facts which are described in section 5.3.4.1.

5.3.4 Protection and Switchgear Data

Switchgear, relay and current transformer (CT) details, and the relationship between these devices are also needed for a network to be diagnosed.

5.3.4.1 Switchgear Details

The existence and type of switchgear equipment is defined by a fact, an example of which follows:

```
sw(  
    ss1,  
    x420,  
    breaker  
).
```

The above defines a circuit breaker 'x420' at substation 'ss1'.

A fact like the latter relates to three single phase switches. The permissible switchgear types are: circuit breaker ('breaker'), remote controlled isolator ('rc_isolator'), remote open isolator ('ro_isolator'), switching isolator ('sw_isolator') and manual isolator ('isolator').

Each circuit breaker which has an autoreclosing capability has an additional fact defined. The fact also gives the associated protection relay and the maximum number of trip/reclose cycles the breaker can go through before being locked out. The following fact illustrates this:

```
auto_reclosing(  
    ssl,  
    rmain1-2,  
    x420,  
    3  
).
```

A similar fact is used to define delayed autoreclosing circuit breakers, namely `del_auto_reclosing/3`.

5.3.4.2 Relay Details

The details required for a relay are dependent on the relay type. The type is not the actual manufacturer's type name, but is the functional type. The top-most classification is as protection or trip relay. Protection relays are further classified as differential (ie. circulating current, 'diff_cc'), directional overcurrent ('dir_oc'), overcurrent ('nondir_oc'), high set overcurrent ('hs_oc'), distance ('dist'), instantaneous earth fault ('zps'), restricted earth fault ('ref') or earth current overcurrent ('ground_oc').

The protection modelling is predominantly qualitative, though non-unit protections are modelled using cumulative impedances ie. effective impedance from a relaying point to faulted plant (see section 5.6.3.2). Consequently, settings are

not required since characteristics are not numerically modelled. The relay details which can be included are type, impedance reach, operational priority (indicated by an integer between 1 and 5), primary/secondary class, and preconditions for operation. Relational details may be included which give the controlled switchgear, the burdened CTs, plant on the reverse side of the relaying point and the controlling relays (ie. relays which control the operation of the described relay). Arguments are available for these in the facts for relays, and are used in different combinations for the various relay types. Irrelevant arguments are set to 'n'. As an example, the following fact describes an instantaneous earth fault relay 'rief2' at substation 'ss2', which does not directly control any switchgear, is the burden for local current transformer 'ct232', has plant item 'ss2::b2' on the reverse side of its relaying point, has the highest operational priority of 1, and is primary so has no controlling relays:

```

relay(
  ss2,
  rief2,
  zps,
  [],
  [ct232],
  n,
  n,
  n,
  ss2::b2,
  1,
  no,
  []
).

```

Moreover, if a relay has an autoresetting capability, it has an additional fact defined like the following:

```
auto_resetting(  
    ss2,  
    rief2  
).
```

5.3.4.3 Current Transformer Details

Current transformers (CTs) and their position are described. The location of CTs is identified by the plant on which the CT is positioned, and the next nearest plant. This is required because otherwise the exact location of a CT may not be discernible when a hypothesised fault and a CT are on the same plant. The next nearest plant is required to be included in order that, in the case of a hypothesised fault being on the same plant as the CT, it may be established which side of the fault the CT is on—this problem is in fact a result of regarding plant subunits as nodes. A CT, as defined in the knowledge base, could be one coil on the single phase plant connected to a star point. It could instead represent three coils at the same position but on separate phases (or four coils, with one on the neutral). These appear the same in the knowledge base CT details—the plant on which the CT is located identifies whether or not it is a single coil.

The connections of CTs on separate phases is defined by the type of the relay which burdens them. As an example, consider the CT referred to by the relay example in the previous section, 'ct232'. The current supplied to an instantaneous earth fault ('zps') relay is the sum of the currents from three coils on the three phases. The single fact which identifies the three phase coils (while the relay fact identifies their connection) is:

```
ct(  
    ss1,  
    ct232,  
    ss1::b2\\s10,  
    ss1::b2\\s11  
).
```

5.3.4.4 Relational Data

Relational data relates plant, switchgear, relays and current transformers with each other. Relationships with plant (ie. location data) for switchgear, relays and current transformers have already been described in sections 5.3.3.17, 5.3.4.2 and 5.3.4.3 respectively. The burdened CTs and the controlled switchgear of a relay have already been described as being included in the relay fact. The CTs related to a relay may be in the same substation, or in a remote station. Both can be included in the CT argument in the relay fact.

Relay to relay relationships are detailed in two ways. Firstly, the enabling of relays by local trip and control wiring is detailed in the relay facts of the controlled device. The last argument in the relay fact is a list of controlling relays with associated additional auxiliary switch logic. Logical statements involving 'and' and 'or' express the logic of the control. For example, consider the trip

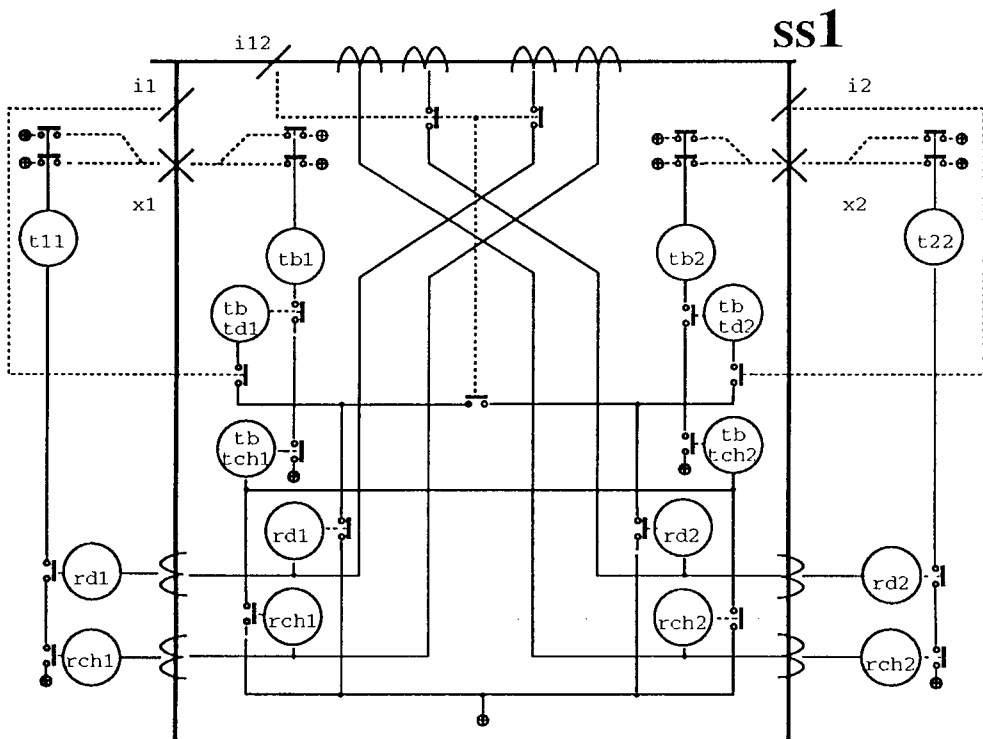


Figure 5.3: An example of some trip and control wiring

and control wiring shown in Figure 5.3 (from reference [78]). This would be

represented in the knowledge base as:

```
relay( ss1, t11, trip, [x1],...,
      [e(rd1,operated(rch1 at ss1)), e(rch1,operated(rd1 at ss1))]).
relay( ss1, tb1, trip, [x1],...,
      [e(tbtd1,operated(tbtch1 at ss1)),
       e(tbtch1,operated(tbtd1 at ss1))]).
relay( ss1, tbtd1, trip, [],...,
      [e(rd1,closed(i1 at ss1)),
       e(rd2,(closed(i1 at ss1) and closed(i12 at ss1)))]).
relay( ss1, tbtch1, trip, [],...,
      [e(rch1,true), e(rch2,true)]).
relay( ss1, rd1, diff_cc, [],...).
relay( ss1, rch1, diff_cc, [],...).
relay( ss1, t22, trip, [x2],...,
      [e(rd2,operated(rch2 at ss1)), e(rch2,operated(rd2 at ss1))]).
relay( ss1, tb2, trip, [x2],...,
      [e(tbtd2,operated(tbtch2 at ss1)),
       e(tbtch2,operated(tbtd2 at ss1))]).
relay( ss1, tbtd2, trip, [],...,
      [e(rd2,closed(i2 at ss1)),
       e(rd1,(closed(i2 at ss1) and closed(i12 at ss1)))]).
relay( ss1, tbtch2, trip, [],...,
      [e(rch2,true), e(rch1,true)]).
relay( ss1, rd2, diff_cc, [],...).
relay( ss1, rch2, diff_cc, [],...).
```

Secondly, relationships between relays may be expressed in terms of protection schemes. Relays belonging to a particular scheme and a description of the scheme are detailed by facts in the knowledge base. Intertrip schemes are represented in a simplified form by this structure. An example of a simple one-way intertrip scheme identified by the name 'it-1-2-12' is shown:

```
relay_scheme(rief2 at ss1,'it-1-2-12',1).
relay_scheme(titr2 at ss2,'it-1-2-12',2).

scheme('it-1-2-12',_,_
      [defn(_,control(rief2 at ss1,titr2 at ss2,true))]).
```

The scheme described by these facts involves a primary protection relay at

'ss1' and a secondary relay at 'ss2' and is such that the primary relay unconditionally controls the secondary relay.

5.3.5 Relative Likelihood Data

The relative likelihood data provides numbers by which a comparable measure can be generated for a scenario. Based on its relative likelihood, a scenario is given a rank in a table of competing scenarios. The scenario with the highest relative likelihood occupies the top ranking position, and is the scenario thought most likely to describe the genuine system state. In order to derive these comparable measures, values must be provided for each possible aspect of a scenario. Typical data is required, and may be supplemented with particular data.

5.3.5.1 Typical Data

Relative likelihood values must be included in facts for each type of aspect which a scenario could contain. The aspects covered include:

- a fault type on a plant type
- fault type detected by protection type
- correct secondary activity
- predicted correct activity
- failure of device type
- false activity
- alarm due to false activity is predicted

5.3.5.2 Particular Data

The ability to add particular data to the knowledge base allows the inclusion of knowledge of deviations from typical performance. The particular relative likelihood offsets the relative likelihood of a scenario aspect: instead of the typical value, the product of the applicable typical and particular values is used. Particular relative likelihood values are identified in `certnty/3` facts containing the appropriate alarm, device or plant. Suppose, for example, that one out of two duplicate main protection relays operates and that, on investigation, it was found to be unjustifiable. The reliability of the relay is now in question so that, if the same alarms were then observed later, a false alarm, rather than fault plus failure, would be the most likely explanation. This knowledge may be incorporated into the knowledge base by the addition of particular data like the following:

```
certnty(  
    operated(rmain1-2 at ss1),  
    correct_opn,  
    0.5  
).
```

Moreover, such details can be related to groups of devices by making use of the flexibility in Prolog's matching technique and using Prolog's rule structure. The chance of failing to receive alarms from a particular substation, for example, can be easily changed.

5.3.6 Telemetry Data

Sufficient details must be provided such that SCADA messages can be read in and interpreted. A protocol fact is required in order to derive the field contents from the message string. The following example is that used to interpret alarms of the type shown previously in section 5.2:

```
telemetry(
  protocol,
  [2,tofday,1,5,1,7,1,6,1,11,1,25,1,6,1,11,1,25,1,6,1,3,1,5,rtn],
  1
).
```

The above fact indicates that the SCADA messages have two unused characters, then the time tag, one character, then a five character field, one character, then a seven character field, and so on, and that they are each contained on one line. Such a fact is used to read in the characters which make up a message and to identify the field contents from the string of characters. Rules which act as templates for operate, reset, trip, reclose, open and close alarms then relate the field contents to specific activity types. Switchgear is usually identified uniquely by its substation and local name. Relays are not referenced in this manner. Instead, a reference table relating the appropriate record contents to specific devices is required. The expert system therefore needs a list of facts relating field contents to particular devices. The following example fact for relay 'rdcoat1' at substation 'dumb' uses the substation, description/function, circuit and legend fields of the message shown earlier (section 5.2):

```
spec_relay_id(
  dumb2,
  '',
  coat2,
  'first main protn optd',
  rdcoat1
).
```

5.4 Qualitative Simulation

5.4.1 Function

The qualitative simulation knowledge is used by the central diagnosis of the hypothesis strategy. Its function is to simulate the effect which a fault and failures hypothesis would have on the system in order to determine the activity which would be caused. Consequently, the SCADA alarms which would be expected to follow the genuine occurrence of the hypothesis are obtained. These results are then the basis of the immediate tests performed to check the validity of the hypothesis.

5.4.2 Reason for Being Qualitative

The diagnosis could have been developed with a numerical, or quantitative, simulation at its core. The simulation of protection activity would require the fault current at various points in the network. Utilities have many sampling points which are continuously polled giving a system-wide pattern of current, voltage and power levels. However, the transducers are typically limited to 125% of full load and therefore could not give any indication of fault levels. Incidentally, their sampling rate is such that it would be unlikely that the sampling instant would correspond with the short period in which fault conditions exist: while a fault could be cleared within 100 milliseconds, analogues are sampled every 5 seconds on the ScottishPower transmission system. Although algorithms for calculating fault levels are well established, the task is not trivial. Significant processing time would be required. One possibility which was considered was to perform a complete network fault level study off-line, and use the results to fill a look-up table. During on-line diagnosis time, the appropriate fault level would be obtained from the look-up table. However, in order that the currents “seen” by remote

backup relays might be known, a power flow analysis would also be necessary. A look-up table of results from this analysis would be unmanageably large for off-line analyses to suffice. Since the purpose of the simulation is to test the validity of hypotheses, several simulations may be required following an alarm. Computation time would be too excessive for an on-line numerical solution using the computing power presently available. (This limitation could be eradicated with further developments in digital signal processing chips. Even then though, it would be inefficient to use only intensive numerical processing if the same point can be reached by reasoning followed by a directed numerical analysis.)

An expert operations or protection engineer would not perform a numerical fault flow in order to determine possible explanations for the received activity. And yet he can successfully perform the diagnosis of system faults by reasoning. This analysis may precede a detailed numerical analysis. The knowledge incorporated in the expert system can likewise perform a qualitative simulation. Having located the problem by this means, the deeper diagnosis could be made to initiate a standard fault flow algorithm for the analysis of the suspected system fault. Finer maloperations, like timing problems, could be recognised by this numerical simulation.

5.4.3 Representation of Activity

Simulated activity can be recorded by the assertion of facts in the knowledge base. One set of facts represents the simulated status of the active devices, while another records the simulated changes of status. By gathering together the simulated change of status facts, the expert system identifies the SCADA alarms expected to follow the genuine occurrence of the simulated hypothesis. Details will be given in the section describing the modelling of activity (section 5.6.4).

Simulated activity does not influence the genuine representation of the network. For example, the facts indicating network connections are not changed by the simulated tripping of a circuit breaker. Rather, the reasoning uses the simulated status facts as an overlay. This overlay knowledge is always used first when checking a device's status. If the check is proved by simulated details, then the tested condition is satisfied. The actual status facts are only checked if the latter test fails and no simulated status facts for the device exist. The use of this overlay prevents any inadvertent changes to the internal representation of the genuine system status.

5.4.4 The Qualitative Simulation Knowledge

The simulation can be viewed as consisting of two major aspects, namely the system simulation and the device simulation. The system simulation resolves the direct effect of a system fault. It therefore determines the primary activity (operation of protection relays) which may be expected. The device simulation then simulates the primary activity, and determines and simulates the consequential activity caused by the protection relay's operation. The device simulation knowledge is also used to determine the activity effected by a false operation. Section 5.6.4 describes the device simulation knowledge.

The protection detection logic models the detection behaviour of devices but not the time characteristics. As the dynamic behaviour of protection is dependent on system currents and voltages, and these are not known, the use of actual timing is not reasonable. Protection schemes are designed with great care given to coordination such that protection nearest the fault location always operates first (given correct operation). This fact has been used to develop the qualitative simulation. Protection nearest the hypothesised fault location is always considered before more remote devices. In addition, a priority classification on each

relay is used by the expert system to discriminate between devices in the same location. The knowledge in the qualitative simulation of fault isolation can best be described by a series of statements. These are listed below:

- a protection operation is required if a fault path exists ie. there exists an electrical connection between the hypothesised fault position and a machine or power infeed through which fault current can flow
- protection is in a substation within the disturbance area
- from the fault location, the plant along the fault path is searched to find CTs
- CTs can be on the fault current carrying plant
- CTs can be on the HV star point to ground connections of a transformer whose HV winding is connected to a junction which is on the fault path (ie. an earthing transformer)
- the highest priority protection relays burdening these CTs have their operation simulated if they can “see” the fault.
- secondary relay and controlled breaker activity is simulated with primary operations
- if the same fault path is not isolated then another protection relay is sought—it is a backup operation
- a backup operation is identified by looking first at the next highest priority relays and then at the CTs on the plant next along the fault path
- when a fault path is isolated, another one is sought

- the simulation is complete when no fault path exists if the alarmed device activity has been included in the simulated activity
- if the device's activity is not included then an additional failure is required

The rule which contains this knowledge is shown:

```
isolation_simulation(DistID,SSsInvolved,LOpdRelay at SS,
  FType fault_on FPos,ExtraDetails):-
  ( route_to_power_infeed(FPos,CBsOfRoute,ZonesOfRoute)->
    ( (isolating_protn(FType,ZonesOfRoute,FoundRys,SSsInvolved),
      not(previously_found_backedup(FoundRys,CBsOfRoute)),
      consequential_opns_to_expect(operated,FoundRys,
        FType fault_on FPos,CBsOfRoute,FoundRys),
      member(CB,CBsOfRoute),
      simul(cb_status,CB,off),
      identify_potblocking_devices(LOpdRelay,SS,FoundRys,
        ZonesOfRoute,CB),
      gather_backup_details(NewSofarBU),
      !,
      isolation_simulation(DistID,SSsInvolved,LOpdRelay at SS,
        FType fault_on FPos,ExistingExtraDetails),
      add_new_bu_details(NewSofarBU,FoundRys,ExtraDetails,
        ExistingExtraDetails))
    ; fault_not_isolated(FType fault_on FPos),
      !,fail
    )
  ; ( simul(telemetered,operated(LOpdRelay at SS),_) ->
    ( reason_output(
      ( writeout('<<<< A '),
        writeout(FType fault_on FPos),
        writeout(' is cleared.'),
        writeout(code(nl))))),
      ExtraDetails=[])
    ; fault_doesnt_explain_operation(FType fault_on FPos),
      !,fail
    )
  ).
```

A description of the workings of the above rule follows.

A. 1A. `isolation_simulation(...)`:-

The `isolation_simulation/5` rule is called to simulate a fault belonging to a disturbance involving certain substations, attempted because of an operated relay, the fault being a particular type and location, and for which extra details are to be gathered.

2A. `route_to_power_infeed(...)`

A route from the fault position to a power infeed along which fault current could flow is identified by the `route_to_power_infeed/3` rule. The power infeed may be a supply point (infeed from undefined part of the network) or a machine. Lists of the fault current carrying plant and circuit breakers along the route define the found fault route.

3A. `isolating_protn(...)`

The list of fault current carrying plant is used to identify the CTs nearest the fault location and hence the potentially isolating protection.

4A. `not(previously_found_backedup(FoundRys,CBsOfRoute))`

Backtracking is forced in order to find another relay if the found protection has previously been found.

5A. `consequential_opns_to_expect(...)`

Consequential activity is simulated together with the operation of the protection relay.

6A. `member(CB,CBsOfRoute), simul(cb_status,CB,off)`

The reasoning then checks if a breaker along the fault current path is now open. If this fails then backtracking returns the reasoning to the `isolating_protn/4` rule to find protection which would operate due to the persistence of the fault current.

7A. `identify_potblocking_devices(...)`

When the fault path is opened by simulated breaker activity, a check is made of the relationship of the isolated fault path to the actually operated relay whose operation the simulation is based on. Suitable facts are stored if it appears the simulated device(s) operated instead of the actually operated relay.

8A. `gather_backup_details(NewSofarBU)`

Details of failed protection or switchgear which have been created in the knowledge base are collected and removed.

9A. `!, isolation_simulation(...)`

The satisfaction of the rule is then attempted again based on the new simulated configuration held in the knowledge base. The cut (`'!'`) prevents the reasoning from backtracking past this call into the previous rules (**1A** ... **8A**).

B. However,

3B. `isolating_protn(...)`

The `isolating_protn/4` rule fails if no protection is found (on the first attempt or after backtracking). In such a circumstance the reasoning progresses to the alternative logic (indicated by `' ; '`).

4B. `fault_not_isolated(FType fault_on FPos)`

The `fault_not_isolated/1` rule asserts the fact of the failed isolation.

5B. `!,fail`

Retrying on backtracking is prevented by the cuts (the `'!'` here and before the recursive call to `isolation_simulation/5`).

The `isolation_simulation/5` rule then fails.

C. On the next attempt to satisfy the `isolation_simulation/5` call after a hypothesised fault has been completely isolated by the simulated tripping of a breaker,...

2C. `route_to_power_infeed(...)`

The `route_to_power_infeed/3` rule will fail. Reasoning then progresses to the appropriate `' ; '` ie. the second one due to the nesting of the logic.

3C. `simul(telemetered,operated(L0pdRelay at SS),_)`

The actual genuine operation on which the simulation is based is then sought among the simulated activity facts.

4C. `reason_output(...)`

Since the genuine operation is included, the fault isolation simulation has been successfully completed. An appropriate comment is displayed to the user if the running commentary has been selected.

5C. `ExtraDetails=[]`

All that remains is to gather the backup details. The extra details list is set initially to be empty. This completes the satisfaction of the immediate `isolation_simulation/5` rule. The reasoning then returns to where the rule was called.

