

EXAMINING THE USE OF EXPERT SYSTEMS
IN
MANAGING FERTILIZER APPLICATION DECISIONS

By

COLIN HO KEUNG LI

A Thesis
Submitted to the Faculty of Graduate Studies
in Partial Fulfillment of the Requirements
for the Degree of

MASTER OF SCIENCE

Department of Computer Science
University of Manitoba
Winnipeg, Manitoba

© April, 1991



National Library
of Canada

Bibliothèque nationale
du Canada

Canadian Theses Service Service des thèses canadiennes

Ottawa, Canada
K1A 0N4

The author has granted an irrevocable non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of his/her thesis by any means and in any form or format, making this thesis available to interested persons.

The author retains ownership of the copyright in his/her thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without his/her permission.

L'auteur a accordé une licence irrévocable et non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de sa thèse de quelque manière et sous quelque forme que ce soit pour mettre des exemplaires de cette thèse à la disposition des personnes intéressées.

L'auteur conserve la propriété du droit d'auteur qui protège sa thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

ISBN 0-315-76917-3

Canada

*EXAMINING THE USE OF EXPERT SYSTEMS IN
MANAGING FERTILIZER APPLICATION DECISIONS*

BY

COLIN HO KEUNG LI

A thesis submitted to the Faculty of Graduate Studies of
the University of Manitoba in partial fulfillment of the requirements
of the degree of

MASTER OF SCIENCE

© 1991

Permission has been granted to the LIBRARY OF THE UNIVERSITY OF MANITOBA to lend or sell copies of this thesis, to the NATIONAL LIBRARY OF CANADA to microfilm this thesis and to lend or sell copies of the film, and UNIVERSITY MICROFILMS to publish an abstract of this thesis.

The author reserves other publication rights, and neither the thesis nor extensive extracts from it may be printed or otherwise reproduced without the author's written permission.

Table of Contents

Acknowledgements	iii
Abstract	iv
Chapter 1: Introduction	1
1.1 What are Expert Systems?.....	3
1.1.1 Human Experts	3
1.1.2 Expert Systems	4
1.1.3 Conventional Computer Systems versus Expert Systems	6
1.2 Fertilization Decisions.....	9
1.2.1 Overview of the Application Domain.....	9
1.2.2 Current System for Making Fertilizer Selection Decisions	10
1.2.3 Making a Fertilization Plan.....	11
1.2.4 The Fertilizer Problem	13
Chapter 2: The Nitrogen-Only Problem	17
2.1 Nitrogen Recommendations	18
2.1.1 Yield Curve Type Recommendations	18
2.1.1.1 Yield Curves.....	19
2.1.1.2 Cost Curves.....	20
2.1.1.3 Net Return on Fertilizer Investment.....	21
2.1.1.4 Complete and Partial Cost Curves	22
2.1.1.5 Multiple Cost Curves	24
2.1.2 Grid Type Recommendations	26
2.2 Implementation of the Nitrogen-Only Module.....	27
2.2.1 Rule Base Processing	28
2.2.2 Pre-Processing.....	30
2.2.3 Main Processing.....	31
2.2.4 Post-Processing.....	34
Chapter 3: Multiple Nutrients	36
3.1 General Multiple Nutrient Fertilizer Problem.....	37
3.2 Overview of General Blackboard Systems	39
3.3 The Fertilizer Advisor Blackboard.....	41
3.3.1 Rule Base Processing Specialist	44
3.3.2 Single Nutrient Recommendation Specialists	46
3.3.2.1 Nitrogen-Only Specialists.....	47
3.3.2.2 Phosphorus-Only Specialist	51
3.3.2.3 Potassium-Only Specialist and Sulphur-Only Specialist	52
3.3.3 NPKS Solution Method	54
3.3.4 Two-Nutrients Specialists and PKS Specialist.....	56
3.3.5 Yield Curve Type NPKS Specialist.....	60
3.3.6 Grid-Type NPKS Specialist.....	66
3.3.7 Other Specialists in the FA System.....	67
3.3.7.1 Explanation Facilities	67
3.3.7.2 Comparison Facility	68
3.3.7.3 Trace Facility	69
3.3.7.4 Multiple Crops.....	72
3.3.7.5 Database Routines.....	72
3.3.7.6 Imperial/Metric Conversions.....	73
3.4 The Fertilizer Advisor Blackboard vs. General Blackboards.....	73