6C. `add_new_bu_details(...)`

The outstanding `add_new_bu_details/4` rules—those after each recursive `isolation_simulation/5` call—gather together the collected information from each stage of the isolation simulation.

The `isolation_simulation/5` rule is then satisfied.

D. If, however, the genuine operation is not included...

3D. `simul(telemetered,operated(L0pdRelay at SS),_)`

The `simul(...)` call will fail; the hypothesis does not account for the actual activity. The reasoning steps to the alternative logic (ie. after the third `' ; '`).

4D. `fault_doesnt_explain_operation(FType fault_on FPos)`

The potential blocking devices (contained in facts asserted during the qualitative simulation) are recorded as failure candidates.

5D. `!,fail`

Retrying on backtracking is prevented by the cuts (the `' ! '` here and before the recursive call to `isolation_simulation/5`).

The `isolation_simulation/5` rule then fails.

Initially the `route_to_power_infeed/3` rule involved a relatively straightforward depth-first search through the system topology. However it was recognised that, having entered a substation and reached a busbar, the search would often progress out of the substation on another circuit without finding the cable to which a generator was connected. Therefore, knowledge was required to bring knowledge of the generator location to the busbar level of the substation topology for early recognition. A rule was developed which traced from each generator to all directly connected busbars, and asserted the routes found. This rule is invoked when the search for a power infeed enters a substation. Then, instead of checking the traced plant for generators, checks are made for the busbars which are known to be connected to generators. In this way the route to power infeed search is much more efficient.

After the fault isolation has been simulated, production rules are fired to simulate the post isolation simulation. This simulated activity may range from simply relay resetting to automatic reclosing activity.

5.4.5 Explanation/Understanding Facilities

The description of the qualitative simulation illustrates the fundamental nature of the knowledge used in this expert system. The qualitative fundamental knowledge based approach means that the expert system can be easily made to include explanatory details which aid understanding of the situation. A separate specific knowledge base would be required for explanation facilities if a shallower level of diagnosis had been used. The approach adopted should give a more flexible and hence more portable solution. The necessary additions have been made to the knowledge in the qualitative simulation.

As well as determining the activity which would follow a system fault, the qualitative simulation also generates explanatory details. The details obtained include protection which should have cleared the fault, reasons why protection does not isolate the fault (eg. relay or breaker failure) and protection which backed up particular failures. The description of a diagnosis conclusion is supplemented by this information. Its output is suppressed, however, unless the expert system has been told to give a detailed description of scenarios. An explanation automatically generated by the expert system for a particular diagnosis is shown:

...In fact because rief2 at dumb will not isolate the fault due to the failure of cb5 at dumb to trip, rdef1 at drum operated, rsbef2 at dumb operated, & rsbef1 at dumb operated to isolate the fault.

This example describes how the failure of its controlled switchgear caused a protection relay ('rief2' at substation 'dumb') to fail to isolate the fault, and how the relays 'rdef1 at drum', 'rsbef2 at dumb' and 'rsbef1 at dumb' operated in a backup capacity, resulting in the isolation of the fault.

This explanation knowledge is derived by reasoning; it is not a static body of text appended to the diagnosis conclusion when appropriate. Its derivation is by the means of the knowledge used in the qualitative simulation, which itself represents the true behaviour of the system. Failure and backup operation relationships are obtained naturally from the rules which identify the nearest protection relays with the highest priority first, and the next highest priority or the next nearest relays on backtracking.

Although the creation of explanatory details is natural to the fundamental knowledge based qualitative simulation, it does entail some extra reasoning work. Hence, additional processing time is used. For example, if during the device simulation a failed relay is encountered, the reasoning will not simply halt. Recognising the failure, the expert system will instead trace through the trip and control wiring to check that the fault path would have been isolated had the relay not failed. The extra rules are minimal since much of the reasoning is included in the rules to determine activity following a fault. The expert system's "understanding" appears deeper than it might otherwise have been.

5.5 Loadflow

5.5.1 Power Infeed Search

As part of the hypothesiser strategy, it is necessary to trace through the network. The reasoning attempts to find a path for fault current by starting at a hypothesised fault location and tracing back to a machine or supply point. Having found such a route, protection along it is identified and simulated, together with controlled relays and switchgear. With the simulated tripping of switchgear along the path, other routes for fault current are likewise traced and dealt with until complete isolation is accomplished. The forming of a fault current path was achieved by successively searching for plant connected to the present plant, and then looking for plant connected to the new plant. A change was made to this basic inference by identifying and looking for a connection to generator feeders instead of machines themselves (see section 5.4.4). This depth-first search worked successfully in most cases.

5.5.2 The Original Problem

In some feeders of some networks, the power flow is only ever in one direction. For such plant, the protection schemes used are effectively open-ended designs since protection relays are only installed at the sending end. According to the above means of identifying a fault current path, the direction in which the search first progressed simply depended on the arrangement of the topology facts in the knowledge base. Therefore, for hypothesised faults on the described feeders, the search sometimes went in the direction of power flow. The result was that the search failed to identify any protection, or wrongly detected backup protection first.

This can be illustrated by considering the qualitative simulation of a ground

fault on the second transformer feeder from 'drum' in the example network in Figure 5.4. Suppose the search for a power infeed commences in the direction




Figure 5.4: Illustration of the failure of the depth-first search

shown by the broken line. The first protection identified is the backup protection (standby earth fault relay) on infeed circuit 1 (to which transformer 't1' is connected). The operation of this protection and its trip relay and the tripping of the substation's infeed circuit breaker are simulated. The second search for a fault current path results in the simulation of the standby earth fault relay on infeed circuit 2 (transformer 't2' is in this circuit) and the consequential isolation of substation 'drum'. The simulation completely isolates the hypothesised feeder fault, but erroneously includes activity of the standby earth fault relays, their trip relays and breakers.

Non-unit protection relays "look" out into the network from a relaying point.

A radial feeder may have such protection at one end since any fault can only be fed from that one end. In an interconnected network, non-unit line protection will often be arranged symmetrically on both ends of the feeder since a fault could be fed from either end: the protection scheme can therefore be described as closed. The depth-first method of searching for a fault current path, as described above, would work correctly in both these cases. The aforementioned problem arises in an interconnected network where, because of the normal use of the system, non-unit protection is not placed symmetrically on both ends of the feeder. Such a scheme is basically open-ended due to the absence of corresponding protection at one end of the feeder.

5.5.3 Possible Solutions

There are considered to be three possible solutions to the problem: a width-first search, the addition of extra system knowledge and the derivation of knowledge from a loadflow.

5.5.3.1 Width-First Search

The aforementioned problem could occur when a longer path to a power source is found before a path which is shorter. A possible solution is therefore to use a width-first search instead of the depth-first search. The search would be impedance based: alternative paths would be traced, with progress in each being maintained at approximately the same cumulative impedance.

Applying this technique to an interconnected network would result in numerous alternative paths being created. Many of these alternatives would appear once the search has moved well away from the locality of the hypothesised fault. As a result, several of these routes would be opened near the fault position when a breaker has its tripping simulated. In contrast, the depth-first approach would

find one route, open that route, and in so doing isolate several alternative paths which shared the tripped breaker but which had not yet been identified. Therefore, much more reasoning would be required for the width-first approach.

5.5.3.2 Addition of Knowledge of System Operating Conditions

The problem arises with plant which is only ever used to supply power in one direction. Plant of this type could be identified in the knowledge base by a fact defining the direction of power flow. For example, the following `one_way_power_flow/2` fact would identify that line 'drum to dumb::fd2' is used for one way power flow, from end 'e1':

```
one_way_power_flow(drum to dumb::fd2,e1).
```

The qualitative simulation would check for these facts. When such a fact exists for the hypothesised fault location, the first fault path search would always be made in the direction of the subunit in the second argument position.

Alternatively, it would be possible to develop rules to recognise these conditions. The arrangement is mainly used on transformer feeders which provide a bulk supply to a lower voltage system. Such cases are identified by the following rule.

```
bulk_supply_transformer_feeder(SS::LTrans,SSs::LFeeder):-
    transformer(SS,LTrans,3,hv(_,,HVoltage),lv(_,,LVoltage),_,_),
    voltage_level(HVoltage,transmission),
    voltage_level(LVoltage,distribution),
    link(LDummy at SS, SS::LTrans\\hv_wdg,SSs::LFeeder\\_,on),
    line(SSs,LFeeder,_),
    not(sw(SS,LDummy,breaker)).
```

This rule may be briefly described as follows. The arguments in the rule head contain two items of plant. The rule body requires the first of these

to be a three phase transformer, whose HV and LV voltage levels can be described as 'transmission' and 'distribution' respectively (according to the `voltage_level/2` rule). The HV winding must be connected to the other item of plant (according to the `link/4` rule), and that other item of plant is a line, while the connection must not be through a circuit breaker. Another rule could assert suitable `one_way_power_flow/2` facts on the basis of satisfaction of the above rule.

5.5.3.3 Derivation of Knowledge from Loadflow

This technique is very similar to the latter solution which derives system-use knowledge. This solution does not exclusively consider the particular case in question; more general knowledge is derived, of which the particular case is a subset. The derived knowledge is the direction of real current flow with normal loading conditions and the network configured with all switchgear closed. A loadflow is performed once with the said network conditions, and the results used to identify the subunits from which real current flows into plant items. Facts like the following are stored:

```
nor_flo_from(ss1::b1,s1).
```

The knowledge provided to solve the original problem (ie. in which direction to search first from the hypothesised fault) is fundamentally the same as that provided by the solution described in section 5.5.3.2. However, the loadflow solution provides added knowledge which guides the search towards a power infeed at every stage of the search, and not just at its start. Having reached a plant item, the search for a power infeed will first progress in the direction of the subunits identified by `nor_flo_from/2` facts.

5.5.4 Solution Adopted

Due to the identified processing inefficiency of the width-first search, the other two solutions were considered preferable. Both of these (system use and load flow knowledge based solutions) would satisfactorily overcome the problems encountered originally. However, the generality of the knowledge derived from the loadflow, provides a means of improving the efficiency of the search for a power infed in general: by following the reverse direction of normal power flow, a power infed will generally be found more quickly than by a “blind” search. Therefore the loadflow knowledge based approach was adopted.

Loadflow analysis is computationally intensive. Therefore, the algorithm was coded in FORTRAN, and later in C. The loadflow analysis is performed off-line; only the knowledge derived from the loadflow is used on-line. Loadflow analysis uses network data in a different form to that held for fault diagnosis. To avoid the user having to enter the configuration of the network twice and in different formats, Prolog rules were developed to translate knowledge base network details into the required node and branch data. Busbar and junction sections, three terminal lines, three power winding transformers and machines are identified as nodes, and appropriate facts asserted. Corresponding branch end facts are asserted for each external subunit. The representation as nodes and branch ends is illustrated by the schematics in Figure 5.5. Branch end facts are checked before a corresponding fact is created. If a new branch end is directly connected to a branch end which is already in the knowledge base (eg. in the case of two connected busbars) then: the adjoining branch ends are removed from the knowledge base; the new node fact is removed with any loads being reassociated with the connected old node; the remaining branch ends are associated with the old node. The number of nodes is significantly reduced by this merging of adjacent nodes. Prolog demonstrated its advantage as a general purpose language in that

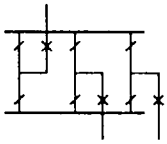

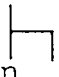
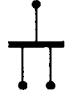

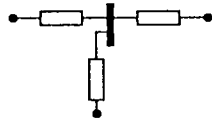
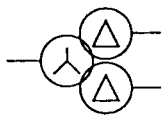
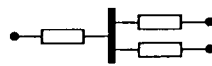
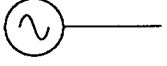
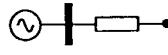


Network Details	Nodes & Branch Ends
 <p>busbars</p>	
 <p>junction</p>	
 <p>3 terminal line</p>	
 <p>3-power-winding transformer</p>	
 <p>generator</p>	
	<p>  node  branch end </p>

Figure 5.5: Schematics of loadflow nodes

the creation of rules to translate the data format was straightforward.

As well as connectivity details, a loadflow requires the parameters of a network. Normal load and generating power levels and generating voltage are used to form node data. Line, 3 terminal line, transformer, autotransformer, 3-power-winding transformer and reactor impedances are used to determine branch impedances. Susceptances were not used due to the difficulty of forming a II-arrangement for transformer feeders.

Quintus Prolog [77] provides a foreign language interface by which foreign code can be called from Prolog, and Prolog can be called from foreign code. This facility is used to initiate the loadflow analysis, but not to transfer data. The translated network data is written to a file, which is in turn read by the loadflow

program. The volume of data was thought to preclude transfer by the direct communication facilities provided.

The loadflow program was developed according to the fast decouple algorithm [79]. The loadflow is commenced once the generated data file has been read in. If a matrix inversion fails, then the solution is stopped and a code is returned to the Prolog rule to indicate such. Likewise, if the solution fails to converge, a different code is returned to the expert system. If the solution converges, then, for each branch, the node at the end from which real current flows is identified. A data file is created which includes, for each branch, the branch number, two end node numbers and the node from which real current flows. This direction of flow data together with a code to indicate successful completion is the only information which is passed to the expert system.

If the success code is received at the Prolog side of the interface, then the file generated by the loadflow is read in. The plant items along a branch have the subunits on the side of the infeeding node stored in `nor_flo_from/2` facts. These are the facts which are used on-line.

5.5.5 Comments

Regard has only been given for real current flow. Sometimes on interconnectors the direction of real and reactive power flow can be different. Such conditions would only exist on interconnectors between systems, on which the problematic open-ended arrangement would not be implemented. Therefore, the knowledge of normal flow direction would not be required to prevent erroneous results, but rather only to improve the efficiency of the search for a power infeed. The difference in the direction of real and reactive power flow is not of significance.

The power flow is a load flow and not a fault flow. The stored facts are no guarantee of the direction in which the largest contribution to fault current will

flow. Switchgear could change status such that plant is fed from the opposite direction. The derived loadflow knowledge is not, however, used to confine the search for a power infeed: when open switchgear is found, backtracking commences in order to find an alternative. In most cases, the derived knowledge is simply used to help the search.

Data from a fault study could be used. Due to the processing time required, fault studies would need to be performed off-line, and the results stored in a look-up table. This would contain fault levels at all points on the network. It would not, however, be sufficient for a numerical-based simulation. The current seen by the main protection may be available, but not that seen by any remote (backup) protection. For this, a fault flow would need to be performed. Fault current levels during a fault are dependent on the type of fault; as a result, a fault flow for each type of fault on each item of plant would need to be performed. This data would form a massive look-up table. Moreover, the data provided would only be for a specific topological configuration. The data could not, therefore, be used alone; the data would be required to be used in conjunction with the topological arrangement. Being dependent on the network configuration, the actual current levels could not be depended upon. Therefore, although containing more information, the fault flow data could not be used to any greater degree than could the loadflow data. The extra processing time, and look-up time would give no extra benefit.

The developments have proved a valid solution to the original problem. And more than this, the use of the knowledge to realise a heuristic-like search has also proved beneficial. Despite this, it may be viewed that the AC fast decouple algorithm is excessive for the limited results required. A DC loadflow may be considered sufficient. If an on-line analysis was needed instead of a one-off analysis, then a DC loadflow would be used.

5.6 Modelling

Some modelling knowledge has already been described. The qualitative simulation determines the activity that would be expected following a system fault given certain failures. It does so in a way which embodies the behaviour of power networks and protection systems. The qualitative simulation makes use of the modelling knowledge described in this section, as do other parts of the “hypothesiser” diagnosis and aspects of the “wait-and-see” diagnosis. The modelling knowledge relates to system faults, plant internal links, protection detection logic, and protection and switchgear activity.

5.6.1 Faults on Plant

There are ten types of short circuit faults which can occur on plant: three single phase to ground faults (red, yellow or blue to ground), three phase to phase faults, three phase to phase to ground faults, and a three phase fault. There are more system faults, including various open circuit faults, but these have not yet been included in the expert system. Alarms of protection relay operation seldom identify the faulted phases. Therefore this distinction is not made by the diagnosis. Separate protection relays are often used to protect ground faults and phase faults. This generic distinction can therefore be made from the alarms describing relay operations. The expert system classifies faults as ‘ground’ faults or ‘phase’ faults. (Three phase faults are classed as ‘phase’ faults.)

Ground and phase faults have different effects on system currents. Those of interest are the effects on the current on the supplied sides of faulted plant. A phase fault on an item of plant causes an increase in the magnitude of currents in the plant feeding the fault. To determine the actual current level requires a numerical analysis, but the general effect is for the current magnitude to rise. Ground faults provide a path for current to flow to ground. The imbalance

in the currents feeding the ground fault contains a residual component. The currents feeding a ground fault also have an increased magnitude, though this is generally less than that following the occurrence of a phase fault. Therefore, the facts which define the effects produced by the different faults make a distinction between levels of these effects. The knowledge in the expert system dictates that:

- 'average' 'phase_effects' are produced by 'phase' faults and 'ground' faults
- 'high' 'phase_effects' are produced by 'phase' faults
- 'neutral_effects' are produced by 'ground' faults

Both the qualitative simulation and the protection detection logic use this knowledge.

In addition to location and generic type, faults are further described by their permanence. Faults can be described as 'permanent', 'temporary' or '(permanent or temporary)'. Fault permanence does not influence the fault isolation simulation, but is involved in the post-isolation simulation. If the post isolation is going to involve more than relay resetting, 'permanent' and 'temporary' faults are dealt with—they only differ in post-isolation activity. Otherwise the permanence is set to '(permanent or temporary)'.

5.6.2 Plant Internal Links

The network topology is required for the expert system to be applied to a power network. This data describes the connections between different items of plant. For example, the connection between a transformer HV winding and a busbar section is defined in the topological data. However, the connectivity between subunits of the same item of plant is generally not defined in the network data as it is a characteristic of the plant type. Therefore modelling knowledge is included

to define the characteristic of the plant type. Characteristic internal links between transformer windings, between line ends, etc. are described. (There is no machine internal link knowledge since machines only have one subunit.)

5.6.2.1 Bases of Links

These internal links are not restricted to physical connectivity. Various aspects of the diagnosis require to search through the network on different bases:

- three phase connection:- Links of this sort are made on the basis of three phase connection. The “connection” is by electrical/physical connectivity and magnetic coupling (ie. between windings of the same transformer, 3-power-winding transformer or autotransformer). It could be used to ascertain the existence of a blackout area.
- one phase connection:- This criterion restricts the involved equipment to single phase plant or single phase subunits. As with the three phase connection, the links include both electrical connection and magnetic coupling. It could be used to find all the plant between a star point and ground.
- 'neutral_effects':- On the basis of 'neutral_effects', links are made according to the route taken by a residual component in the system current. A link from one subunit to another exists if a residual component in the current flowing through the first subunit would result in a corresponding effect on the current in the other subunit. It is used to find the path of propagation of a residual component.
- 'phase_effects':- According to the 'phase_effects' criterion, links are made between subunits if an increased system current magnitude in one would result in increased magnitude current flowing through the other subunit. It traces the propagation of an increased magnitude system current.

Whenever the plant internal link knowledge is to be used, as well as specifying a subunit (“find a link”) or two subunits (“check a link”), at least one basis of the link must also be specified.

5.6.2.2 Derivation of Model

The different bases on which links can be made clearly divide into two: the connections and the effect propagation. The connections-based links are straightforward. They are defined by the existence of electrical connections or magnetic coupling, and by the one/three phase nature of plant subunits. The “three phase connection” criteria is included in the links between all subunits of plant types other than transformers, 3-power-winding transformers and autotransformers. The “three phase connection” for the three types of three phase transformer is restricted to the links between windings. Similarly the links between the windings of single phase transformers include the “one phase connection” criterion. “One phase connection” is also included in the connection between the windings and their neutrals of three phase transformers, autotransformers and 3-power-winding transformers. Busbars have no “one phase connection” links. As single phase junctions and reactors can be connected between a star point and ground, all the internal links for these plant types include the “one phase connection” criterion.

The links for effect propagation are based on electrical connection and magnetic coupling. The links between plant subunits based on three phase electrical connection do not change the nature of the currents flowing through the plant. Therefore, for such, both effect propagation criteria are trivial since they are included in the links between all subunits. This applies to all but the three types of transformer.

Since the propagation of currents through transformers, 3-power-winding trans-

formers and autotransformers is due to magnetic coupling, the links for 'phase_effects' and 'neutral_effects' propagation are generally non-trivial and dependent on the winding arrangements. The only trivial case is the single phase plant item. Links exist between the windings, and between each winding and its neutral. Since one current induces one field which induces one current, the nature of the currents is identical, and both effect criteria are included in all subunit links.

By a similar token, the applicability of the 'phase_effects' criteria to the links between windings of three phase transformers, autotransformers and 3-power-winding transformers will be shown. An increased current magnitude in any winding will induce an increased magnitude field. By the ampere turns rule, all windings coupled by this field will experience a corresponding increase in current (ie. $I_1N_1 = I_2N_2 = \dots$). Hence, the 'phase_effects' propagate between winding subunits for all winding arrangements. The 'phase_effects' are defined as an increase in the magnitude of some of the three phase currents, without any effect on the neutral current (ie. there is no zero phase sequence component). Therefore, the 'phase_effects' in the currents of a winding are not propagated to the winding neutral. Also, the neutrals of all other windings do not experience the 'phase_effects' since the net field is zero.