Table of Contents

Chapter 4: Evolution of Fertilizer Advisor	76
4.1 Incremental Development.....	76
4.2 MacIntosh Prototype.....	77
4.3 IBM Prototype.....	78
4.3.1 Nitrogen-Only Version.....	78
4.3.1.1 Hardware and software requirements.....	78
4.3.1.2 User-interface, rule base component and explanations.....	78
4.3.1.3 Objectives for the Next Phase of Development.....	82
4.3.2 Nitrogen-and-Phosphorus Version.....	83
4.3.2.1 Knowledge Acquisition and Knowledge Engineering.....	83
4.3.2.2 System Enhancements.....	84
4.3.2.3 Objectives for the Next Phase of Development.....	86
4.3.3 NPKS version.....	87
4.3.3.1 Controlling the Complexity of the Problem.....	87
4.3.3.2 More System Enhancements.....	89
 Chapter 5: Conclusion	 91
 References	 98

Acknowledgements

I sincerely thank Dr. Mark Evans for his guidance throughout my years in the Master's program at the University of Manitoba.

I am also grateful to other members of my committee, Dr. David Scuse and Dr. Don Flaten for their valuable suggestions toward my thesis.

Thanks to the experts, Dr. Don Flaten, Mr. John Ewanek, and Mr. Ken McGill who have generously contributed their knowledge and experience to the Fertilizer Advisor. Without them, this expert system would not exist.

I would like to express my appreciation to the members of the Solomon Sinclair Farm Management Institute, Dr. Rae Josephson, and Mrs. Bonnie Luterbach, for their continuing efforts in managing this project.

Abstract

Making fertilizer application decisions is a complex task. A large number of inter-related factors must be considered in the decision making process. Ideally, a farmer should have an agronomist analyze the field conditions and recommend a cost effective fertilization plan which outlines the fertilizer product(s) to be used, and the rate, method and timing of application. Choosing the most cost effective fertilization plan involves constructing many alternative plans and selecting the economic optimum plan. Such a strategy will result in an enormous number of plans being constructed and searched before the best plan can be found. We call this the *fertilizer problem*.

An agronomist does not use any formal algorithm for the selection of the economic optimum plan. Instead, agronomists rely on facts known about the problem, their previous experience, intuition, and heuristics in screening out the irrelevant information during the decision making process. These problem characteristics suggested that the fertilizer problem was a good candidate for expert systems technology.

The result is the development of a prototype expert system called the *Fertilizer Advisor*. It takes a knowledge-based approach to perform a constraint-directed search of the problem space. Like a human expert, it applies the constraints known about the problem to limit its search to only the relevant information.

The problem solving strategy used in the Fertilizer Advisor will be presented in this thesis. First we will examine how a specialist determines the economic optimum fertilization plan for applying nitrogen fertilizers. Then we will extend the scope of the problem and use several specialists for determining the economic optimum plans for applying each of the other fertilizer nutrients. The contributions made by all of these specialists are then integrated to form the economic optimum solution for the fertilizer problem. The problem solving strategy is incremental in nature and thus is suitably implemented using an incremental development strategy common to building expert systems. By incrementally extending and refining the system's knowledge, a complex problem such as the fertilizer problem can be solved.