The propagation of 'neutral_effects' between subunits of transformers, 3-power-winding transformers and autotransformers could also be proved by argument. However, since 'neutral_effects' propagation describes the flow of a residual or zero sequence component, the internal link knowledge can be derived from the zero phase sequence representation of the various winding arrangements. These representations have themselves been derived by argument in relation to the currents and fields. The modelled links together with the zero sequence equivalent networks are shown in Figure 5.6. Corresponding connections exist

Winding Arrangement		ZPS Equivalent	Knowledge of internal links for neutral effects
HV	LV		
			hv_wdg – lv_wdg hv_wdg – hv_ntrl lv_wdg – lv_ntrl
			hv_wdg – hv_ntrl
			hv_wdg – hv_ntrl
			X
			X
			X

Figure 5.6: Transformer winding arrangements with their zero sequence equivalents and knowledge of internal links for neutral effects

with the HV and LV windings interchanged. A grounded star provides a path for a residual component to flow to ground. Currents can only flow if ampere turn balance is possible. Therefore residual components will only flow to ground if another winding provides a route for zero sequence currents to flow to ground or circulate. A link for 'neutral_effects' between a winding and its neutral is conditional on this ampere turn balance being maintained by another winding. A 'neutral_effects' link between different windings exists if there is a path to ground on both windings (grounded star/grounded star). This knowledge can be seen in Figure 5.6, and is incorporated into the expert system. Although the figure relates specifically to normal transformers with two power windings, the

above knowledge can be applied to the three types of transformer, and has been in the expert system.

The following rules exemplify the above knowledge. They define the conditions for a (2 power winding) transformer HV winding to be linked to its neutral. A brief description follows the rules.

```
tx_link_model(hv_wdg,hv_ntrl,3,star,ntrl_con is _,LVWdg,LVE,Ter,
    Conditions,imped(Z,_),Z):-
    not((LVWdg=star,LVE=no,Ter=no)),
    intersection(Conditions,[propagation is neutral_effects,
        connect_1ph]).
```

```
tx_link_model(hv_wdg,hv_ntrl,3,star,ntrl_con is _,star,no,no,
    Conditions,imped(Z,_),Z):-
    intersection(Conditions,[connect_1ph]).
```

The first two arguments of the `tx_link_model/11` rule head contain the linked subunit names. The next argument holds the number of phases, and is followed by the HV winding arrangement, the HV neutral connection, the LV winding arrangement, the LV neutral connection, and tertiary winding existence. The ninth argument holds the basis of the link being tested. The final two arguments hold the transformer impedance term and the positive sequence impedance between the linked subunits. Clearly then the first clause shown relates to three phase transformers with an earthed star HV winding. The rule body of this clause requires that the transformer is not one with an unearthed star LV winding and no tertiary. The basis for this connection being established can be either 'neutral_effects' propagation or one phase connection. Satisfaction of this clause corresponds with the cases in the first three rows in Figure 5.6. The second clause defines the same link on the basis of one phase connectivity provided the LV winding is an unearthed star and there is no tertiary. The absence of a 'neutral_effects' propagation based link in this case is shown on the fourth row of Figure 5.6.

5.6.2.3 Additional Tier of Model

The described model, by which internal links via individual subunits are used in the connectivity searches, is satisfactory for most types of plant. However, a busbar can have many, highly interconnected subunits, some of which may be without external connection. Also, the interconnection is by manually controlled isolators (by definition of plant items) whose status needs to be checked when performing connectivity searches. Such knowledge is used repetitively during the qualitative simulation. It was recognised that the processing time spent tracing through busbars was undesirable due to its significance, and because the status of the manually controlled isolators would remain unchanged.

The addition of a less detailed hierarchical level to the model was realised as a solution. The models for all plant types except the busbar and junction types are equivalent in the two models. Since it was felt unreasonable to be repetitively checking the manually controlled isolators' status in disturbance diagnosis time, an undetailed representation of a busbar (and junction) is maintained. A busbar's undetailed representation consists of one or more facts listing together those subunits which are connected by manually controlled isolators. The detailed and undetailed representations of a busbar are shown in Figure 5.7. (This




Figure 5.7: Illustration of the two levels of busbar model

hierarchical distinction has also been used in other problem solving tasks such as

reliability [80].)

The status of isolators does not always need to be checked in disturbance diagnosis time. The detailed representation is used to resolve a new undetailed representation only when manual control has been exercised and an isolator has been opened or closed—generally, this will not be during a disturbance. The undetailed representation indicates that subunits are connected to others, but does not define the route. Therefore, in order to find a busbar’s unit protection CTs, the detailed representation needs to be used since the relevant CT may not always be on the externally connected terminals. However the benefit of the undetailed model can be seen in the search for a power infeed.

The undetailed model is used when initially searching for a power infeed. Thereby a path to a machine or infeed from undefined plant can be found quicker. The detailed model is used to fully define the route only when the fault is on the busbar and a fault path has been found.

5.6.3 Protection Detection Logic

For each protection type, corresponding modelling knowledge has been developed. The scope of the modelling knowledge encompasses two aspects of protection behaviour. Knowledge is required to determine the area of the network protected by a particular device. Also, the measure by which a device detects a fault condition is defined as either the 'phase_effects' or 'neutral_effects', and a level of such ('average' or 'high') is specified as appropriate.

The models for determining the protected area are predominantly qualitative, and therefore possess varying degrees of realism.

5.6.3.1 Unit Protection Models

Unit protections basically operate on the principle of operating for (internal) faults between associated CTs, and not operating for all others (external faults). Therefore unit protections are well defined since their detection logic is relatively straightforward. The same principle applies to unit protection modelling knowledge: determine the plant contained within the area bounded by the CTs burdened by the relay.

5.6.3.1.1 Differential: `diff_cc` There are different means of achieving a protection scheme with differential characteristics, such as circulating current and balanced voltage arrangements. The detection circuit in a circulating current scheme is supplied with the sum of the secondary currents from CTs. The CTs enclose some part of the network and so, by Kirchhoff's current law, will normally have a net zero current. A fault on the enclosed plant will upset this balance, causing a current to flow in the detection circuit. Balanced voltage schemes behave similarly, but with a CT secondary voltage imbalance causing a current in the detection circuit. In addition, restrain circuits and summation/interposing transformers are used by various schemes. Since they all function in a similar way, one relay type is used for all. Every burdened CT in the scheme is a three phase CT. A simplified protection arrangement is shown in Figure 5.8 with an




Figure 5.8: Protection scheme using the differential relay type: (a) schematic; (b) knowledge base schematic

illustration of its representation in the knowledge base. The knowledge base contains the fact that 'average' 'phase_effects' are the detected measure.

5.6.3.1.2 Restricted Earth Fault: rest_zps The arrangement of restricted earth fault protection is very similar to the differential scheme. It is used to protect a transformer winding or that of a 3-power-winding transformer, or the connected windings of transformer banks. The feature which distinguishes restricted earth fault from differential protection is the fact that the former has only one detection coil. Consider the arrangement shown in Figure 5.9. Since there is only




Figure 5.9: Protection scheme using the restricted earth fault relay type: (a) schematic; (b) knowledge base schematic

one detection circuit in the shown arrangement, 'phase_effects' will not be detected. The detection circuit is supplied by a neutral current and by the sum of three phase currents (ie. the residual component in the three phase currents). Clearly then the relay can detect 'average' 'neutral_effects'. Like the differential scheme, the detection circuit is supplied with a difference of currents. Any residual component in the phase currents under normal conditions would be summed at the star point to produce an equivalent neutral current. The net current supplied to the detection circuit would therefore be zero. However, with a fault between the CTs, the resulting imbalance would supply a current to the detection circuit.

5.6.3.2 Non-unit Protection Models

In order to accurately model non-unit protections, extensive numerical computations in the form of fault flows would be required. Due to the time constraints on the diagnosis, and the number of times the calculations would be required, performing these calculations was considered unreasonable.

Protection relays burdening CTs on three phase plant can detect faults if the measured current, sum of measured currents, or ratio of VT secondary current to CT measured current satisfy certain conditions. This has been represented by two conditions. The first is where the fault is located within a protected area of the network, and the second is that CTs are between the fault location and a power infeed. The second condition is met in the qualitative simulation by only considering protection once its CTs have been found on a route for fault current. The derivation of the protected area is performed by modelling knowledge, and is described in the following paragraphs. Protection with a CT on the plant between a star point and ground (ie. a neutral overcurrent relay) is the exception to this general rule. The transformer whose neutral current these relays detect does not need to be on the fault current path for a fault to be detected. This characteristic is included in the qualitative simulation knowledge. The protected area is described in the modelling knowledge. What was required was a means of gauging the protected area approximately, and doing so with minimal numerical computation—as an expert protection engineer might do.

Non-unit protections are generally used for feeder protection. The settings of the relays are often designed to cover faults on the immediate feeder, and/or to protect the feeder beyond the adjacent substation in a backup capacity. Therefore the first means of determining non-unit protected zones to be developed was by measuring substations away from the relaying point. Each protection relay type had a reach (allocated for the type or set by a specific argument in the relay

fact) which indicated the extent of the protected area. If the reach was not a whole number then the protected area included the substation beyond the furthest “reached” feeder. Thus a protected area stretching from the relaying point in one substation to the adjacent substation would be indicated by a reach of 1, and an area including the remote substation would be produced by a reach of 1.5.

This means of deriving the extent of non-unit protected areas was satisfactory up to a point. In an interconnected network there will be several lines beyond the adjacent substation for which a relay acts as backup protection. These various lines would have different lengths and hence different impedances, and will therefore be protected by the remote non-unit relay to varying degrees. By measuring in substations from the relaying point, this fact is ignored.

The use of cumulative impedances rather than substations was adopted. An example of an area defined by cumulative impedances is illustrated in Figure 5.10. Impedances are already available for transformers, lines, etc. since they were re-




Figure 5.10: Illustration of a protection reach defined by cumulative impedances
 quired for the loadflow. Connected plant is progressively identified away from the relaying point, and the impedances of the connected plant added. The progress

is halted when the cumulative impedance satisfies a set condition. Backtracking then commences in order to identify alternative paths. Therefore the cumulative impedances is the impedance along a path ignoring all other paths. This effect is in evidence in Figure 5.10.

5.6.3.2.1 Distance: dist Distance protection is the protection type which is most closely related to the developed method of modelling. The detection logic of a distance relay is to “measure” the impedance seen looking into the network from the relaying point. The detection logic is therefore based on ‘average’ ‘phase_effects’. Three phase voltage transformer and current transformer signals are used to derive an impedance measure by which the relay ascertains the existence of fault conditions. Various characteristics are used in distance relays, the simplest of which is the impedance characteristic. The impedance characteristic is a straightforward impedance magnitude comparison eg. $|Z| > 5 \times 10^{-3}(pu)$. Although the circuitry is complex, the basic schematic of a distance relay scheme is trivial. It is none-the-less shown in Figure 5.11.




Figure 5.11: Protection scheme using the distance relay type: (a) schematic; (b) knowledge base schematic

The modelling knowledge of cumulated path impedances could comply with actual characteristics if faults were 3 phase faults, and contributions via branching points were negligible. Ignoring the effect of load current and topological changes off branching points curtails the reasoning of the expert system. However such effects would require a numerical fault flow in order to be dealt with properly.

Having roughly located the fault by reasoning, a numerical analysis could be performed. This may identify reasons for strange patterns of behaviour caused by unforeseen or uncommon network configurations or load levels.

5.6.3.2.2 Directional Overcurrent: dir_oc Directional overcurrent relays recognise fault conditions by the magnitude of the current supplied by its CT(s). Every CT is on three phase plant. The device therefore detects fault conditions on the basis of 'average' 'phase_effects'. If the current flow is in the relevant direction, the relay will operate when the detection circuit "sees" a current which is above its set characteristic. Clearly, therefore, a fault flow analysis is required to accurately model the detection characteristics of such devices. Since this was not a valid option, the cumulative impedance technique is used—the impedance to the fault is clearly related to the fault current which will flow. A simple impedance characteristic (ie. boundary condition of the form $|Z| > 5 \times 10^{-3}(pu)$) is used. The direction in which faults are detected is identified by an argument in the relay fact with the plant on the reverse side of the relaying point. A directional overcurrent arrangement which involves two CTs is shown in Figure 5.12. This arrangement protects both the feeder and the feeder circuit 'T' point.

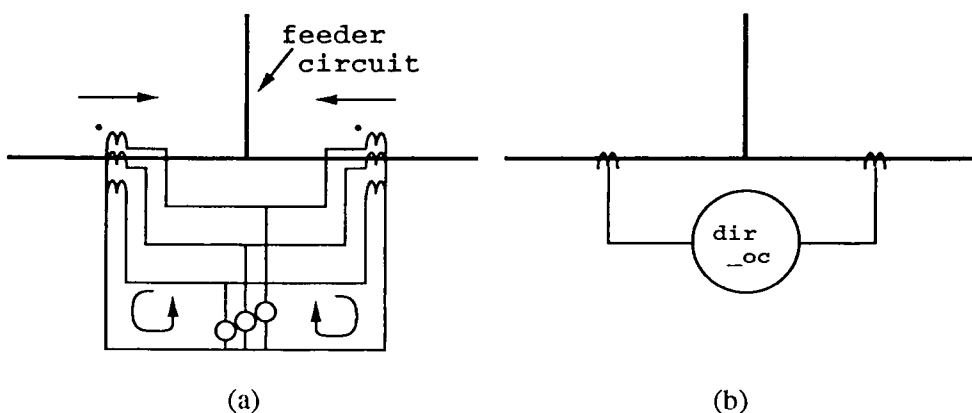


Figure 5.12: Protection scheme using the directional overcurrent relay type: (a) schematic; (b) knowledge base schematic

The resolution of an effective impedance reach can be rigorously derived from

the overcurrent setting. However, the function of this expert system is not to audit the utility's protection scheme; though with some modifications it could fulfill this role. The effective reach can therefore be derived from the functional specification of the relay eg. protect to a distance of two substations away from the relaying point on the shortest adjacent feeder. Thereby the cumulative impedance reaches could be set manually, or automatically by rules involving a topological search.

For some characteristics, the limit to which they could detect a fault is unclear. The Inverse Definite Minimum Time (IDMT) characteristic (an example is illustrated in Figure 5.13) would appear to indicate an almost limitless reach




Figure 5.13: IDMT current-time characteristic

due to its asymptotic nature at low currents. For example, point *A* on the curve would probably represent a fault on plant several substations away from the relaying point. Setting the reach this far can be avoided due to other relays in the protection scheme. Backup protection would normally exist at every substation, and have settings such that protection nearest the fault would operate first. Therefore, for the said protection to operate at point *A* on its characteristic would

probably require several simultaneous device failures. Simultaneous failures are relatively unlikely; an arbitrary limit on the number of simultaneous failures to be considered by the expert system is set. Based on this, and the minimum number of backup relays per substation, a maximum substation reach can be applied to all protection relays. Characteristics which appear to exceed this limit can be confidently curtailed.

5.6.3.2.3 Overcurrent: nondir_oc Overcurrent relays are equivalent to the latter relay type, except that they can detect faults on both sides of the relaying point. The detection means are 'average' 'phase_effects'. Therefore cumulative impedance reaches are set in the same way but are applied to both directions from the relaying point. A scheme is shown in Figure 5.14.




Figure 5.14: Protection scheme using the non-directional overcurrent relay type: (a) schematic; (b) knowledge base schematic

5.6.3.2.4 High Set Overcurrent: hs_oc High set overcurrent relays are also similar to the directional overcurrent type. The difference is that such high set overcurrent relays have high settings, with the result that ground faults are not detected by them (except for close up faults). The expert system has this knowledge in the form that 'high' 'phase_effects' are detected. The knowledge concerning close up faults could be contained in the deeper diagnosis knowledge. A scheme is shown in Figure 5.15.

5.6.3.2.5 Neutral Overcurrent: ground_oc The actual relaying device identified by this relay type may have equivalent characteristics to that on one phase




Figure 5.15: Protection scheme using the high set overcurrent relay type: (a) schematic; (b) knowledge base schematic

of a 'dir_oc'. However, because of its functional difference, it is attributed to a different type. The burdened CT is on plant between a star point and ground, and therefore supplies a proportion of the neutral current to the relay. The detection of faults is consequently based on 'average' 'neutral_effects'. A representation of a neutral overcurrent protection arrangement is shown in Figure 5.16.




Figure 5.16: Protection scheme using the neutral overcurrent relay type: (a) schematic; (b) knowledge base schematic

Being a zero sequence current, the fault current "seen" by such relays is dependent on zero sequence impedance. However, since the reach is specified based on the device's function rather than its setting, it was considered reasonable to specify the reach in terms of an equivalent positive sequence impedance. The same technique of cumulative impedances as used by the directional overcurrent type is adopted.

5.6.3.2.6 Instantaneous Earth Fault: zps Instantaneous earth fault relays are also basically overcurrent devices measuring the 'average' 'neutral_effects'. However the residual component measured by the detec-

tion circuit is the sum of the currents from CTs on the three phases, and not the current from a CT on a neutral connection. Figure 5.17 contains schematics of such a protection scheme. Again, cumulative impedances are used to define a representation of the protected area.




Figure 5.17: Protection scheme using the instantaneous earth fault relay type: (a) schematic; (b) knowledge base schematic

5.6.4 Protection and Switchgear Activity

The protection detection logic is used in the qualitative simulation to ascertain if a protection relay would operate. If the test is satisfied, and the relay would be expected to operate, then the operation is simulated. The secondary activity which would be caused by such primary activity must also be simulated, and this includes relay and switchgear activity. Likewise the ability to simulate the activity of devices following a false operation is required. The knowledge used to simulate this device activity is described in this section.

There are two main categories of activity:

- “actual activity”:- This relates to activity which, it is believed, would have taken place under the given hypothesised conditions.
- “possible activity”:- “Possible activity” is that whose occurrence is a possibility.

In certain cases, the classification as actual or possible activity may not have a sound basis; a class called “default activity” is also used. “Default activity”

is applicable where the system response to an event is as yet unclear due to there being alternative processes which may occur. The “default activity” and its alternative are related in a fact added to the knowledge base—this fact is part of the extra knowledge contained in scenario facts—and the “default activity” is simulated as “actual activity” on the basis that it would occur.

The “possible activity” category was introduced in order to cope with the uncertainty involved in the clearance of ‘temporary’ faults. Since a ‘temporary’ fault may be extinguished during any of the dead times of an autoreclose scheme, there is possible activity as well as expected activity. (The dead time is the period when all the breakers are open and the faulted plant is temporarily isolated.) Considering how these activity categories are modelled:

- “actual activity”:- A fact describing the simulated status change is asserted. A new simulated status fact is also added to the knowledge base, with any obsolete fact being removed.
- “possible activity”:- This activity is not simulated. A fact grouping possible activity alarms is simply stored in a similar way to default and alternative details.

This knowledge is initiated by the isolation simulation rules, the false activity rules, and by the post-isolation rules. The application of this activity modelling to specific device types will be included in the following descriptions of the modelling of relay activity and switchgear activity.

5.6.4.1 Relay Activity

When a relay operation is simulated, the operation is dealt with as “actual activity”. In addition, if the relay has an automatic resetting capability, the relay resetting will also be simulated when the post-isolation simulation rules are fired (resetting occurs after any fault has been isolated).

As well as its own activity, the relay activity model simulates the activity caused by the relay operation. Trip or secondary relays may be controlled by the relay whose operation was simulated. Details of the local trip and control wiring are provided in the relay data together with the logic of any additional auxiliary switches in the wiring. The trip and control wiring is traced in a depth-first manner. A secondary relay controlled by the simulated operation is identified as is the logic of any intervening auxiliary switches. The logic is tested against the simulated state of the network. If the logic is satisfied, the trip relay has its operation likewise simulated, and any trip relays it controls also tested for simulation. If the auxiliary switch logic is not satisfied, then another trip relay is sought. When switchgear is reached, its activity is simulated before the reasoning traces back along the wiring in search of more secondary activity.

Before any relay which should operate has its operation simulated, the knowledge base is checked to see if the relay is a known or hypothesised failed device. If the device is in a fit state, the procedure progresses as described above. If the device is a “failure” then its operation is not simulated. However, the reasoning proceeds through the trip and control wiring, without simulating activity, in order to determine if the failed device would have caused a breaker to trip had it not failed. This latter feature was introduced in order to relate individual failures with backup activity.

Originally, intertrip schemes were dealt with in the same way as local trip and control schemes. However it was recognised that local and remote processes can race against each other in some schemes, and that the resulting activity following a fault (with all devices perfectly operational) can sometimes be different. Consider, for example, a two-way intertrip scheme. A fault on the feeder may cause the protection at one substation to operate, sending an intertrip signal to the remote end. The remote end breaker may then trip due to the operation of the intertrip

receive relay or before this due to its local protection operating.

This scheme level knowledge is used during a simulation once a protection relay has been identified via its CT but before it is simulated. This is required since the protection relay found might not operate: for the described case, the path found may be through the substation which will clear the fault due to the intertrip receive signal. The related alarms which have already been received govern the activity which is simulated by the expert system. If there are no such received alarms, then the process defined as the “default activity” is simulated and the default and its alternative are recorded for future reference. By this means, the simulation is responsive to the available data, but makes assumptions when the data has not yet been received.

To illustrate this knowledge, a particular hypothetical two-way intertrip scheme will be looked at. The scheme shown in Figure 5.18 is represented by the scheme details shown:

```
relay_scheme(p at ss1,'it-1-2',1).
relay_scheme(r at ss1,'it-1-2',2).
relay_scheme(p at ss2,'it-1-2',1).
relay_scheme(r at ss2,'it-1-2',2).

scheme('it-1-2',Fault,
  [defn(ss1,racing(control(p at ss1,r at ss2,true),
    simultaneous(p at ss1,p at ss2,Fault),not_exclusive,1)),
    defn(ss2,racing(control(p at ss2,r at ss1,true),
    simultaneous(p at ss2,p at ss1,Fault),not_exclusive,1))]
).
```

The scheme knowledge would be used in the qualitative simulations following the operation of the protection relay 'p at ss1'. Suppose this operation has occurred and that, during a subsequent simulation, a fault path through 'ss1' is found first. The operation of the protection relay ('p') at substation 'ss1' is simulated since its alarm has been received. As a result, the operation of the

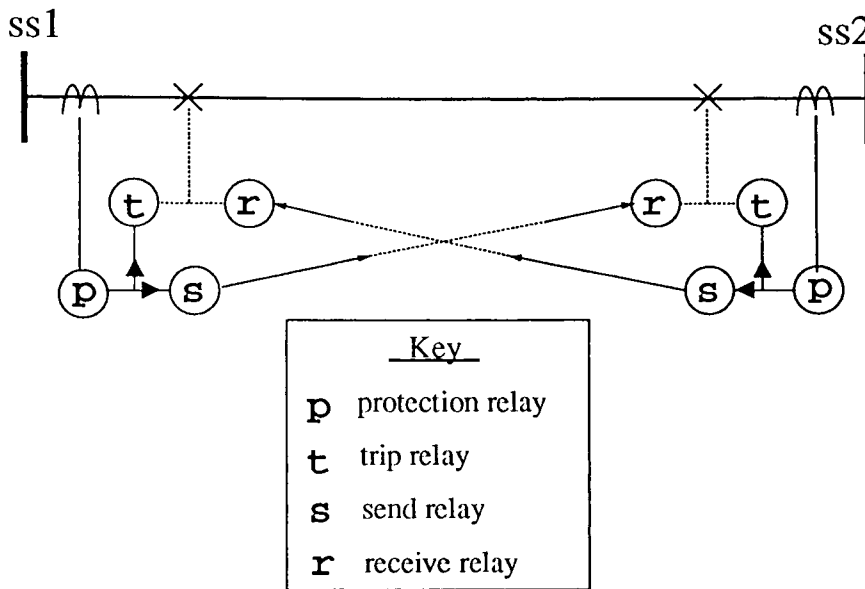


Figure 5.18: The abstract schematic of a two-way intertrip scheme

trip relay ('t') and breaker at substation 'ss1', and, by default, that of the intertrip receive relay ('r') and breaker at substation 'ss2' are simulated. The fact that the intertrip receive and protection relays at 'ss2' are non-exclusive default and alternative respectively is also stored. Suppose the alarm describing the operation of the protection at 'ss2' is subsequently received. Then, because the received activity is an alternative to a default, the fault is resimulated. This time the 'ss2' protection relay operation is included, and the intertrip received trip relay operation is included as possible since the alternatives are not mutually exclusive.

The knowledge to simulate the resetting of automatically resetting relays is contained in the post-isolation production rules.

Having simulated the operation of any relay, as well as looking for controlled relays, the reasoning also looks at the relay details for controlled switchgear. The activity of controlled switchgear is simulated.

5.6.4.2 Switchgear Activity

The behaviour of switchgear can also be simulated in terms of “actual activity” and “possible activity”. The simulated activity relates to switchgear trip/open and reclose/close operations. Like relays, if a breaker or isolator is known to be a failed device or a hypothesised failure, then its activity is prevented.

The simulated tripping/opening and reclosing/closing of switchgear is as “actual activity” and involves the assertion of both status and status change facts. Circuit breakers with an automatic reclosing capability are also simulated by post-isolation production rules.

For the particular case of autoreclosing breakers, the permanence of the fault influences the simulated activity. The operation of an autoreclosing breaker may involve several alarms: following a ‘permanent’ fault, the breaker would go through all its trip/reclose cycles until it locks out; following a ‘temporary’ fault, the breaker would go through one or more trip/reclose cycles (ending with a reclosure) since the fault may be cleared during any of the dead times. The simulated activity knowledge generates the facts to represent this in the expert system. If the hypothesised fault is ‘permanent’, then a trip is simulated as “actual activity” during the isolation simulation, and tripped and reclosed activity is simulated for every cycle the breaker can go through by post-isolation rules. A trip is also simulated as “actual activity” for the case of a ‘temporary’ fault. In addition, however, one reclosure is simulated as “actual activity”, and tripped and reclosed alarms for the remaining cycles (ie. one less than the total cycles) are stored as “possible activity”. Tripping due to false activity will result in a reclosure and no more activity since there is not (and never was) any fault current. The false activity rules therefore simulate this as “actual” tripped, and post-isolation rules simulate the reclosed activity.

5.7 Chapter Summary

To summarise, this chapter has shown the significant network knowledge contained within the fault diagnosis expert system. In the first instance, data is required on-line which describes protection and switchgear activity. The SCADA system is the source of this dynamic data. Details of the power network and of the protection and switchgear equipment to which the expert system is being applied is also required. This knowledge is largely static and is presently contained in an internal database. A result of the research undertaken by the author has been to develop a hierarchical structure for the power system data, and this was described in the chapter.

This chapter has also demonstrated the extent of network knowledge contained in the expert system which relates to the behaviour of power systems in general. The qualitative simulation knowledge provides the means to ascertain the protection and switchgear activity which would be expected to follow a particular fault given the system configuration. It has been shown how the qualitative simulation makes use of two other areas of knowledge. The need for and the nature of the loadflow knowledge is described as are the considered alternatives to that knowledge. Finally the modelling knowledge for system faults, plant internal links, protection detection logic, and protection and switchgear activity have been described.

Chapter 6

Use of the Fault Diagnosis

6.1 Case Studies

6.1.1 Test Power System and Cases Used

The early developments in the fault diagnosis expert system reported in this thesis were tested on a small hypothetical transmission/distribution network. Data was later obtained detailing the electrical power network and the protection system for a South of Scotland Electricity Board nuclear power station (now operated by Scottish Nuclear). The station's main electrical power network is not large. However, the trip and control wiring is extensive due to the amount of plant in such a large electrical-mechanical-nuclear system as well as the significance of disconnecting it and the important safety aspects. Unit protections predominate in the system, and therefore the location of possible electrical faults, given an operated protection relay, is generally not a particularly onerous task. The protection relay alarm can, however, be hard to find in the voluminous alarms caused by extensive secondary plant and control activity. Additionally, maloperations and missing alarms are both difficult to identify in a system producing so many alarms. The alarms are gathered by a Protection Performance Monitor, and are of the form shown:

[[[[10 18 55.812 31/05/1989]]]]
Generator No.2.
CBX820 400KV GRID CIRCUIT BREAKER X820
Contact opened.

Due to corporate changes, the continuity of this particular application became inappropriate. Attention then shifted to the ScottishPower transmission/sub-transmission system. The portability of the core diagnostic knowledge was demonstrated in that little change was required in the central diagnosis module. In contrast with the power station system, this network is very extensive, includes many unit and non-unit protections and could have disturbances involving a large

area of the network. The alarms which are available for the expert system are provided by the SCADA system. They are time-tagged (with the event time) at the outstation and are of the form shown by the following example:

15:47:58:02	TORN4	X220	OPEN	TORN4
-------------	-------	------	------	-------

Some network and protection details have been supplied for the power system. Since the system data has all been provided in paper format, the process of realising the system details in the knowledge base was very tedious. Consequently, only part of the power system details have been represented. For a utility which was seriously interested in implementing the expert system, it would be possible to translate the majority of the necessary details from files used by other power applications software or from corporate databases. Moreover it is becoming possible to extract some suitable data from Computer Aided Design (CAD) drawing files; this is a means which will bring utilities' diverse engineering and business functions much closer together.

Ideally, the expert system would be validated by testing its response to genuine records of many diverse system disturbances. This was not however possible. Only paper records of SCADA alarms are maintained by ScottishPower. The control centre does not normally have the means to store SCADA activity on a magnetic medium. Attempts were made by control centre staff to capture a record of SCADA activity on floppy disks. This was done non-intrusively by connecting a personal computer to the line between the SCADA computer and the printer, but failed to capture any significant disturbances.

Many cases were tested on different portions of network during the development of the expert system. Both hypothetical and genuine networks were used. A utility protection engineer helped in the creation of some of these studies.

The case studies presented in this thesis are not the most common system incidents but realistic in terms of the situations they represent. Disturbances are relatively simple and straightforward in the majority of cases. However, the aim

of the work reported in this thesis was not merely to be able to diagnose these common simple fault events. Rather, the primary goal was to develop a system which would, as well as diagnose simple faults, provide help where it was most needed. The system required the ability to diagnose the less common but more complex (and hence more serious) disturbances. The case studies which are included in this chapter are intended to demonstrate the ability of the expert system and illustrate some of the differences in the developed diagnostic strategies.

6.1.2 Case 1

Consider the portion of transmission network shown in Figure 6.1—a small part of that held in the knowledge base. The number 2 275kV interconnector between

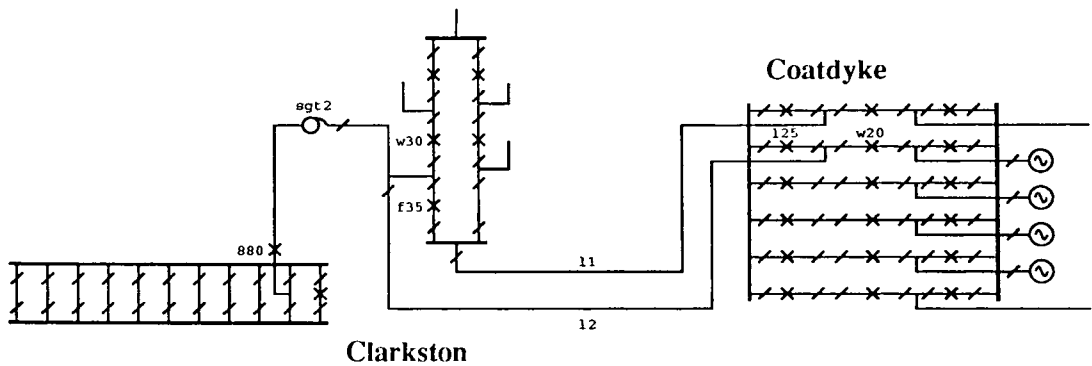


Figure 6.1: A portion of a transmission network

the Coatdyke and Clarkston substations has two main protections: pilot wire and distance protection with acceleration. This first study relates to a relatively common type of disturbance: a fault on a line (number 2) which is fed from both sides and is isolated with all equipment functioning correctly. Following a phase fault on the middle of the line between Coatdyke ('coat') and Clarkston ('clar'), the two main protections operate, acceleration signals are sent by both distance relays, trip relays at the two substations operate, and the line is isolated by the tripping of the breakers at the two ends of the line. The telemetry which is listed below shows that which would be generated by a fully functioning

SCADA system:

12:45:24:53	COAT2			CLAR2 NO2	FIRST MAIN PROT OPTD	ON	COAT2	i
12:45:24:54	COAT2			CLAR2 NO2	SECOND MAIN PROT OPTD	ON	COAT2	
12:45:24:54	CLAR2			COAT2 NO2	FIRST MAIN PROT OPTD	ON	CLAR2	
12:45:24:54	CLAR2			COAT2 NO2	SECOND MAIN PROT OPTD	ON	CLAR2	
12:45:24:54	COAT2			CLAR2 NO2	TRIP RELAYS TO BE RESET-E	ON	COAT2	
12:45:24:55	CLAR2			COAT2 NO2	TRIP RELAYS TO BE RESET-E	ON	CLAR2	
12:45:24:55	CLAR1			SGT2	TRIP RELAYS TO BE RESET-E	ON	CLAR1	
12:45:24:56	COAT2			CLAR2 NO2	ACCELN RECEIVED	ON	COAT2	
12:45:24:56	CLAR2			COAT2 NO2	ACCELN RECEIVED	ON	CLAR2	
12:45:24:56	COAT2	W20	OPEN				COAT2	
12:45:24:56	COAT2	L25	OPEN				COAT2	ii
12:45:24:57	CLAR2	F35	OPEN				CLAR2	
12:45:24:57	CLAR2	W30	OPEN				CLAR2	
12:45:24:58	CLAR1	880	OPEN				CLAR1	iii
12:45:24:62	COAT2			CLAR2 NO2	SECOND MAIN PROT OPTD	OFF	COAT2	
12:45:24:63	CLAR2			COAT2 NO2	SECOND MAIN PROT OPTD	OFF	CLAR2	
12:45:24:63	COAT2			CLAR2 NO2	FIRST MAIN PROT OPTD	OFF	COAT2	
12:45:24:63	CLAR2			COAT2 NO2	FIRST MAIN PROT OPTD	OFF	CLAR2	
12:45:24:63	COAT2			CLAR2 NO2	ACCELN RECEIVED	OFF	COAT2	
12:45:24:63	CLAR2			COAT2 NO2	ACCELN RECEIVED	OFF	CLAR2	

The labelling down the right-hand-side of the above identifies the sections of the SCADA listing as follows:

- i. protection and trip relay operations
- ii. switchgear tripping
- iii. relay resetting

These SCADA alarms are read in one at a time by the expert system. It is assumed that there were no disturbances before the above alarms. The process of diagnosis according to both the “wait-and-see” and the (third) “hypothesiser” strategies shall be detailed for this case study.

6.1.2.1 Wait-and-See Diagnosis

On receipt of the first alarm, the “gut reaction” recognises a primary alarm which is not related to an old disturbance and is neither wrong (ie. the relay was in the OFF state according to the knowledge base) nor definitely expected. Consequently, the central diagnosis uses the relay’s protection modelling knowledge to store the location and type of faults which the relay can detect. Checks are then done on the electrical status of the plant where the protection’s current transformers (CTs) are located, which reveal it to be live and thus indicate that the

operation of the relay is reasonable. The next three alarms are dealt with in an equivalent fashion. Following the receipt of the trip relay operation alarm from 'coat2', the "gut reaction" attempts to account for the activity by looking for an operated controlling device (from among the first and second main protections at both substations). The first main protection at 'coat2' fulfills this role, and the relationship is asserted in a fact. The "central diagnosis" simply updates the records such as that of expected breaker trips. The alarms describing the operation of the corresponding trip relays at 'clar2' and 'clar1' are dealt with in the same way.

On receipt of the alarm relating to the tripping of breaker 'w20 at coat2', the "gut reaction" checks for justification for the breaker trip. The breaker trip is found to be caused by the trip relay operation at the same substation. Checking the electrical status of both of the breaker's isolated terminals reveals them both to be 'live'. The "central diagnosis" therefore performs no significant reasoning. Nor does it perform significant reasoning for the next three breaker trip alarms since they all evoke the same "gut reaction" judgement as did that for 'w20 at coat2'.

With the tripping of breaker '880 at clar1', the topological change results in the creation of a blackout area. This is recognised by the "gut reaction" which resolves the status of the breaker terminals to be 'live' and 'dead'. Since the blackout area is not adjacent to another, a new disturbance is judged to have taken place. The "central diagnosis" now makes its major contribution to the reasoning process. The related relays are identified and their recorded lists of detectable faults gathered. The intersection of these fault lists is obtained, and is found to include a phase and a ground fault on the number 2 line between Coatdyke and Clarkston—both '(permanent or temporary)'. No failed devices are identified; both scenarios only include a fault, the supportive alarms and predicted reset

alarms. The remaining reset alarms merely strengthen these scenarios so that the conclusion (as output by the expert system) is:

* Messages to 11 jan 1990 12:45:24:63 read. Disturbance -19999

1. There was a ground fault_on coat2 to clar2::l2 which_is (permanent or temporary), which the protection equipment cleared.
2. There was a phase fault_on coat2 to clar2::l2 which_is (permanent or temporary), which the protection equipment cleared.

2 possibilities.

(The ordering of these two scenarios is based on the comparative relative likelihoods of a 'phase' or 'ground' fault occurring.)

6.1.2.2 Hypothesiser Diagnosis

The hypothesiser-based expert system comes to its conclusion in a different way. Having received the first alarm, indicating the operation of a pilot wire relay at Coatdyke, the "gut reaction" judges the alarm to relate to a new disturbance, and creates a disturbance area of involvement around Coatdyke. As a result the "central diagnosis" attempts to generate hypotheses which account for the received alarm. Firstly, the relay's protection modelling knowledge is used to obtain the location and type of faults which the pilot wire relay can detect. Hypotheses of a phase fault and a ground fault on the line are thereby generated. To these are added the permanence 'permanent or temporary' because the relay does not belong to an autoreclose scheme (and therefore no difference between 'permanent' and 'temporary' faults will be apparent). With no failures added, each hypothesis is simulated in turn. Both result in the hypothesised faults being isolated by the simulated operation of the line's first and second main protection relays and their associated trip relays and circuit breakers. In the process, acceleration signals are also generated by the distance relays. A further scenario is added which is based on the hypothesis that the alarmed activity was false. The three scenarios have their relative likelihoods calculated and are ranked accordingly.

The next alarm is judged by the "gut reaction" to relate to the existing distur-

bance. Therefore, the “central diagnosis” updates those existing scenarios which predicted the operation of the second main protection relay. Both fault-based scenarios are strengthened thereby. Since a scenario remains to be updated, the expert system attempts to find a reason for the latest operation which would form the basis of partial scenarios. All of the detectable faults are already in scenarios and so no fault-based partial scenarios are created. However a false operation partial scenario is successfully generated, and combined with the only remaining old scenario. Again three scenarios exist.

The following two protection operation alarms evoke the same response. Additionally, all three scenarios predict the subsequent trip relay operations and the breaker trips. Therefore the scenarios are all strengthened following the alarms, and the appropriate “predicted” alarms are changed to “supportive” accordingly. Following the change of the 'tripped(880 at clar1)' alarm in the two fault based scenarios to supportive, the expert system recognises that all breakers whose tripping was predicted have opened and remained so; an argument in the scenario fact is therefore changed from 'not_isolated' to 'isolated'. The reset alarms also strengthen all the scenarios, after which, the resulting scenarios are displayed as shown below:

- | |
|---|
| <p>* Messages to 11 jan 1990 12:45:24.63 read. Disturbance -19999</p> <ol style="list-style-type: none">1. There was a ground fault_on coat2 to clar2::l2 which_is (permanent or temporary), which the protection equipment cleared.2. There was a phase fault_on coat2 to clar2::l2 which_is (permanent or temporary), which the protection equipment cleared.3. iacc21 at clar2 reset, iacc21 at coat2 reset, rm1_coat at clar2 reset, rm1_clar at coat2 reset, rm2_coat at clar2 reset, rm2_clar at coat2 reset, rm2_coat at clar2 operated, rm1_coat at clar2 operated, rm2_clar at coat2 operated & rm1_clar at coat2 operated, all falsely.
3 possibilities. |
|---|

The preliminary “deeper diagnosis” does not change the top two scenarios, but suggests ignoring the false operation scenario on the basis of its low likelihood and the close timing of the primary alarms.

6.1.2.3 Brief Comparison

This case demonstrates the “wait-and-see” and “hypothesiser” strategies coming to the same conclusion. Both expert systems correctly concluded a fault on the line between substations 'coat2' and 'clar2'—the identification of the type and permanence of the fault cannot be discerned from the SCADA alarms alone in this case.

6.1.2.4 Consideration of Missing Alarms

Suppose now that, instead of all the alarms listed in section 6.1.2 being received, the trip relay operated and breaker tripped alarms from substation 'clar1' are lost. The “wait-and-see” expert system does not recognise the creation of a blackout since a corresponding isolated area is not now created in the network representation. The diagnostic reasoning in the “central diagnosis” for identifying faults is not even attempted as the triggering condition is not met. No conclusion is reached except that all the reset alarms and primary operation alarms are recorded as false activity.

In contrast, the conclusions of the hypothesiser remain unharmed. The same conclusions, with the additional recognition of the missing alarms, are reached as shown:

- * Messages to 11 jan 1990 12:45:24.63 read. Disturbance -19999
1. There was a ground fault_on coat2 to clar2::l2 which_is (permanent or temporary), which the protection equipment is attempting to clear.
(The alarms for 880 at clar1 tripped & iart at clar1 operated are missing.)
 2. There was a phase fault_on coat2 to clar2::l2 which_is (permanent or temporary), which the protection equipment is attempting to clear.
(The alarms for 880 at clar1 tripped & iart at clar1 operated are missing.)
 3. iacc21 at clar2 reset, iacc21 at coat2 reset, rm1_coat at clar2 reset, rm1_clar at coat2 reset, rm2_coat at clar2 reset, rm2_clar at coat2 reset, rm2_coat at clar2 operated, rm1_coat at clar2 operated, rm2_clar at coat2 operated & rm1_clar at coat2 operated, all falsely.
(The alarms for 880 at clar1 tripped & iart at clar1 operated are missing.)
- 3 possibilities

In fact, the correct system fault was among the scenarios generated by the

hypothesiser after the first alarm; this demonstrates how the strategy is not dependent on receiving every alarm. This, though, is not always so; the correct scenario may not be generated or reach the top ranking positions until later in the development of a disturbance.

This simple case study demonstrates the main advantage of the hypothesiser approach over the wait-and-see approach, namely its ability to cope with incomplete telemetry. The appearance of the above circumstance does not have to be due to telemetry malfunctioning. It is feasible that the devices in 'clar1' might not have been wired into the SCADA system, being associated with lower voltage plant. The dependence on receiving every breaker trip alarm is a major limitation of the described wait-and-see strategy.

6.1.3 Case 2

Figure 6.2 shows a hypothetical portion of network containing 132kV, 33kV and 11kV plant. Substations ss01, ss02 and Drumgelloch are defined




Figure 6.2: A portion of hypothetical network

as supply points. Both 33kV transformer feeders between substations Dumbeth ('dumb') and Mosside ('moss') are protected against ground faults by in-

stantaneous earth fault relays at substation Dumbeth which have intertripping to the remote end breakers. Standby earth fault relays measure the current in the neutrals of earthing transformers connected to the infeed circuits of substation Dumbeth. There is a directional earth fault relay at substation Drumgelloch ('drum'). Clearly, this is not a description of the complete protection system, but rather is some details relevant to this case study.