Chapter 1

Introduction

Successful crop production begins with good crop management practices. Each farmer must make crop management decisions concerning seedbed preparation, seed selection, seed treatment, seeding rate, depth and methodology, fertilizer application, and weed, insect and disease control. It is desirable for farmers to optimize production techniques so as to increase the efficiency of their operations. The Manitoba Department of Agriculture has proposed several crop management methodologies to assist Manitoba farmers in achieving this goal. These methodologies are designed to help farmers maximize their net returns on crop production investment. To attain the highest possible net return, farmers must ensure that each management input is applied at the optimum rate and time. In all likelihood, many farmers do not have all the necessary knowledge or expertise to make effective decisions about each and every one of these management inputs. Furthermore, the detailed knowledge needed to help farmers and advisors make effective crop management decisions is complex and scattered. To rectify this problem, researchers at the University of Manitoba from the Department of Computer Science, the Faculty of Agriculture, and the

Solomon Sinclair Farm Management Institute have been examining the application of expert system technology to crop management decision making.

The goal of this research is to formalize the crop production expertise so it can be applied in the form of a computer program. Encapsulating the expertise in a program has a number of advantages. It will provide a permanent record of expertise. Furthermore, the captured expertise can be replicated easily and be distributed to a wide population of the farming community. This distribution will benefit both Manitoba Agriculture and farmers. First, Manitoba Agriculture can ensure that the crop production expert knowledge being distributed is consistent across the province. Secondly, the expert knowledge can assist less experienced agricultural representatives in providing expert level advice to farmers. In the process, junior personnel are in effect being trained by an expert. With the support of a knowledgeable expert system, the inexperienced agricultural representatives will be able to provide farmers in the province a high quality of service.

The initial attempt at achieving this goal is the development of a prototype expert system, known as the *Fertilizer Advisor (FA)*. It is designed to provide assistance to farmers in making decisions about rates, sources, timing, and methods of applying fertilizers. The research on the computerization of an expert system for making fertilizer recommendations has been motivated by the growing popularity of computer applications in the agricultural industry. The ability of the computer to efficiently and quickly perform repetitive tasks such as database management and statistical analysis have helped in improving human productivity. There is a growing desire in the agricultural industry to utilize computer technology to perform tasks of higher complexity. The advent of expert systems provides a solution for this.

1.1 What are Expert Systems?

Expert systems research is an area of research in *artificial intelligence* (AI). AI is “the branch of computer science that is concerned with the automation of intelligent behaviour” [Luger et al., 1989]. For example, in natural language processing, which is another research area in AI, researchers are concerned with the design and construction of programs which simulate humans' ability to accept, understand and produce language for the purpose of communicating effectively. Likewise in robotics research, another research area in AI, robots are designed to duplicate the human senses of sight and touch to perceive and react to changes in a dynamic environment. However, by and large the most successful area of research in AI has been in expert systems.

Expert systems are computerized systems designed to imitate the abilities of experts in solving problems in a narrow problem area. In general, an *expert* is a person who has a significant depth of understanding in some specific field and a system designed to simulate the expert in solving a problem in his/her area of expertise is called an *expert system*. To gain a better fundamental understanding of expert systems, it would be beneficial to first consider the characteristics of human experts.

1.1.1 Human Experts

An expert is “a person who, because of training and experience, is able to do things the rest of us can not; experts are not only proficient but also smooth and efficient in the actions they take” [Waterman, 1986]. Unlike a novice, experts can perform at high levels when solving difficult problems. A high level of performance is the result of combining a strong theoretical background in the problem domain with effective **heuristic**¹ rules of problem-

¹ A heuristic is "rule of thumb or simplification that limits the search for solutions in domains that are difficult or poorly understood" (Waterman 1986).

solving that have accumulated from experience in dealing with domain problems [Luger et al., 1989]. All human experts have several common characteristics.