This second case study relates to the occurrence of a ground fault on the second transformer feeder of substation Dumbeth for the case where breaker 'cb5 at dumb' remains stuck closed. Under this circumstance, the instantaneous earth fault relay operates, its trip relay operates and attempts to trip 'cb5', an intertrip is sent to the remote end which results in the tripping of the breaker at Mossie. The fault persists because breaker 'cb5 at dumb' fails to trip. As a result, the directional earth fault relay at Drumgelloch and the standby earth fault relays at Dumbeth operate, and the fault is isolated when their controlled breakers trip. If the SCADA system functioned perfectly, the following alarms would be received:

10:16:00:32	DUMB			MOSS NO2	INST EF PROT OPTD	ON	DUMB	i
10:16:00:33	DUMB			MOSS NO2	TRIP RELAYS TO BE RESET-E	ON	DUMB	
10:16:00:39	MOSS			T2	INTERTRIP RECEIVE OPTD	ON	MOSS	ii
10:16:00:42	MOSS	CB2	OPEN				MOSS	
10:16:00:84	DRUM			DUMB	DIRECTIONAL EF OPTD	ON	DRUM	ii
10:16:00:85	DRUM			DUMB	TRIP RELAYS TO BE RESET-E	ON	DRUM	
10:16:00:88	DRUM	CB1	OPEN				DRUM	iii
10:16:04:44	DUMB			SS02	STBY E/F PROT OPTD	ON	DUMB	
10:16:04:44	DUMB			SS01	STBY E/F PROT OPTD	ON	DUMB	iii
10:16:04:45	DUMB			SS02	T2 TRIP RELAYS TO BE RESET-E	ON	DUMB	
10:16:04:45	DUMB			SS01	T1 TRIP RELAYS TO BE RESET-E	ON	DUMB	iii
10:16:04:48	DUMB	CB2	OPEN				DUMB	
10:16:04:48	DUMB	CB1	OPEN				DUMB	iv
10:16:04:52	DUMB			MOSS NO2	INST EF PROT OPTD	OFF	DUMB	
10:16:04:52	DRUM			DUMB	DIRECTIONAL EF OPTD	OFF	DRUM	iv
10:16:04:53	DUMB			SS02	STBY E/F PROT OPTD	OFF	DUMB	
10:16:04:53	DUMB			SS01	STBY E/F PROT OPTD	OFF	DUMB	iv

The labelling at the side of the SCADA listing above indicates the following:

- i. main operations
- ii. remote backup operations
- iii. local backup operations
- iv. relay resetting

The diagnostic reasoning according to the three stages of developed hypothesising strategies shall be described for this “fault plus failure” disturbance. No disturbances existed before the above alarms.

6.1.3.1 First Hypothesiser Diagnosis

The first hypothesiser uses a fixed failure limit and always attempts to create new scenarios. Therefore, when the alarm indicating the operation of the instantaneous earth fault relay is received, the expert system begins generating original hypotheses and scenarios with up to 1 simultaneous failure. Firstly, no-failure hypotheses are generated. The protection modelling knowledge identifies the faults which the instantaneous earth fault relay could have detected. The winding arrangement of the transformer at the Mosside end of the feeder is found to “block” neutral effects with the result that the relay can only detect a ground fault on the feeder or on the transformer primary winding. Two hypotheses are simulated and two scenarios are therefore produced. A sub-hypothesis is then derived by hypothesising the failure of a device whose activity is predicted in a scenario. Eight scenarios are produced based on sub-hypotheses. A false operation scenario is also generated.

Since they describe predicted secondary activity, the next three alarms strengthen those scenarios which predicted them and retract those based on the respective devices’ failures. When the directional earth fault relay alarm is received, those scenarios which predicted the activity are strengthened. No

scenarios included the device as a failure because of the limited definition of the network. New scenarios are attempted based only on this one alarm. In total, twenty six partial scenarios are successfully created, and these are merged with the remaining (un-strengthened) old scenarios to form complete scenarios. The remaining primary and secondary alarms are dealt with in the same way as that described for the earlier alarms of the corresponding types. When all the alarms have been received, the ranking table of scenarios looks like the following:

* Messages to 13 feb 1989 10:16:04.53 read. Disturbance -19999

1. There was a ground fault_on dumb to moss::fd2 which_is (permanent or temporary), which the protection equipment cleared.
However cb5 at dumb failed.
2. There was a ground fault_on dumb to moss::fd1 which_is (permanent or temporary), which the protection equipment cleared.
However rief1 at dumb failed.
There was also a ground fault_on dumb to moss::fd2 which_is (permanent or temporary), which the protection equipment is attempting to clear.
(The alarm for cb5 at dumb tripped is missing.)
3. There was a ground fault_on moss::tx2\\hv_wdg which_is (permanent or temporary), which the protection equipment cleared.
However cb5 at dumb failed.
80 possibilities

6.1.3.2 Second Hypothesiser Diagnosis

The second hypothesiser strategy also repetitively attempts to create new scenarios. The difference in this strategy is that disturbance failure limits are initially zero, and are incremented when an unexpected, but possibly related, alarm is received. Therefore, in the described case, the false alarm scenario and the two no-failure scenarios (faults on feeder and transformer HV winding) are the only scenarios to be generated and maintained up until the receipt of the alarm from substation Drumgelloch.

The “gut reaction” then identifies the directional earth fault relay alarm as unexpected but potentially related. The “central diagnosis” increments the simultaneous failure limit. It then proceeds to generate and simulate hypotheses (including sub-hypotheses) based on the directional earth fault relay operation using

the new failure limit. Hypotheses of a ground fault on 'dumb to moss::fd2' and on 'moss::tx2\\hv_wdg' are among those generated. Of these, partial scenarios are only generated with 'cb5 at dumb' as a failure—the no-failure hypotheses' simulations did not include the directional earth fault relay operation. These two successful partial scenarios were copied as complete scenarios since no disturbance alarms were unaccounted for. Every other complete scenario is formed by merging the other partial scenarios with the old scenarios. Eighty scenarios exist at this stage. The remaining alarms strengthen the two single-fault-breaker-failure scenarios and some of the other scenarios; the remaining scenarios are updated to include the primary alarms as false. The final ranking table of scenarios is shown:

- * Messages to 13 feb 1989 10:16:04.53 read. Disturbance -19999
1. There was a ground fault_on dumb to moss::fd2 which_is (permanent or temporary), which the protection equipment cleared.
However cb5 at dumb failed.
 2. There was a ground fault_on dumb to moss::fd1 which_is (permanent or temporary), which the protection equipment cleared.
However rief1 at dumb failed.
There was also a ground fault_on dumb to moss::fd2 which_is (permanent or temporary), which the protection equipment is attempting to clear.
(The alarm for cb5 at dumb tripped is missing.)
 3. There was a ground fault_on moss::tx2\\hv_wdg which_is (permanent or temporary), which the protection equipment cleared.
However cb5 at dumb failed to trip.
80 possibilities

6.1.3.3 Third Hypothesiser Diagnosis

In contrast with the latter strategy, the third version of the hypothesiser strategy does not simply generate original hypotheses after incrementing the number of simultaneous failures ie. following unexpected but potentially related activity. Instead, the existing old scenarios are first adapted. Therefore the diagnosis of the described case is equivalent to that by the second hypothesiser version up until the (unexpected but potentially related) operation of the directional earth fault relay.

When the directional earth fault relay alarm is received, the existing scenarios include a ground fault on the feeder, a ground fault on the transformer HV winding, and a false alarm scenario. Only breaker 'cb5 at dumb' is a possible failure since it is the only device with predicted activity (disregarding resetting and reclosing alarms). The two existing single fault hypotheses therefore have the breaker failure added to form new hypotheses which are then simulated. Successful partial scenarios are formed from the simulation results. The partial scenarios are each copied as complete scenarios since every disturbance alarm is already accounted for. A false alarm partial scenario is then generated, and combined with the old false alarm scenario to form an updated false alarm scenario and with the two single-fault scenarios to form two "fault-plus-false-alarm" scenarios. The remaining alarms strengthen the two single-fault-single-failure scenarios, while the two remaining primary alarms are also included as false activity in the other scenarios. After all the listed alarms have been received, the resulting ranking table appears as:

- * Messages to 13 feb 1989 10:16:04.53 read. Disturbance -19999
1. There was a ground fault_on dumb to moss::fd2 which_is (permanent or temporary), which the protection equipment cleared.
However cb5 at dumb failed to trip.
 2. There was a ground fault_on moss::tx2\\hv_wdg which_is (permanent or temporary), which the protection equipment cleared.
However cb5 at dumb failed to trip.
 3. There was a ground fault_on dumb to moss::fd2 which_is (permanent or temporary), which the protection equipment is attempting to clear.
rsbef1 at dumb reset, rsbef2 at dumb reset, rdef1 at drum reset,
rsbef1 at dumb operated, rsbef2 at dumb operated & rdef1 at drum operated, all falsely.
5 possibilities

6.1.3.4 Brief Comparison

This case study demonstrates the pros and cons of the different hypothesiser strategies. Firstly, it should be noted that all three had two common conclusions in the top three ranking positions. The first strategy, however, involved much more unreasonable effort since many scenarios were created only to be removed a few alarms later. This gives evidence of the deliberate design philosophy adopted: the aim in developing the hypothesising expert system was to firstly ensure the correct solution was made, and then to devise the means by which unreasonable effort could be curtailed.

The second hypothesiser demonstrates the improvement achieved in that fewer scenarios are created during the disturbance. The third hypothesiser gives the correct result with much less effort without sacrificing generality. When displayed in detailed form, the top scenario also conveys the explanation details contained in the scenario:

1. There was a ground fault_on dumb to moss::fd2 which_is (permanent or temporary), which the protection equipment cleared - t5 at dumb operated, rief2 at dumb operated, cb2 at moss tripped, tirrt2 at moss operated, cb1 at drum tripped, t1 at drum operated, rdef1 at drum operated, cb1 at dumb tripped, t1 at dumb operated, rsbef1 at dumb operated, cb2 at dumb tripped, t2 at dumb operated, rsbef2 at dumb operated, rief2 at dumb reset, rdef1 at drum reset, rsbef1 at dumb reset & rsbef2 at dumb reset. However cb5 at dumb failed.
In fact because rief2 at dumb will not isolate the fault due to the failure of cb5 at dumb to trip, rdef1 at drum operated, rsbef2 at dumb operated, & rsbef1 at dumb operated to isolate the fault.

Although the case studied would not occur frequently, it is not an unrealistic incident, and demonstrates the type of help which would be made available to an operator.

6.1.4 Case 3

This case study relates to an incident on the portion of transmission network shown in Figure 6.3. Two main protections are provided on the Petersburn

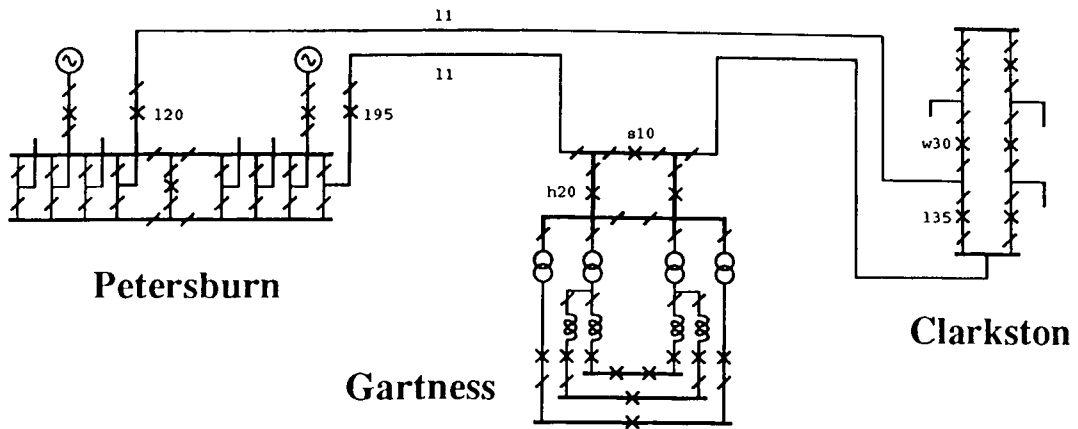


Figure 6.3: A portion of a transmission network

('pete') to Clarkston ('clar') interconnector by two distance protection schemes with acceleration. A pilot wire scheme and distance scheme with acceleration provide the protection on the interconnector between Petersburn and Gartness ('gart'). Delayed autoreclosing is provided on both feeders. In these schemes one attempt is made at reclosing at the Clarkston and Gartness ends, while the Petersburn end is delayed and is conditional on the breaker terminals being synchronised. Suppose that these two lines out of Petersburn experience a simultaneous lightning strike on the section on which the lines share towers and that the fault is cleared. The two sets of duplicate main protections would operate, those relays at Clarkston operating by acceleration. As a result, the trip relays at the three substations would operate and breakers would open. The breakers at Clarkston and Gartness reclose, and stay closed since the fault has been cleared. The Petersburn end breakers subsequently reclose. The alarms which would be expected to be generated by the SCADA system are as shown:

11:03:55:02	PETE2		GART2	FIRST MAIN PROT OPTD	ON	PETE2	i
11:03:55:02	PETE2		CLAR2	FIRST MAIN PROT OPTD	ON	PETE2	
11:03:55:02	GART2		PETE2	FIRST MAIN PROT OPTD	ON	GART2	
11:03:55:03	PETE2		GART2	SECOND MAIN PROT OPTD	ON	PETE2	
11:03:55:03	PETE2		CLAR2	SECOND MAIN PROT OPTD	ON	PETE2	
11:03:55:04	GART2		PETE2	SECOND MAIN PROT OPTD	ON	GART2	
11:03:55:04	CLAR2		PETE2	FIRST MAIN PROT OPTD	ON	CLAR2	
11:03:55:04	CLAR2		PETE2	SECOND MAIN PROT OPTD	ON	CLAR2	
11:03:55:04	PETE2		GART2	TRIP RELAYS TO BE RESET	ON	PETE2	
11:03:55:04	PETE2		CLAR2	TRIP RELAYS TO BE RESET	ON	PETE2	
11:03:55:04	GART2		PETE2	TRIP RELAYS TO BE RESET	ON	GART2	
11:03:55:05	CLAR2		PETE2	TRIP RELAYS TO BE RESET	ON	CLAR2	
11:03:55:05	PETE2		GART2	DAR IN PROGRESS	ON	PETE2	
11:03:55:05	PETE2		PETE2	DAR IN PROGRESS	ON	GART2	
11:03:55:05	PETE2		CLAR2	DAR IN PROGRESS	ON	PETE2	
11:03:55:05	CLAR2		PETE2	DAR IN PROGRESS	ON	CLAR2	
11:03:55:06	PETE2	L20	OPEN			PETE2	
11:03:55:06	PETE2	L95	OPEN			PETE2	
11:03:55:07	GART2	H20	OPEN			GART2	
11:03:55:07	GART2	S10	OPEN			GART2	
11:03:55:07	CLAR2	W30	OPEN			CLAR2	
11:03:55:07	CLAR2	L35	OPEN			CLAR2	
11:03:55:12	PETE2		GART2	FIRST MAIN PROT OPTD	OFF	PETE2	
11:03:55:12	PETE2		CLAR2	FIRST MAIN PROT OPTD	OFF	PETE2	
11:03:55:13	GART2		PETE2	FIRST MAIN PROT OPTD	OFF	GART2	
11:03:55:13	CLAR2		PETE2	FIRST MAIN PROT OPTD	OFF	CLAR2	
11:03:55:14	PETE2		GART2	SECOND MAIN PROT OPTD	OFF	PETE2	
11:03:55:14	PETE2		CLAR2	SECOND MAIN PROT OPTD	OFF	PETE2	
11:03:55:14	GART2		PETE2	SECOND MAIN PROT OPTD	OFF	GART2	
11:03:55:14	CLAR2		PETE2	SECOND MAIN PROT OPTD	OFF	CLAR2	
11:03:55:14	PETE2		GART2	TRIP RELAYS TO BE RESET	OFF	PETE2	
11:03:55:14	PETE2		CLAR2	TRIP RELAYS TO BE RESET	OFF	PETE2	
11:03:55:15	GART2		PETE2	TRIP RELAYS TO BE RESET	OFF	GART2	
11:03:55:15	CLAR2		PETE2	TRIP RELAYS TO BE RESET	OFF	CLAR2	
11:03:55:20	GART2	H20	CLOSED			GART2	
11:03:55:21	GART2	S10	CLOSED			GART2	
11:03:55:21	CLAR2	W30	CLOSED			CLAR2	
11:03:55:21	CLAR2	L35	CLOSED			CLAR2	
11:03:55:32	PETE2	L20	CLOSED			PETE2	
11:03:55:32	PETE2	L95	CLOSED			PETE2	
11:03:55:37	PETE2		GART2	DAR IN PROGRESS	OFF	PETE2	
11:03:55:37	GART2		PETE2	DAR IN PROGRESS	OFF	GART2	
11:03:55:38	PETE2		CLAR2	DAR IN PROGRESS	OFF	PETE2	
11:03:55:38	CLAR2		PETE2	DAR IN PROGRESS	OFF	CLAR2	

The labels along the right-hand-side of the above SCADA listing indicate the following:

- i. initial relay operation
- ii. switchgear first tripping
- iii. protection and trip relay resetting
- iv. switchgear reclosure attempt
- v. autoreclose relays resetting (no re-tripping)

Assuming there were no disturbances in existence before the above alarms are received, the (final) hypothesising expert system's diagnosis of this case shall be presented.

6.1.4.1 Hypothesiser Diagnosis

A disturbance area is created by the “gut reaction” in response to the first alarm. Both ‘permanent’ and ‘temporary’ fault hypotheses are generated since the operated protection belongs to an autoreclose scheme. Use of the protection modelling knowledge and the qualitative simulation result in the generation of scenarios involving a ‘permanent’ and a ‘temporary’ ground fault, a ‘permanent’ and a ‘temporary’ phase fault, and a false alarm. The operation of the second main protection is found to relate to the same disturbance, and strengthens the four fault-based scenarios since they predicted the operation. A false alarm partial scenario is combined with the existing false alarm scenario.

The third alarm is also related to the same disturbance area. However, the alarmed operation is not predicted by any of the existing scenarios and neither can the relay detect any of the hypothesised faults. Therefore the expert system presumes the operation to have been caused by a different event and proceeds to generate and simulate hypotheses with the same simultaneous failure limit (ie. 0 failures). The resultant five partial scenarios, four fault-based and one based on a false operation, are merged with the five old scenarios to produce a total of twenty five scenarios. The subsequent protection operated alarms cause the expert system to change the scenarios in a similar way to the changes which followed the receipt of the first “second main” alarm from Petersburn. All the scenarios predict the trip relay, acceleration and breaker alarms and are therefore strengthened by the reclassification of alarms as supportive—no failures have been hypothesised. The reset alarms also strengthen the scenarios, as do the breaker closed alarms from Gartness and Clarkston. However, the two breaker closed alarms from Petersburn were not predicted by the ‘permanent’ fault scenarios, with the result that the ‘temporary’ fault scenarios appear at the top of the ultimate ranking table as shown:

- * Messages to 28 jan 1989 11:03:55.38 read. Disturbance -1999
1. There was a ground fault_on clar2 to pete2::l1 which_is temporary, which the protection equipment cleared.
There was also a ground fault_on gart2 to pete2::l1 which_is temporary, which the protection equipment cleared.
 2. There was a ground fault_on clar2 to pete2::l1 which_is temporary, which the protection equipment cleared.
There was also a phase fault_on gart2 to pete2::l1 which_is temporary, which the protection equipment cleared.
 3. There was a phase fault_on clar2 to pete2::l1 which_is temporary, which the protection equipment cleared.
There was also a ground fault_on gart2 to pete2::l1 which_is temporary, which the protection equipment cleared.
- 25 possibilities

The preliminary “deeper diagnosis” knowledge makes some additional observations. One of these is that the faults in the top scenarios are probably caused by a single incident since the lines share the same towers. It is expressed in the following way:

- * Messages to 28 jan 1989 11:03:55.38 read. Disturbance -1999
(Based on scenarios: 1-4)
Possibly, the fault_on clar2 to pete2::l1 which_is temporary
and the fault_on gart2 to pete2::l1 which_is temporary
were caused by one fault on the shared tower length near pete2.

6.1.4.2 Observations

This study demonstrates the ability of the expert system to cope with multiple faults. If two faults were more remote such that they were contained in different disturbance areas, then the diagnosis of each would simply be done independently of the other and two ranking tables would be maintained. The case described is of greater significance. The type of knowledge which could be available from a “deeper diagnosis” is also illustrated. Deeper concepts but not deeper diagnostic knowledge is made use of in this particular case.

6.2 Evaluation

6.2.1 Some Factors

There are various factors by which the developed expert system can be evaluated. The range of measures is a result of the complexity of the diagnosis task and the diversity of requirements. In this section, the performance of the expert system shall be viewed with respect to analysis timing, reliability, robustness, indecisiveness and comprehensiveness. Moreover, “analysis timing” can relate to two aspects of a fault diagnosis system: the time taken to perform the diagnosis and the point in time (ie. the stage in the disturbance) when the diagnosis is started.

6.2.1.1 Analysis Timing–Completion Time

This factor relates to the time which the expert system takes to complete its diagnosis of a disturbance. It is important for both simple and complex disturbances for different reasons.

Following simple disturbances, the expert system should complete a correct diagnosis within a reasonable time period in order to evoke confidence in its ability. An aid which produces its answer long after the engineer has reached the same conclusion does not evoke much confidence, except perhaps to strengthen the engineer’s confidence in himself. Even this is unnecessary for simple disturbances—the presentation of diagnostic conclusions would simply create a nuisance and would waste the operator’s time. The majority of incidents on power systems are straightforward: a fault occurs on an item of plant which is isolated together with a small part of the network by the correct operation of protection and switchgear equipment, and the SCADA system provides the correct alarms. Such simple disturbances would likely form the backbone of any feasibility-testing field trial;

the response of the expert system to these incidents would be the means by which operations and protection engineering staff would judge the system. The acceptance of expert systems by end-users has often proved a delicate task due to inherent suspicion and scepticism. If a prolonged delay precedes the presentation of the expert system's diagnosis of a straightforward disturbance, the merits of the system in diagnosing more complex cases might not be considered.

This is the major strength of the “wait-and-see” expert system.

Strategy	Completion Time (CPU seconds)
wait-and-see	51.72
hypothesiser 1	64.6
hypothesiser 2	39.58
hypothesiser 3	26.7

Table 6.1: Comparison of analysis times for Case Study 1

For example, refer to the analysis times for Case Study 1 (Section 6.1.2) when performed on a Digital VAX8650, which are shown in Table 6.1. While the “wait-and-see” diagnosis in Case Study 1 took 51.72 CPU seconds (39.22 CPU seconds of which was used in searching for generating sources or infeeds), the first hypothesising expert system took 64.6 CPU seconds. This exemplifies a general observation; the wait-and-see expert system completed its diagnosis of simple disturbances in less time than the hypothesising expert system. This is in part due to the hypothesiser commencing its main diagnosis effort earlier in the alarm stream: more time is spent updating scenarios and recalculating their uncertainty measure. Also, when non-unit relays are involved at the start of the disturbance, the hypothesiser will often consider possible fault locations which are ultimately outside the isolated area and therefore not considered by the “wait-and-see” diagnosis. Also, by considering up to one or two failures immediately on the receipt of the disturbance's first alarm, the initial hypothesiser delays reaching a conclusion to simple disturbances. The later hypothesiser strategies suppress the consideration of more complex hypotheses until there are indications that

the disturbance is not simple, that is, involves failures. This is again evident in the results shown in Table 6.1. Clearly, for this particular case study, the “wait-and-see” strategy does not have a completion time advantage over the latest hypothesiser. This is partially due to the combination of the complexity of the network and the number of searches made for generating sources or infeeds by the “wait-and-see” strategy.

The time taken by the expert system to diagnose more complex disturbances is also important. The expert system is presented as a fault diagnosis aid to the operator or protection engineer which copes with the complexities of real system performance such as device maloperations and telemetry malfunctions. Following a disturbance, a window of opportunity for completion of the diagnosis exists in the control centre before the dead-line for a standard response or implementation of a restoration procedure. Similarly, off-line use by a protection engineer will demand a reasonable, though less stringent, response time. A favourable cost-benefit analysis must be in evidence: the cost includes the implementation cost and processing time cost; the benefit is the value of the diagnostic conclusion at the time when it becomes available.

Case Study 2 demonstrates again an advantage of the developments to the hypothesising strategy. For the sake of comparison, the “wait-and-see” expert system concluded with the top two scenarios of the third hypothesiser (but without the explanation details). On a Digital VAX8650 the CPU times for complete diagnosis were as shown in Table 6.2. In terms of the analysis timing (completion

Strategy	Completion Time (CPU seconds)
wait-and-see	12.6
hypothesiser 1	367
hypothesiser 2	321
hypothesiser 3	28.3

Table 6.2: Comparison of analysis times for Case Study 2

time) factor, the “wait-and-see” strategy is clearly better than the hypothesising

strategy. Moreover, the improvements in the analysis timing resulting from the modifications to the hypothesiser strategy are obvious. Some comments should however be added to the above statistics.

- Figure 6.2 shows the complete defined network. With this being so small and simple, searching for a generating source or infeed does not use significant processing time. As a result, the “wait-and-see” analysis for this case is, in contrast with Case Study 1, faster than that of the third hypothesiser.
- Although the same fault and failure diagnosis was concluded by the “wait-and-see” strategy, the conclusion lacked the added value of the explanation of the backup activity.
- The “wait-and-see” expert system provided no benefit until a blackout area was seen to be created by the receipt of one of the disturbance’s later alarms. The hypothesiser identified the system fault after the first alarm. However SCADA alarms are generated so rapidly that such information would need to be withheld for further confirmation of the conclusions before being made available to the control centre engineers.

This latter characteristic stems from the second aspect of analysis timing, the commencement time.

6.2.1.2 Analysis Timing–Commencement Time

This aspect of analysis timing is concerned with the stage in a disturbance’s progress at which the main diagnosis is started. Whereas the “wait-and-see” expert system performs its main diagnosis near the completion of the disturbance, the main diagnosis of the hypothesiser is commenced following the first received alarm. By waiting until a system fault appears to be isolated, the “wait-and-see” expert system has more evidence on which to base its diagnosis. Consequently,

the diagnosis is reasonably well directed to the correct possibilities. There is generally less protection activity after an area of the network has been isolated. Therefore the “wait-and-see” diagnosis performs the majority of its reasoning when there are fewer SCADA alarms still to be received.

The diagnosis of the hypothesiser approach is started earlier, but is consequently primarily based on less evidence. Some of the scenarios developed may therefore be later proved to be unreasonable by subsequent alarms. The cost of this unprofitable reasoning is dependent on the nature of the device which has operated—non-unit relays with large reaches, like overcurrent types, can detect many faults for which many scenarios may be created. Significant modifications to scenarios may be required since scenarios are generated before many alarms are received. Following a single fault, the first hypothesiser mostly removes or strengthens scenarios (its predicted alarms changed to supportive). However, the alterations to the hypothesiser strategy have shifted some of the reasoning load to later in the disturbance’s progress: failures are included in hypotheses only when there is evidence to support them. This has resulted in fewer fruitless scenarios being generated.

Therefore, regarding analysis timing, the “wait-and-see” strategy is better than the hypothesising strategy. Also, in terms of analysis timing, the later hypothesiser strategies exhibit an improvement over the earlier one. However the other evaluation factors must also be borne in mind and will be discussed in the following sections.

6.2.1.3 Reliability

Reliability measures the guarantee which can be given that the scenario which occupies the top ranking position does indeed describe the actual disturbance.

Reliability is crucial to an inexperienced operator who may rely more heavily

on the exact diagnostic conclusions. Not only so, but under severe conditions a skilled operator may have very little time to take in information. To reliably order alternative diagnostic conclusions is therefore a valuable attribute. It is equally important that the expert system is capable of providing alternative conclusions since these alternatives are merely less likely and may be the correct description of a disturbance. The robustness factor is concerned with the expert system at least including the correct diagnosis among the possible alternatives produced. Reliability is therefore largely dependent on the means of dealing with uncertainty.

The “wait-and-see” and hypothesiser expert systems both use the relative likelihood measure to deal with uncertainty. The case studies presented in the previous section demonstrate the reliability of this measure. Applying relative likelihoods to different aspects of system behaviour permits relative judgements to be made between different faults and different maloperations and malfunctions. The different ranking positions following disturbances will sometimes be due to scenarios expecting different numbers of alarms. Such an obvious difference would be recognised and suitably used by most means of dealing with uncertainty. The comparison of alternative diagnostic conclusions is often less obvious. It may be impossible to distinguish between alternatives from the SCADA alarms alone. Scenarios 1 and 2 produced by the third hypothesiser for Case Study 2 are of this type. A transformer fault would normally be more serious, while a line fault could well have been a self restoring fault. This is clearly of significance to the subsequent restoration procedures adopted. Complex disturbances can involve a large area of the network and hence result in a large number of possible diagnostic conclusions. Discrimination between alternatives on the basis of different numbers of missing alarms and the frequency of occurrence of certain events, which is provided by the relative likelihood, will improve the expert system’s reliability and thus enable better restorative action to be taken.

Although the hypothesiser and “wait-and-see” expert systems use the same means of dealing with uncertainty, their overall reliabilities are not equivalent. This is because both of them do not always include the correct diagnosis as a possibility—they do not have the same robustness.

6.2.1.4 Robustness

Unlike consultative expert systems, the fault diagnosis system does not have control over the supply of symptoms to it. The availability of protection and switchgear activity alarms is dependent on the performance of the SCADA system. The fault diagnosis system must be robust enough for its diagnostic ability to remain unharmed under adverse conditions. That is, the correct explanation of a disturbance must be included among the alternative diagnostic conclusions even when alarms of equipment activity are lost or corrupted. This quality is here given the term robustness. Where robustness is lacking, the diagnosis will be characterised by “brittleness” (this term is used by Kirschen et al [58]). It has been observed by utility engineers that the occurrence of missing alarms is much more common than protection or switchgear equipment failures. Therefore robustness is an important factor.

The second circumstance in Case Study 1 demonstrates how the main failing of the “wait-and-see” strategy is its level of robustness. The expert system failed to generate any diagnostic conclusions simply because one breaker trip alarm was not received at the control centre. This dependence on a complete knowledge of switchgear activity causes this strategy to be lacking in robustness. Moreover, the consideration of only those faults which can be detected by all protection relays related to a disturbance area is not robust: if a disturbance causes a relay to falsely operate then the genuine fault might not be identified.

In contrast, the hypothesising strategy does not rely on receiving every alarm.

Its robustness stems from the prediction of protection and switchgear activity. The activity will generally not be predicted before it occurs, but rather before the activity alarm has been read in by the expert system. For a simulated hypothesis to be used as the basis for a scenario, it need only include the received alarm on which the hypothesis was based; a perfect match with all the received alarms is not required. The discovery of lost alarms in the disturbance in Case Study 1 demonstrates the robustness of the hypothesising expert system. Section 6.1.2.4 shows how the diagnosis of that disturbance is achieved despite the loss of two alarms: the correct diagnosis is still made and the missing alarms are identified.

The adaptations made to the hypothesiser strategy in creating the second and third hypothesising strategies have in fact deteriorated the robustness. While the initial hypothesiser would include a single-fault-single-failure scenario despite all backup activity alarms being lost or corrupted, the second hypothesiser would not since the inclusion of failures is suppressed until there is evidence to support it. Further, the third hypothesiser could fail to include a double-fault-missing-alarm scenario if the protection which detected the second fault (the only relay to do so) appeared to be operating as backup for the first system fault. While these developments are retrogressive, it should be stated that device failures are relatively unlikely, as are simultaneous system faults.

In order to cope with the loss of a protection relay operated alarm, knowledge was added to the hypothesiser to respond to a trip relay being received first: scenarios are generated based on the operation of each relay which could have caused the telemetered activity, with its alarm classed as missing. This therefore presents a problem, not of lacking robustness, but rather of indecisiveness.

6.2.1.5 Indecisiveness

Whereas robustness relates to the expert system's ability to include an explanation of the genuine disturbance among its alternative diagnostic conclusions, indecisiveness indicates the problem of producing too many alternatives. Indecisiveness is therefore the opposite of "brittleness". The indecisiveness of a diagnostic strategy is closely related to its reliability. A lack of reliability in the ranking system necessitates the provision of more than the top ranking alternative in order that the correct diagnosis is among those displayed. Otherwise, the correct diagnosis would be missing from the alternatives displayed in an unacceptable number of cases. At the same time, it is important that the expert system does not replace a voluminous list of alarms with an equally incomprehensible list of alternative conclusions. Moreover, indecisiveness is more than a problem of how many alternatives to display. Indecisiveness caused by preserving too many alternatives for too long, even if their display is suppressed, results in wasted analysis time.

The "wait-and-see" strategy combats indecisiveness by only considering fault locations inside blackout areas, and only identifying faults which account for the operation of all related protection relays. For example, the expert system's diagnosis of the disturbance described in Case Study 2 included, and only ever included, two alternative conclusions.

Indecisiveness is more of a problem for the hypothesiser, but more so with the earlier hypothesising strategies. The response of the three strategies to the disturbance in Case Study 2 illustrates this. Indecisiveness in the final suppression of alternative conclusions shall not be considered since the three strategies have the same potential for restricting the displayed conclusions. Table 6.3 shows the number of scenarios in the ranking tables of the three hypothesising expert systems at different stages of the disturbance: after the first alarm, after the disturbance

Hypothesiser	Number of Scenarios		
	first alarm	last alarm	minimum
1	11	80	5
2	3	80	3
3	3	5	3

Table 6.3: Total scenarios at various stages of the diagnosis of Case Study 2

and the minimum ever held in the ranking table. Comparing the initial hypothesiser's results with the two scenarios only ever considered by the "wait-and-see" expert system illustrates the unfavourable indecisiveness which this strategy exhibited. The second hypothesiser exhibits a decreased indecisiveness because the inclusion of failures is suppressed until there appears to be evidence of backup activity. The indecisiveness in the third hypothesising expert system is now at a reasonable level.

Lost alarms can adversely affect the indecisiveness exhibited following a certain disturbance, as they do the robustness. For example, a breaker failure with a missing trip relay alarm will produce the same alarms as would failure of the trip relay—this can result in the hypothesiser almost doubling the number of possible scenarios. Data skewing is also a problem in some SCADA systems. In such, alarms from one substation may be received at the control centre before other alarms generated earlier at another substation. If this causes backup activity alarms to be received before any main protection alarms, the hypothesiser might generate many more scenarios than it would without data skewing since backup protection can generally detect faults in a larger part of the network. Increased indecisiveness would therefore be exhibited.

6.2.1.6 Comprehensiveness

Comprehensiveness indicates the extent of the knowledge which is available from the expert system. Rules have been included in both the "wait-and-see" and hypothesising expert systems which form literal messages for display from the

frame-like scenarios facts. The conversion process is not as involved as it might otherwise have been due to the nature of the knowledge in the expert systems. These messages describe system faults and device maloperations along with the alarm details. More than one of each can be included.

The hypothesising expert system, however, exhibits greater comprehensiveness. It has the intrinsic ability to provide explanatory details of failures and associated backup activity. This is in evidence in the detailed output shown at the end of Case Study 2. By providing such additional information, the system lends itself to other varied uses such as operator training systems. Moreover, the knowledge is intrinsic in the diagnostic strategy and requires little extra reasoning to retrieve it. The “wait-and-see” expert system would require dedicated rules to provide such additional explanatory knowledge. (The provision of suitable explanatory knowledge is a large and complex problem. An indication of this is given in Chandrasekaran et al [81].) The deeper diagnosis will introduce many more concepts which may be included in diagnostic conclusions. This will greatly add to the system’s comprehensiveness.

6.2.2 Knowledge Nature

In evaluating an expert system, the nature of the knowledge must be considered. The soundness of the developed expert systems, especially that of the hypothesiser, is evident from the earlier discussions. However, a more objective standpoint needs to be taken in order to further judge the validity of the research in terms of the suitability of the developed techniques. Therefore, the nature of the knowledge representation, its closeness to corresponding human knowledge, its depth and the level of understanding it contains shall be considered.

6.2.2.1 Nature of Knowledge Representation

Details of the power system to which the expert system is applied are contained in frame-like Prolog facts. This is a very suitable format which allows both details and relationships to be recorded in a natural form. In many respects this database is like the relational databases which are being extensively adopted by utilities. Indeed, a frame-based database management system was implemented in Prolog and is described by Sheblé [44] as a relational database that includes a knowledge based system.

The logic rule paradigm has proved to be a suitable means of representing the knowledge of a fault diagnosis system. Primarily, these advantages relate to the development stage, rather than implementation. The evolutionary development cycle facilitated by Prolog enabled knowledge to be represented for small-scale initial prototypes. The diagnostic knowledge was proved with small test cases, and built upon as the diagnostic abilities were expanded. In addition, Prolog provided the flexibility which permitted the representation of some knowledge as production rules. Although rule bases of only production rules were traditionally used (comment made on p12 in CCTA [82]), their suitability has limitations. Some degree of structuring (eg. sectionalising) of the rules is required in order to restrict the rules which are invoked at any time and thereby meet the performance criteria required for on-line systems. Prolog logical rules have provided structuring which is suited to the fault diagnosis problem without excessively detracting from the understandability.

There are some disadvantages of representing the knowledge in Prolog. Firstly there is the issue of implementation in an Energy Management System environment. The discussion in section 3.3.5.1 demonstrated that this is becoming less significant with the extended use of communication networks and workstations. Such an environment is suited to expert systems implemented in Prolog. A sec-

ond disadvantage of Prolog is its relatively slow speed of operation. The efficiency issue has been and is being addressed by developers of Prolog systems, and the parallelisation of reasoning using concurrent Prolog could lead to much faster reasoning.

Nevertheless, it should be stated that, in the application-oriented research reported in this thesis, Prolog has been viewed as a tool. The aforementioned attributes of Prolog matched certain needs in the fault diagnosis system, and the expert system's development took advantage of this suitability. However, the developed knowledge is not critically dependent on Prolog; the diagnostic strategy and the principles of the inference in Prolog could be retained in a translation to another language. One large consulting company is known to develop expert systems in Prolog to meet the task's reasoning demands, and then manually translate the code into FORTRAN for delivery in order to meet the performance criteria. In such a translation, performance is gained at the expense of losing some understandability and maintainability.

6.2.2.2 Closeness to Human Knowledge

Another question which may be asked of the knowledge nature is:

“How close is the knowledge in the fault diagnosis systems to that of a human expert?”

There are different types of knowledge which an engineer might use in his diagnosis: relational, particular and essential.

An engineer may apply his experience of previous disturbances to a present incident. Specific symptoms would be related to specific diagnoses in what amounts to a look-up table. Neither the “wait-and-see” nor the hypothesising expert systems contain such knowledge. This omission of relational knowledge is justified on two counts:

- power systems are so large and diverse that any reasonably comprehensive diagnosis system would require a massive look-up table of unmanageable proportions
- the fallibility of SCADA systems means that equivalent disturbances might not generate equivalent symptoms thus adding to the possible relationships

Knowledge of particular relevance to one item of plant or one type of plant may be used by an engineer in disturbances involving the said plant. An expert system for the fault diagnosis of power systems would require several particular knowledge bases in order to adequately cover the complete system. Moreover, knowledge of a suitably extensive nature would be required to efficiently select the particular knowledge which should be invoked. The developed expert systems do not contain particular plant knowledge, but perform diagnoses to the point of locating faults on certain parts of plant items. The “deeper diagnosis” part of the diagnosis structure will later include the plant-particular knowledge.

The expert system has instead been developed with “essential” knowledge. That is, the fundamental knowledge required of an engineer to diagnose faults has been represented. This essence of fault diagnosis knowledge makes available the expertise required to diagnose disturbances on any power system. The “wait-and-see” strategy involves a connectivity-based identification of isolated areas which an engineer might use. The location of faults within the blackout area is performed using a connectivity-based protection representation. Such a reasoned view of protection behaviour is also used in the hypothesising strategy. The knowledge realising both strategies is therefore reasonably close to “essential” or fundamental human knowledge.

6.2.2.3 Depth of Knowledge

Consideration of the depth of knowledge in the developed expert system is a corollary of the previous section. A shallow diagnosis could be adopted for the sake of speed of operation and could simply directly relate symptoms to conclusions. Conversely, a deep diagnosis would give a more detailed conclusion longer after the receipt of the symptomatic data. It could be said that a middle course has been taken in the research described in this thesis. Modelling knowledge, though not deep plant knowledge, has been included. The “wait-and-see” strategy involves the shallower knowledge of the two developed philosophies. The modelling knowledge contained in the “wait-and-see” expert system covers the detection logic of protection relays and the internal plant connectivity, while that of the hypothesiser includes protection detection logic, system behaviour, internal plant connectivity, and protection and switchgear activity. Only following the inclusion of the detailed modelling of transformers, generators, etc. will the expert system strongly merit the adjective “deep”. In assessing the ultimate depth to which the diagnosis might go, the value of the knowledge at the time of its derivation must be borne in mind.

6.2.2.4 Level of Understanding in Knowledge

The level of understanding is closely coupled with the type of knowledge which is incorporated—whether or not significant depths of modelling knowledge have been represented in the system. A significant level of understanding of the system behaviour is inherent in the hypothesiser strategy. This provides for good explanation facilities. The potential has been used to provide the explanation of backup activity, an example of which was given in Case Study 2.

The nature of the knowledge derived by both strategies should be commented upon. The output of the developed expert systems contains a significant amount

of “understandable knowledge”. For example, the top conclusion in Case Study 1 includes:

```
'ground fault_on coat2 to clar2::12 which_is (permanent or temporary)'
which consists of Prolog variables and operators. In addition,
```

```
'rm2_clar at coat2 operated'
```

has a similar make-up. The output of the expert system is not just knowledge to the engineer to whom it is presented; it could be used for further reasoning by the same, or another, expert system. In contrast, an indivisible textual block, for example,

```
'rm2_clar at coat2 operated'
```

(ie. the string of ASCII codes representing this) falsely appears to represent detailed understanding. On this basis, a significant level of understanding is exhibited. The “representation trap” described by Lenat [83], by which systems can appear to have significant knowledge and yet possess little understanding, has been avoided.

6.2.3 Comparison with Other Systems

Section 3.4 reviewed other expert systems which have been developed for the fault diagnosis of power systems. All but one of these systems are similarly intended for location in control centres. (The expert system for implementation in substation computers has not been developed sufficiently for it to be compared with control centre implementation.) The telemetered protection and switchgear activity is the most common dynamic knowledge source among the other expert systems; it is used in the research reported in this thesis also. Some of the fault diagnosis systems rely on switchgear alarms since they were developed for SCADA systems which did not alarm relays. Continued developments in hardware, eg. in Remote Telemetry Units (RTUs), will cause relay information to be increasingly available

from SCADA systems. Digital disturbance recorder data has instead been used in SEPT [69]. Such data is of unquestionable benefit to fault diagnosis. However, for it to be used for the fault diagnosis of a power system, reasoning is required to diagnose the disturbance to such a degree that the need for disturbance recorder data can be assessed and the appropriate recorders identified. This has only been implemented to a limited extent in SEPT, which reflects its limited application to one substation.

The hypothesising strategy described in this thesis compares favourably with those strategies developed by other authors. The other strategies are generally more restrictive. Some can only deal with one system fault, and device failures have been ignored or restricted to one per conclusion. The expert system described in this thesis has been developed with a view to complex and simultaneous disturbances—real world problems which complicate the fault diagnosis task, but which strengthen the case for implementation of a fault diagnosis system. Therefore, the diagnosis of multiple faults with multiple device failures has been included as described in this thesis.

The use of a generate-and-test philosophy is not unique to the hypothesising expert system described in this thesis. D2, FAUST, DSF, GESTAL, AHFA, SEPT, and the system developed by Amantegui et al (see section 3.4) all use a generate and test philosophy. However, each has certain limitations:

- The modelling used in D2 appears excessively detailed for the limited knowledge derived. As a result it is extremely inefficient.
- FAUST uses very shallow system modelling: fault modelling is only at a connectivity level.
- Hypotheses are generated by DSF following the identification of a black-out and therefore the strategy does not give reasonable robustness against missing telemetry.

- AHFA also uses shallow modelling in its generate and test strategy and consequently does not identify device failures.
- It would appear that SEPT lacks selectivity in the hypotheses it generates since “votes” are determined for a fault on every high voltage item of plant.
- The strategy developed by Amantegui et al is similar to that described in this thesis. However, although it generates and immediately tests hypotheses, it performs no subsequent testing. It must therefore await the completion of a disturbance before the diagnosis is commenced.

It is claimed that the strategy described in this thesis strikes the correct balance in the depth of system modelling. The “deeper diagnosis” knowledge will contain more accurate models, but these will only be used after the central diagnosis has curtailed the number of hypotheses to be considered. Moreover, protracted disturbances can be confidently dealt with because of the inclusion of rules for subsequent testing in addition to those for testing hypotheses immediately following their generation. That is, the expert system does not “mark time” for an unacceptable length of time while waiting for the cessation of protection and switchgear activity.

The nature of the knowledge described in this thesis differs from that adopted in many of the other expert systems. For example, the depth of knowledge ranges from a relatively shallow connectivity-based diagnosis described by Fukui and Kawakami [50], to the detailed time-related simulation in D2; in the former a quick shallow conclusion is given, while, at the other end of the depth scale, the latter needs 3 Microvaxes to produce two hypotheses per minute. This thesis describes the development of an expert system which concludes diagnoses of reasonable content in a time scale which is reasonable for on-line use in a control centre. In addition, the use of “essential” knowledge has resulted in a very portable system. The frame-like facts whose use is described in this thesis are also expected to

improve portability since databases describing the system should be able to be rapidly built up. It may even be possible for a significant proportion of the database to be virtually automatically obtained from marked-up Computer Aided Design (CAD) drawing files.

It is evident from the publication dates that the work in this thesis was not simply a series of corrections to these other expert systems. The work of Fukui and Kawakami [50] did provide a launch pad for this research. However many of the aspects in the developed strategy were a direct result of limitations observed in the prototypes developed in the course of this research. The fact that the development of several expert systems has resulted in the adoption of a generate-and-test philosophy corroborates the claim that the developed strategy is suitable.

Many means of realisation have been adopted. Prolog has been used in the development of at least six of the expert systems cited in section 3.4. Of the remainder, OPS has been used in at least four. In addition, C, Lisp, ART, KEE, Poplog, Smalltalk and a shell called M1 are also reported to have been used. Prolog systems lack some of the development aids provided by the other languages and tools. This does not invalidate the suitability of Prolog, the grounds of which are given in this thesis.

Of those expert systems described in section 3.4, only six (FAUST, GESTAL, SEPT, CRAFT, and those by Hsu et al and Amantegui et al) have had periods of testing in a control centre. This is an obvious advantage over the expert system described in this thesis, with which there has been no installation experience. Field trials give a good appreciation of an expert system outside a laboratory and provide a means of gaining valuable operator appraisal. However, such periods of installed use sometimes do little to test the system more rigorously than in a laboratory since disturbances, especially complex ones, are relatively rare events. For example, the field trial of FAUST did not involve any fault events [36], while

CRAFT had seen one fault before the developers wrote their paper [65].

The timing of analysis in the developed expert system is preferable to that published by other authors. Provided that a buffering program can be written, the developed hypothesiser has help available from the early stages in a disturbance's progress. Many of these other expert systems await more evidence before starting their diagnosis. However this is risky since triggering conditions, such as the identification of a blackout area or the existence of a quiet period, might not be satisfied or may be a long time in being fulfilled. The completion time of the developed expert system appears to be reasonable in comparison with other systems.

The reliability of the developed expert system is due to the relative likelihood based ranking. It is better than having no ranking and also uses more knowledge than merely counting the numbers of failures and/or false alarms—the two most common approaches among other authors. Conversely, a mathematically sound technique is not suitable since sufficiently sound data is not available; moreover the mathematically based techniques adopted in expert systems in general are ill matched to this domain.

Robustness is the major factor which sets the developed expert system apart from many of the others. The expert system can correctly diagnose missing switchgear or protection alarms. It will diagnose failed relays and switchgear, and explain relationships with backup activity. Simultaneous disturbances do not pose a problem to this expert system. In addition, the expert system considers alternative explanations and not just the first solution found—a few of the expert systems cited in section 3.4 take the latter approach.

Whereas the earlier versions of the hypothesiser described in this thesis exhibited an unreasonable degree of indecisiveness, this has to a large extent been overcome. Moreover, the high resolution provided by the relative likelihood mea-

sure gives the confidence to suppress the display of alternative conclusions in the lower end of a ranking table. This makes the output of the expert system appear more decisive.

The comprehensiveness of the expert system is also favourable in comparison with other reported expert systems. The robustness of the expert system, which permits it not only to cope with but also to identify maloperations and multiple faults, likewise illustrates how comprehensive the diagnostic ability is. It must be noted, however, that GESTAL and AHFA also make use of automatic recovery system activity. Knowledge of the post-fault activity of auto-reclosing equipment has been included in the expert system described in this thesis, but no inference from post-fault isolator activity is yet made.

6.3 Anticipated Use

The research on which this thesis is based was not intended to be a purely theoretical exercise. The objective is to have this expert system operational in a utility's control centre, providing useful diagnostic knowledge to system operators and also to protection engineers. It is envisaged that by implementing this expert system, a utility's management of their network would be significantly improved. Realistic use of this expert system in the short and long terms will be briefly discussed.

6.3.1 Short Term

6.3.1.1 Non-intrusive Implementation

At present, most working energy management system (EMS) computers are not best suited for implementing expert systems. Furthermore, utility engineers would not permit access to their SCADA computer. It would be unwise to venture to directly implement such a relatively new technology and consequently risk interfering with such a vital function as the operation of the power system. A non-intrusive implementation would be advantageous to the security of the EMS functions and also for reasons of access to the expert system for maintenance. The expert system could run on a workstation with its own internal representation of the network, or it could be given strictly read-only access to the topology database. A connection into the link between the SCADA computer and the alarm logging printer would provide access to the SCADA alarms. In order to avoid delaying normal functionality of the EMS, a buffering program on the workstation would hold messages until the expert system was ready for them.

6.3.1.2 Data Collection

Clearly, significant effort must be undertaken to represent the power system data in the required form. (Even if access to existing databases is given, the necessary data will often not be available since utilities in general have as yet not built up such databases, having not yet seen the benefits of these.) One approach to lessen this task is to limit the proportion of the network whose data is realised in the knowledge base. This, however, lessens the opportunity to help the operator in a genuine disturbance.

This process of data collection would be helped by partially automating it. Attributised Computer Aided Design (CAD) network drawing files could provide a significant level of initial data which would be manually added to. Utilities are using CAD packages with interfacing facilities for producing network diagrams, and are also reviewing their corporate data structures and requirements. Consideration might be given to symbolic data representations of the network which could integrate many corporate functions. Thereby changes made in the utility drawing office could be automatically downloaded to analysis databases, echoing the implementation of the modifications on the actual network. Implementation and maintenance of expert systems for power system applications in general could thus be achieved more rapidly.

6.3.1.3 Return for Implementation Effort

Once implemented on a workstation, the expert system would be an extra aid to the operator. Situated beside the operator console, the workstation monitor would automatically display the diagnostic conclusions. The diagnostic conclusions would simply be displayed and made available for perusal. The dedicated workstation's keyboard and mouse would only be used if necessary, for example, to manipulate the display in order to show more details of conclusions or

to show more alternative conclusions. In addition to being non-intrusive of the EMS hardware, the implementation would therefore not intrude on the operator's activities.

Moreover, the expert system could help the protection engineer by making historical electronic records of its diagnoses. The disturbance's diagnostic conclusions would supplement operator logs in giving valuable help to the protection engineer. The resultant documentation (operator log, on-line diagnostic conclusion, SCADA listing and investigation conclusion) would be a valuable means of proving the expert system. Corporate communications networks which are increasingly being installed could provide a means of delivering diagnostic conclusions to the protection function headquarters.

6.3.2 Long Term

6.3.2.1 Integration at Control Centres

Integration of the fault diagnosis expert system in the control centre Energy Management System (EMS) computing environment could be achieved. The point at which the SCADA activity is picked up by the expert system could be shifted to earlier in the data path than that for short term use. By accessing the SCADA data before its conversion to textual messages, the fault diagnosis process could be speeded up. However, this could only be done provided the data checking routines were not by-passed and hence require duplication in the expert system.

Although intrusive, the expert system must not be obstructive to other functions. For this reason, a distributed, open environment holds the best prospect for successful implementation: a dedicated or shared workstation could be connected to the control centre Local Area Network (LAN). Software such as X-windows provides the means of presenting the results in a window on one of the monitors

on the operator's console.

6.3.2.2 Integration in a Corporate Information System

It is expected that operator fault reports will become electronic rather than remaining paper based. Fault reporting software could be more than just a simple editor by being interfaced with the fault diagnosis expert system. The expert system could therefore form a report skeleton to which the operator could add comments.

Interfacing the fault diagnosis expert system with a Geographic Information System (GIS) would also improve the management of power systems. Diagnostic conclusions could be related to topological and geographical system representations of the network. Not only could this improve the conveyance of information to the operator, but also the geographical information would improve investigation procedures. Better directions could be conveyed to investigation teams. Management of field engineering resources could be improved through the provision of geographical knowledge of manpower location, fault location and substation location.

Manipulation of maintenance schedules in line with diagnoses of device performance would also improve efficiency of that aspect of utility operations. Not only would maloperations in diagnostic conclusions be used, validated activity could be used to delay routine maintenance in order to free resources for more necessary investigation.

One of the acknowledged shortcomings in utility operations is the accurate upkeep of statistical data. Comprehensive statistical data and lifetime plant records are of increasing interest to utilities. The fault diagnosis performance monitor could automatically update the statistical records (following operator confirmation if necessary).

In addition to the above uses of the diagnostic conclusions, the expert system diagnostic process could itself be used to automatically manage the transmission of digital disturbance recorder data. If the fault diagnosis expert system was interfaced with the replay software presently available (eg. that produced by CSD [41]), data transfer would be intelligently selective and also rapidly initiated.

6.3.2.3 Diversification

The knowledge in the fault diagnosis expert system has more diverse potential uses in addition to its immediate role in control centres.

The core knowledge could be used in an operator training system. Both instructor and trainee could benefit from the availability of fault diagnosis expertise. Due to the scarcity of complex disturbances, artificial disturbances are often created and realistically presented to the trainee on a control console simulator. Fault diagnosis knowledge could be part of a scenario-building aid by which an instructor could recreate realistic and varied disturbances to enhance the training portfolio. A simulator of the type described by Boutacoff [32] provides the structure for the straightforward implementation of such an aid. With such flexibility, a simulator could be used for the further training and development of existing operators. This will become more significant with the changes in control centre technology: experienced operators could also become familiar with the new technology by using a simulator.

In the longer term, the current research in the field of coordinated control and protection will result in the resources being available to disperse much of the fault diagnosis effort to main substations. A degree of parallelisation would be realised. Effort has been directed towards achieving a distributed open architecture for use in substations. The field bus researched by EDF ([84]) is an example of this. This would make the status of many more contacts available to the fault diagnosis:

information which is only presently available by remote data transmission from disturbance recorders, if at all. If suitable operational speeds could be achieved, the fault diagnosis expert system could even suppress the transmission of SCADA alarms and instead generate more densely detailed messages. An expert system at the control centre would use these “sub-diagnosis” conclusions to obtain complete explanations of disturbances. In this way, control centre staff would not lose sight of what was happening out on the system.

Chapter 7

Conclusions and Future Work

7.1 Conclusions

The work which was undertaken and described in this thesis has not concluded research in this field. However several conclusions can be drawn from the completed research. These are listed below:

- This research has proved the feasibility of a fault diagnosis expert system. The produced system can cope with real-world problems such as simultaneous disturbances and failed equipment. In this sense it is robust. At the same time, efforts have been made to curtail the indecisiveness. It has been demonstrated how many of these performance factors oppose each other. However, a suitable and reasonable compromise is exhibited in the final version of the hypothesising expert system.
- Although others were considered, two main diagnostic strategies have been realised and tested: the “wait-and-see” and the “hypothesiser”. The “wait-and-see” approach, which is based on the identification of blackout areas, has been shown to exhibit brittleness due to its need to receive every switch-gear trip alarm. Since the “hypothesiser” approach uses a qualitative simulator, it is less susceptible to SCADA system fallibility. The hypothesiser strategy has been shown to be preferable for this and other reasons.
- In contrast with many other expert systems, a significant level of network knowledge is included in the developed expert system—especially that using the hypothesising strategy. In this sense it may be described as being “model-based”. The developed expert system has a significant level of understanding due to this knowledge. Moreover, the knowledge is fundamental, essential or generic in nature, with the result that the expert system is very portable. In addition, the level of understanding has provided the means for a limited explanation of fault-and-failure diagnoses without the

addition of dedicated explanatory rules. Furthermore, the level of understanding included means that more knowledge can be easily added in order to progress the reasoning further eg. into restorative advice.

- The importance of adequately structuring the knowledge has been recognised. This was recognised during the research when the problems of the scale of genuine power systems were considered. A hierarchical structure has been used for the diagnostic knowledge, diagnosis conclusions and the specific network details. This structuring has made the diagnosis more focussed.
- The research has proved the worth of the logic rule based paradigm, and also of Prolog as a development tool. The knowledge is transparent since the language is symbolic and can be recognizably close to natural language. An evolutive development cycle can therefore be used, and this has been employed. The feasibility of Prolog as an implementation language has yet to be fully evaluated. However, the developments in control centre computing facilities have independently provided an environment suitable for the implementation of expert systems. What remains to be established is whether or not Prolog can provide the speed of performance required for a real time application.
- The management of uncertainty in a fault diagnosis expert system is problematic. Many systems count the unverified evidence: the uncertainty is the total number of maloperations, sometimes including the number of system faults. This technique lacks selectivity, however, especially when telemetry malfunctions are also included in the diagnosis. At the same time, the available mathematically based techniques are unsuitable due to a lack of necessary statistics and also because of inappropriate assumptions which

do not match the domain. “Relative likelihood” has been developed and is presented as a viable alternative.

- It has been found that the building of the databases necessary for a demonstration of a working expert system takes a significant part of the development time. The benefits to be had from maintaining a complete protection database, a symbolic connectivity database, plant history database, etc. are only truly realised when used with or integrated with an expert system such as that developed. The development of an expert system will therefore of necessity involve at least the partial creation of system databases—this is a time-consuming task.
- The prototyping method of development does hold benefits over a more traditional design-coding-debugging approach as was claimed by Wollenberg and Sakaguchi [49]. However it must be used with care. Although tested on small test systems, the development must not progress without bearing the ultimate implementation in mind. For example, data structures were influenced by the problems of the scale of genuine power systems. If all attention had been given to the prototype system, problems with unsuitable data structures may have appeared later in the development.
- The development and use of the loadflow routine as described in this thesis demonstrates the potential for algorithmic analysis routines to be used co-operatively with knowledge-based systems. Automatic routines were developed to build up the mathematical system details from the symbolic system database. The system data was passed to the loadflow routine via a file before any diagnosis began. The output of the loadflow was processed to derive the data of interest to the fault diagnosis. Three observations may be made regarding this aspect of the work:

- The data required by knowledge and algorithm based systems may relate to the same application system but be of different types—automatic translation routines can be developed to minimise duplicated effort and ease maintenance.
- File-based communications is suitable for once-off large volume data transfer, however inter-process communications (IPC) techniques such as UNIX sockets or shared memory would be better for on-line cooperation (such as overload checking of restoration plans).
- Seldom is the vast amounts of data generated by the likes of loadflow routines of interest: rather, certain features in that data are of importance. Where possible, placing the filtering routines on the algorithm side of the interface can dramatically reduce the volume of data to be transferred.
- Through the giving of demonstrations, it has been found that literal messages are particularly powerful. Although they are longer than corresponding messages which merely express the information they contain, the literal messages output by the fault diagnosis have been well received and the information quickly discerned by engineers. The main limitation is in expressing the location of a disturbance or maloperated device; a role for which graphics are particularly suited. It is believed that a graphical representation coupled with literal explanation would be the most user friendly and effective means of informing power system operators of diagnosis conclusions.

7.2 Future Work

The work described in this thesis has successfully produced a fault diagnosis expert system. There still remains some effort to be given to a detailed consideration of implementation issues. In addition, the knowledge which can be concluded by the present system could be used as a starting point for further valuable reasoning. Guidelines to these areas of fruitful future work follow:

- Work has to be performed to develop an implementation interface by which the expert system could access SCADA alarms on-line and yet not be a hindrance to existing SCADA computer activities. Marathe et al [65] describe the interfacing system they developed for a fault diagnosis expert system called CRAFT.
- Detailed modelling knowledge should be created for items of plant, such as transformers and generators. This knowledge will form part of the “deeper diagnosis” level of the fault diagnosis and thereby permit the generation of more detailed conclusions. With the addition of deeper modelling knowledge, the maloperation of the varied protection used on such plant can be more effectively recognised. Indeed, rather than incorporate this knowledge in the system-wide fault diagnosis expert system, this expertise could be made available through communicating with plant-specific fault diagnosis and condition monitoring expert systems. Researchers in CEPE [71] are developing an expert system for turbo-alternators, whose integration could provide the desired depth of knowledge for such items of plant.
- Establishing interfaces between the expert system and other software systems could be greatly beneficial. An interface with digital disturbance recorder replay software would provide more plant indications and would permit the deeper plant modelling knowledge to be used to greater effect.

The same interface would also enable the performance monitor to recognise timing irregularities which could precede outright failure. Integrating the fault diagnosis expert system with Geographic Information Systems (GIS) would provide the former with a valuable man-machine interface and the latter with the ability to provide much more detailed and beneficial information. By enhancing the existing interface with a loadflow and extending it to other power application software like fault flows, state estimators, etc., the “deeper diagnosis” could verify conclusions more confidently. The provision of these interfaces would allow the role of the expert system to move from a purely supervisory one into some simple executive activity such as automatically initiating disturbance recorder data transfer and running analyses.

- Having realised a working fault diagnosis system, a natural progression would be to add restoration knowledge. Several researchers have been looking into this area of applying expert systems [1, 13, 60, 65, 85, 86, 87, 88, 89, 90]. Initially the expert system would only have responsibility for providing restorative advice. However, with access to loadflows for careful validation, some restoration procedures could ultimately be employed automatically. The movement towards significant executive authority for the expert system holds great potential but requires careful investigation because of safety issues.
- An area requiring investigation is the parallelisation of the fault diagnosis system. Both local and dispersed parallelisation hold potential for improved performance. Local parallelisation would be achieved by the use of parallel processors in the control centre to share the symbolic processing load—both the 3-tier diagnosis and the consideration of individual hypotheses hold the potential for parallelisation. Dispersed parallelisation would be achieved by

implementing fault diagnosis expert systems in substations. Some diagnosis could be performed at a substation without any reference to external events eg. trip relay, enabled by protection relay, operates and causes tripping of breaker. Telemetry malfunctions would not be problematic. The expert systems could possibly take over an executive role of suppressing SCADA alarms provided performance speeds are reasonable: alarms would be filtered at source and a more compact message transmitted. As well as reducing the data volume displayed to the operator, this would reduce the data volume transmitted. Coordinated control and protection substation technology is providing the environment by which such processing could be performed at the substation level.

Appendix A

Introduction to Prolog

A.1 Convention Adopted

In the following description, the semantics of Quintus Prolog [77], which are compatible with the Edinburgh Prolog convention, shall be used. Comments follow a percentage symbol on a line, or are contained within `/*...*/`.

A.2 Brief Description of the Prolog Language

The basic data object in Prolog is a term. A term can be a constant, a logical variable or a structure. `relay1`, `1.5`, `'Substation1'` are constants. A logical variable is identified by having an upper case first letter, like `Relay`; don't care logical variables, those whose value is not used, are denoted by an underscore. A structure is a relationship (the predicate name), with one or more arguments inside parentheses as shown: `predicate_name(arg1, arg2, ..., argn)`. The predicate is defined by its predicate name and the number of arguments thus: `predicate_name/n`. The predicate name is a constant, and each argument is a term.

A Prolog program consists of clauses, each of which comprise a head and a body and is terminated with a period. The head of a clause is a structure or constant. A fact is a clause which has an empty body. Some facts are shown:

```
relay_class(distance, non_unit).
relay_type('Relay 1', differential_cc).
relay_type('Relay 2', distance).
```

A rule is a clause with a non-empty body, and is of the form

```
head:- body.
```

which reads "if body then head" or, more literally, "head if body". The body is a series of conditions separated by commas. For example, the rule

```
a_relays_class(Relay,Class):-
    relay_type(Relay,Type),
    relay_class(Type,Class).
```

expresses the knowledge that a relay has a certain class if the relay has a certain type and the type is of that class.

Clauses in the knowledge base are used when a query (or goal) is specified. A query can be a simple query, or a compound query. A compound query is a series of two or more simple queries separated by commas, and is satisfied when all the simple queries have been satisfied in turn. To satisfy a simple query, an attempt is made to match it with clause heads in the knowledge base. If no match is found then the simple query fails; if that simple query is part of a partially satisfied compound query, backtracking commences to find alternative solutions to the previously satisfied simple queries. (Backtracking is prevented by a cut; the ! (exclamation) operator is simply satisfied, but when backtracking reaches it, the reasoning does not proceed to the earlier terms in the compound query.) However, if a matching clause head is found, and that clause is a fact, then the simple query is satisfied; if part of a compound query, the next simple query is attempted. When the matching clause head belongs to a rule, then the conditions in the rule body must be satisfied for the head to be satisfied. The conditions in the rule body act like a series of simple queries.

For example, with the four clauses

```
relay_class(distance, non_unit).
relay_type('Relay 1', differential_cc).
relay_type('Relay 2', distance).
a_relays_class(Relay,Class):-
    relay_type(Relay,Type),
    relay_class(Type,Class).
```

suppose a simple query is given thus

```
| ?- relay_class(distance, Class).
```

This query would match with the first clause above, the variable `Class` being instantiated (or set) to `non_unit`, and the response would be given

```
Class = non_unit
```

Likewise, suppose the following simple query is expressed

```
| ?- a_relays_class(Ry, non_unit).
```

The query matches the rule, with the rule variable `Class` assuming the constant `non_unit` (ie. the variable is instantiated). The rule variable `Relay` matches query variable `Ry` and remains uninstantiated. The rule conditions now represent a compound query to be satisfied. The first condition in the rule body is attempted: `relay_type(Relay, Type)`. The first `relay_type/2` clause matches this and instantiates both arguments thus: `Relay = 'Relay 1'`, `Type = differential_cc`. These apply to the rest of the rule. Therefore the next condition whose satisfaction is attempted, is now `relay_class(differential_cc, non_unit)`. This condition fails to be satisfied, matching neither facts nor rules. Backtracking therefore commences.

The previous condition is reattempted, with the search for a matching clause commencing after the point in the knowledge base where it was last satisfied ie. after the second clause. As `Relay` and `Type` had been instantiated by this condition, they are returned to their uninstantiated state. The structure `relay_type('Relay 2', distance)` matches this condition and instantiates `Relay` and `Type` to new constants. The next condition in the rule's body is again attempted: `relay_class(distance, non_unit)` matches the first clause in the knowledge base. All the rule conditions are satisfied with `Relay` having assumed the constant `'Relay 2'`. Therefore the query is satisfied and the response would be given:

```
Ry = 'Relay 2'
```

Not all conditions are separated by commas; the operators ; (semicolon) and -> are also used between conditions. While the comma can be viewed as logical AND, the semicolon can be viewed as representing logical OR. For example, consider the compound query (which could equally well be a rule body)

```
| ?- a, (b;c), d.
```

To satisfy this, suppose a is tested and satisfied. Then b is attempted. If b is satisfied, then only d remains to be satisfied for the query to be satisfied. If however b fails, then c is attempted, and again d remains to be satisfied. Likewise, if a and b have been satisfied but d fails, then backtracking returns to b for an alternative solution. If that fails then c is attempted. Only if c fails will backtracking extend back to a.

The -> operator is used in the following format:

```
a -> b;c.
```

This reads as “if a then b otherwise c”. That is, if a is satisfied then b is attempted. If a fails then c is attempted. If b fails, then c is not attempted. In addition, a is never reattempted on backtracking.

Prolog also has the flexibility to use modified semantics. For example, describing a fault as being a type of fault on an item of plant, a fault could be defined by a predicate `fault_on/2` in the standard way as illustrated:

```
Fault = fault_on(ground, line1).
```

However, Prolog allows the definition of operators. By defining `fault_on` to be an infix operator, it is permissible to express the above as:

```
Fault = ground fault_on line1.
```

Thus, without translation, easily understood expressions can be reasoned with. Using this feature can allow complex structures like

```
Fault = which_is(fault_on(ground, ::(to(ss1,ss2), line1)), permanent).
```

to be expressed as

```
Fault = ground fault_on ss1 to ss2::line1 which_is permanent.
```

The knowledge base can be modified. Clauses can be added or deleted from the knowledge and thereby change the working of the Prolog program. Clauses are added to the knowledge base using the built-in predicate `assert/1`, and removed by `retract/1`. Therefore, the addition or deletion of facts from the knowledge base are often referred to as asserting or retracting.

References

- [1] T. Sakaguchi and K. Matsumoto. "Development of a knowledge based system for power system restoration". *IEEE Transactions on Power Apparatus and Systems*, Vol. PAS-102, No. 2, pp 320–329, February 1983.
- [2] T.E. Dy Liacco and T.J. Kraynak. "Processing by logic programming of circuit-breaker and protective-relaying information". *IEEE Transactions on Power Apparatus and Systems*, Vol. PAS-88, No. 2, pp 171–175, February 1969.
- [3] D. Swade. "It will not slice a pineapple". *IEE Review*, pp 217–220, June 1991.
- [4] T. Moore. "Artificial intelligence: Human expertise from machines". *EPRI Journal*, pp 6–15, June 1985.
- [5] Z.Z. Zhang, G.S.Hope, and O.P. Malik. "Expert systems in electric power systems—A bibliographical survey". *IEEE Transactions on Power Systems*, Vol. PWRS-4, No. 4, pp 1355–1362, October 1989.
- [6] P. Mason. "Embedded expert systems". Presented at an IEE meeting, February 1990.
- [7] I.J. Cox. "Expert systems". *IEE Electronics & Power*, pp 237–240, March 1984.
- [8] A.M. Turing. "Can a machine think?". In E.A. Feigenbaum and J. Feldman, editors, *Computers and Thought*. McGraw-Hill, 1963.
- [9] E.H. Shortliffe. *Computer Based Medical Consultations: MYCIN*. American Elsevier, 1976.
- [10] R. Duda, J. Gaschnig, and P. Hart. "Model design in the PROSPECTOR consultant system for mineral exploration". In D. Michie, editor, *Expert*

Systems in the Microelectronic Age, pp 153–167. Edinburgh University Press, 1979.

- [11] T. Flintham. “Making the rules”. *IEE Review*, pp 235–239, June 1991.
- [12] P.A. Subrahmanyam. “The ‘software engineering’ of expert systems: Is Prolog appropriate?”. *IEEE Transactions on Software Engineering*, Vol. SE-11, No. 11, pp 1391–1400, November 1985.
- [13] M.M. Adibi and R.J. Kafka. “Power system restoration issues”. *IEEE Computer Applications in Power*, pp 19–24, April 1991.
- [14] R.S. Engelmores. “Expert systems: A glimpse into the 1990s”. In J.A. Naser, editor, *Expert Systems Applications for the Electric Power Industry*, volume 1 of 2, pp 1–11. Hemisphere Publishing, 1990.
- [15] M. Yazdani. “Shells versus toolkits”. *Expert Systems User*, 5(1):13, January 1989.
- [16] D.G. Bobrow. “If Prolog is the answer, what is the question? or What it takes to support AI programming paradigms”. *IEEE Transactions on Software Engineering*, Vol. SE-11, No. 11, pp 1401–1408, November 1985.
- [17] T. Taylor and D. Lubbkeman. “Implementation of heuristic search strategies for distribution feeder reconfiguration”. *IEEE Transactions on Power Delivery*, Vol. PWRD-5, No. 1, pp 239–246, January 1990.
- [18] J. Schreurs and F. Bednar. “On-line valve monitoring and diagnosis”. *IEEE Computer Applications in Power*, pp 25–29, January 1990.
- [19] Y. Freundlich. “Knowledge bases and databases”. *IEEE Computer*, pp 51–57, November 1990.

- [20] L.A. Zadeh. "Knowledge representation in fuzzy logic". *IEEE Transactions on Knowledge and Data Engineering*, Vol. 1, No. 1, pp 89–100, March 1989.
- [21] K.-C. Ng and B. Abramson. "Uncertainty management in expert systems". *IEEE Expert*, pp 29–47, April 1990.
- [22] C. Marcus. *Prolog Programming: Applications for Database Systems, Expert Systems, and Natural Language Systems*. Addison-Wesley, 1986.
- [23] P.H. Winston. *Artificial Intelligence*. Addison-Wesley, 1977.
- [24] I. Bratko. *Prolog Programming for Artificial Intelligence*. Addison-Wesley, 1986.
- [25] A. Abu-Hanna and Y. Gold. "An integrated, deep-shallow expert system for multi-level diagnosis of dynamic systems". In J.S. Gero, editor, *Artificial Intelligence in Engineering: Diagnosis and Learning*. Computational Mechanics, 1988.
- [26] D.S. Prerau, A.S. Gunderson, R.E. Reinke, and M.R. Adler. "Maintainability techniques in developing large expert systems". *IEEE Expert*, pp 71–79, June 1990.
- [27] P.S. Sell. *Expert Systems: A Practical Introduction*. Macmillan, 1986.
- [28] L. Johnson and E.T. Keravnou. *Expert Systems Architectures*. Kogan Page, 1988.
- [29] C. Clark. "In the know". *CADCAM*, p 15, November 1990.
- [30] R.P. Schulte, S.L. Larsen, G.B. Sheble, J.N. Wrubel, and B.F. Wollenberg. "Artificial intelligence solutions to power system operating problems". *IEEE Transactions on Power Systems*, Vol. PWRS-2, No. 4, pp 920–926, November 1987.

- [31] R.D. Christie and S.N. Talukdar. "Expert system for on line security assessment. A preliminary design". *IEEE Transactions on Power Systems*, Vol. PWRS-3, No. 2, pp 654-659, May 1988.
- [32] D. Boutacoff. *EPRI Journal*, pp 27-35, June 1990.
- [33] B.F. Wollenberg. "Feasibility study for an energy management system intelligent alarm processor". *IEEE Transactions on Power Systems*, Vol. PWRS-1, No. 2, pp 241-247, May 1986.
- [34] J.-M. Arès and P. Girouard. "GESTAL: A specialised tool to build real-time alarm processing and fault diagnosis expert systems for power network control centres". In J.A. Naser, editor, *Expert Systems Applications for the Electric Power Industry*, volume 1 of 2, pp 679-689. Hemisphere Publishing, 1990.
- [35] P.H. Ashmole. "Power system alarm analysis and fault diagnosis using expert systems". pp 144-150.
- [36] D. Cockburn, J.R. McDonald, G. Burt, J.R. Brailsford, J. Beaton, and K.L. Lo. "Expert systems for on-line fault diagnosis in electrical power networks". In *11th International Conference on Electricity Distribution (CIRED)*. Liège, April 1991.
- [37] J.R. McDonald, G.M. Burt, and D.J. Young. "Alarm processing and fault diagnosis using knowledge based systems for transmission and distribution network control". *IEEE Transactions on Power Systems*, Vol. PWRS-7, No. 3, pp 1292-1298, August 1992. Also in *IEEE PES 1991 Transmission and Distribution Conference*, Dallas, September 1991.
- [38] H.T. Roman and F.A. Marian. "The year 2000 power plant". *IEEE Computer Applications in Power*, pp 19-22, April 1989.

- [39] P. Bornard. *EDF Bulletin de la Direction des Etudes et Recherches*, B(4):33–43, 1988.
- [40] H.L. Smith and T.J. Modzelewski. “Enhancing energy management systems with advanced RTU capabilities”. *IEEE Computer Applications in Power*, pp 26–29, October 1989.
- [41] Circuits & Systems Design (CSD). *Replay System Manual*.
- [42] P.V.F. Beardow, J.A. Barber, and J.E. Tame. “Implementation of a low cost rural telecontrol and monitoring system”. In *3rd International Conference on Power System Monitoring and Control*, pp 22–27. London, June 1991.
- [43] D. Rumpel, S.S. Venkata, K. Pandji, C.C. Liu, and N.M. Nagdy. “Real-time database for power systems using language oriented data structure”. *IEEE Transactions on Power Systems*, Vol. PWRS-5, No. 3, pp 993–1000, August 1990.
- [44] G.B. Sheblé. “Distributed distribution energy management system (DDEMS) exploratory experiments”. *IEEE Computer Applications in Power*, pp 37–42, January 1991.
- [45] A Working Group of the Relaying Practices Subcommittee of the IEEE Power System Relaying Committee. “Summary of the special publication ‘Applications of fault and disturbance recording devices for protective system analysis’”. *IEEE Transactions on Power Delivery*, Vol. PWRD-4, No. 3, pp 1625–1630, July 1989.
- [46] W.R. Prince, B.F. Wollenberg, and D.B. Bertagnolli. “Survey on excessive alarms”. *IEEE Transactions on Power Systems*, Vol. PWRS-4, No. 3, pp 950–956, August 1989.
- [47] B.M. Weedy. *Electric Power Systems*. Wiley, 1984.

- [48] J.A. Naser. In J.A. Naser, editor, *Expert Systems Applications for the Electric Power Industry*, volume 1 of 2, pp xiii–xiv. Hemisphere Publishing Corporation, 1990.
- [49] B.F. Wollenberg and T. Sakaguchi. “Artificial intelligence in power system operations”. In *Proceedings of the IEEE*, volume 75 No. 12, pp 1678–1685. December 1987.
- [50] C. Fukui and J. Kawakami. “An expert system for fault section estimation using information from protective relays and circuit breakers”. *IEEE Transactions on Power Delivery*, Vol. PWRD-1, No. 4, pp 83–90, October 1986.
- [51] D.C. Berry and D.E. Broadbent. “Expert systems and the man/machine interface. Part2: The user interface”. *Expert Systems*, 4(1):18–28, February 1987.
- [52] E. Cardozo and S.N. Talukdar. “A distributed expert system for fault diagnosis”. *IEEE Transactions on Power Systems*, Vol. PWRS-3, No. 2, pp 641–646, May 1988.
- [53] S.N. Talukdar, E. Cardozo, and L. Leao. “TOAST: The power system operator’s assistant”. *IEEE Computer*, 19(7), July 1986.
- [54] S.N. Talukdar, E. Cardozo, and T. Perry. “The operator’s assistant—An intelligent, expandable program for power system trouble analysis”. *IEEE Transactions on Power Systems*, Vol. PWRS-1, No. 3, pp 182–187, August 1986.
- [55] M.A. Bramer, J.S. Pearce, D. Muirden, J.C. Platts, and D.L. Vipond. “FAUST—An expert system for diagnosing faults in an electricity supply

system". In B. Kelly and A. Rector, editors, *Research and Development in Expert Systems 5*. Cambridge University Press, 1988.

- [56] A.A. Girgis and M.B. Johns. "A hybrid expert system for faulted section identification, fault type classification and selection of fault location algorithms". *IEEE Transactions on Power Delivery*, Vol. PWRD-4, No. 2, pp 978–985, April 1989.
- [57] K. Watanabe and N. Oobe. "On-line data analysis system with remote data transmission". In *Colloquium of CIGRE Study Committee 34 - CIGRE24 Preferential Subject on Development and Applications of Expert Systems to Protection and Local Control Systems*. Brazil, September 1991.
- [58] D.S. Kirschen, B.F. Wollenberg, G.D. Irisarri, J.J. Bann, and B.N. Miller. "Controlling power systems during emergencies: The role of expert systems". *IEEE Computer Applications in Power*, pp 41–45, April 1989.
- [59] B. Jeyasurya, S.S. Venkata, S.V. Vadari, and J. Postforoosh. "Fault diagnosis using substation computer". *IEEE Transactions on Power Delivery*, Vol. PWRD-5, No. 2, pp 1195–1201, April 1990.
- [60] K. Hotta, H. Nomura, H. Takemoto, K. Suzuki, S. Nakamura, and S. Fukui. "Implementation of a real-time expert system for a restoration guide in a dispatching centre". *IEEE Transactions on Power Systems*, Vol. PWRS-5, No. 3, pp 1032–1038, August 1990.
- [61] R. Balakrishnan and A. Pahwa. "A computer assisted intelligent storm outage evaluator for power distribution systems". *IEEE Transactions on Power Delivery*, Vol. PWRD-5, No. 3, pp 1591–1597, July 1990.
- [62] H. Okamoto, A. Yokoyama, and Y. Sekine. "A real-time expert system for fault section estimation using cause-effect network". In *Proceedings of the*

10th Power Systems Computation Conference, pp 905–912. Graz, Austria, August 1990.

- [63] J. Partanen, T. Koppanen, and P. Jarventausta. “An expert system for fault diagnosis of medium voltage distribution networks”. In *Proceedings of the 10th Power Systems Computation Conference*, pp 1080–1087. Graz, Austria, August 1990.
- [64] Y.-Y. Hsu, F.C. Lu, Y. Chien, J.P. Liu, J.T. Lin, H.S. Yu, and R.T. Kuo. “An expert system for locating distribution system faults”. *IEEE Transactions on Power Delivery*, Vol. PWRD-6, No. 1, pp 366–372, January 1991.
- [65] H.Y. Marathe, C.-C. Liu, M.-S. Tsai, R.G. Rogers, and J.M. Maurer. “An on-line operational expert system with data validation capabilities”. *IEEE Transactions on Power Systems*, Vol. PWRS-6, No. 2, pp 882–889, May 1991.
- [66] K. Tomsovic, C.-C. Liu, P. Ackerman, and S. Pope. “An expert system as a dispatchers’ aid for the isolation of line section faults”. *IEEE Transactions on Power Delivery*, Vol. PWRD-2, No. 3, pp 736–743, July 1987.
- [67] D.G. Esp and E.T. Cahill. “A real time expert system for fault diagnosis on a transmission network”. In *3rd International Conference on Power System Monitoring and Control*, pp 237–240. London, June 1991.
- [68] J. Claus, F. de Viron Dongier, M. Baudrux, P. Lienart, P. Van Miegroet, and J.C. Maun. “Processing the indications of operation and alarm in the electrical networks”. In *Colloquium of CIGRE Study Committee 34 - CIGRE24 Preferential Subject on Development and Applications of Expert Systems to Protection and Local Control Systems*. Brazil, September 1991.

- [69] P. Fauquembergue and J.V. Esquevin. "An expert system as an on-line aiding tool for the maintenance of protection systems". In *Colloquium of CIGRE Study Committee 34 - CIGRE24 Preferential Subject on Development and Applications of Expert Systems to Protection and Local Control Systems*. Brazil, September 1991.
- [70] J. Amantegui, J. Corera, J. Echávarri, J. Barandiarán, I. Laresgoiti, and J. Pérez. "Expert system for fault analysis in electrical networks. Operating experience and further developments". In *Colloquium of CIGRE Study Committee 34 - CIGRE24 Preferential Subject on Development and Applications of Expert Systems to Protection and Local Control Systems*. Brazil, September 1991.
- [71] B.D. Gemmell, J.R. McDonald, R.W. Stewart, and B.J. Weir. "An on-line vibration monitoring system: An embedded expert system for condition monitoring and fault diagnosis on turbo-alternators". In *3rd International Congress on Condition Monitoring and Diagnostic Engineering Management (COMADEM)*, pp 52-59. Southampton, July 1991.
- [72] J.J. Keronen. "An expert system prototype for event diagnosis and real-time operation planning in power system control". *IEEE Transactions on Power Systems*, Vol. PWRS-4, No. 2, pp 544-550, May 1989.
- [73] J.J. Keronen. "EKA: An expert system for real-time operation planning and event analysis in electric power networks". In J.A. Naser, editor, *Expert Systems Applications for the Electric Power Industry*, volume 1 of 2, pp 447-463. Hemisphere Publishing, 1990.
- [74] S.D. Kaminaris, A.V. Machias, and B.C. Papadias. "An intelligent tool for distribution substations troubleshooting and maintenance scheduling".

IEEE Transactions on Power Delivery, Vol. PWRD-6, No. 3, pp 1038–1044, July 1991.

- [75] C.A. Protopapas, K.P. Psaltiras, and A.V. Machias. “An expert system for substation fault diagnosis and alarm processing”. *IEEE Transactions on Power Delivery*, Vol. PWRD-6, No. 2, pp 648–655, April 1991.
- [76] D. Curtice. “Review of expert systems in power system operations electrical systems division”. In J.A. Naser, editor, *Expert Systems Applications for the Electric Power Industry*, volume 1 of 2, pp 55–62. Hemisphere Publishing Corporation, 1990.
- [77] Quintus Computer Systems. *Quintus Prolog User’s Guide and Reference Manual*.
- [78] L.C.W. Frerk. “Busbar protection”. In The Electricity Council, editor, *Power System Protection*, volume 3, pp 81–148. Peter Peregrinus Ltd, 1986.
- [79] J. Arrillaga, C.P. Arnold, and B.J. Harker. *Computer Modelling of Electrical Power Systems*. Wiley, 1983.
- [80] R. Billinton and R.N. Allan. “Power-system reliability in perspective”. *IEE Electronics & Power*, pp 231–236, March 1984.
- [81] M.C. Tanner B. Chandrasekaran and J.R. Josephson. “Explaining control strategies in problem solving”. *IEEE Expert*, pp 9–24, Spring 1989.
- [82] CCTA. *Appraisal and Evaluation Library: Knowledge Based Systems Volume*. HMSO, July 1991.
- [83] D.B. Lenat. “Ontological versus knowledge engineering”. *IEEE Transactions on Knowledge and Data Engineering*, Vol. 1, No. 1, pp 84–88, March 1989.

- [84] O. Grob and A.-M. Germain. "Numerical control in HV/MV substations—EDF considerations". In *3rd International Conference on Power System Monitoring and Control*, pp 7–9. London, June 1991.
- [85] Y. Kojima, S. Warashina, M. Kato, and H. Watanabe. "The development of power system restoration method for a bulk power system by applying knowledge engineering techniques". *IEEE Transactions on Power Systems*, Vol. PWRS-4, No. 3, pp 1228–1235, August 1989.
- [86] C.-C. Liu, S.J. Lee, and S.S. Venkata. "An expert system operational aid for restoration and loss reduction of distribution systems". *IEEE Transactions on Power Systems*, Vol. PWRS-3, No. 2, pp 619–626, May 1988.
- [87] D.S. Kirschen and T.L. Volkmann. "Guiding a power system restoration with an expert system". *IEEE Transactions on Power Systems*, Vol. PWRS-6, No. 2, pp 558–566, May 1991.
- [88] S.-M. Wang, Z.-Z. Dong, Q.-H. Sun, and D.-Z. Xia. "An decision-support expert system for bulk power system restoration". In *Proceedings of the 10th Power Systems Computation Conference*, pp 966–971. Graz, Austria, August 1990.
- [89] B. Valiquette, G.L. Torres, and D. Mukhedkar. "An expert system based diagnosis and advisor tool for teaching power system operation emergency control strategies". *IEEE Transactions on Power Systems*, Vol. PWRS-6, No. 3, pp 1315–1322, August 1991.
- [90] S. Shah and S.M. Shahidehpour. "A heuristic approach to load shedding scheme". *IEEE Transactions on Power Systems*, Vol. PWRS-4, No. 4, pp 1421–1429, October 1989.