- 1) Human experts are able to solve elementary and difficult problems in their area of expertise.
- 2) A human problem solver is able to explain the reasoning processes which led him/her to a solution.
- 3) Human problem solvers have the ability to learn and acquire new knowledge and skills in an attempt to broaden their problem solving capabilities. Humans can also dynamically modify or extend previously learned knowledge to conceive new concepts for solving problems in a novel situation.
- 4) Experts are “good at plowing through irrelevant information in order to get to the relevant issues” [Waterman, 1986].
- 5) In solving problems where it is difficult to arrive at a solution with complete confidence, the human expert often can provide some measure of solution reliability [Miller, 1986].
- 6) Expert problem solvers possess the power of discernment—the ability to recognize when a problem is outside their area of expertise.
- 7) Human experts are known to be apt at balancing the effort expended on deriving a solution to the quality of that solution. When an acceptable answer to a problem is produced, no extra effort will be expended on deriving an optimum solution unless that effort can be reasonably justified.

Experts systems researchers attempt to replicate, to some degree, the features of a human expert outlined above. A closer analysis of the features and components of expert systems will provide insight into what extent these goals have been achieved.

1.1.2 Expert Systems

Expert systems have a number of distinguishing features. The following is a list of these features analogous to the features of human experts listed above.

- 1) Every expert system contains a base of knowledge about some particular domain. When the knowledge is for solving problems which are not considered difficult enough to require an expert, the system is known as a knowledge-based (or simply knowledge) system. Similarly, the knowledge contained in an expert system can be used to solve expert-level problems. An expert system is a specialized instance of a knowledge-based

system.² The knowledge in an expert system refers to the information that has been extracted from experts in the problem domain. The expert's knowledge about a domain can exist as facts or heuristics. Facts are propositions that are known, either by observation or experience, to be true. Heuristics are subjective rules of good judgement ("rules of thumb") that characterize expert-level decision making in a particular field. Generally, heuristics are established through experience and can not be learned from the standard theory presented in textbooks and classes. The knowledge base is what gives an expert system its high-level problem solving capability. "The accumulation and codification of knowledge is one of the most important aspects of an expert system" [Waterman, 1986].

- 2) Knowledge in an expert system is explicit and accessible, and hence, can be easily understood and reproduced. Expert systems have the quality of being "open to inspection" [Firebaugh, 1988]. Users can interact with the expert system at various intermediate stages of execution and request the expert system to explain the reasoning processes used to derive its decisions. Such a capability has proven extremely valuable for expert system developers responsible for maintaining and extending knowledge bases.
- 3) Like human experts, expert systems acquire their knowledge gradually over time. The building of an expert system is of an exploratory nature. Generally, a knowledge engineer extracts knowledge from one or more domain experts, and implements the knowledge in a concise and efficient knowledge base. A knowledge engineer³ is a "person who designs and builds the expert system. This person is usually a computer scientist experienced in applied artificial intelligence methods" [Waterman, 1986]. Preferably, a domain expert is someone with a deep understanding of the domain and able to solve difficult problems in the domain quickly and concisely. The experts provide the bulk of a system's knowledge (additional knowledge can be gathered from empirical data, case studies, texts and other sources). The scope of the problem being covered by an expert system is constantly growing as the knowledge engineer transfers the experts' knowledge to the expert system.
- 4) The quality of the domain experts and the conciseness with which the experts' knowledge has been represented will be reflected in the performance of the resulting expert system. Experts usually use heuristics and shortcuts to perform efficient searches for the problem solution. The ensuing expert system should model these optimizing skills.
- 5) An expert system is not limited to dealing with factual knowledge which it has complete confidence in being true. For example, heuristic knowledge is imprecise and can be represented in an expert system by associating each bit of knowledge with some measure of confidence. In general, this measure of confidence is called a *certainty factor*. The expert system manipulates and combines the certainty factors using some prescribed formalisms until eventually a measure of certainty is propagated to the final solution. The association of a confidence measure with a solution suggests two important points. First, if a solution can not be suggested with complete confidence, then there is a chance that the solution is incorrect, implying that expert systems are prone to making mistakes. Second, if expert systems are capable of making mistakes, those people using them

² The term expert system is often abused by those impressed with the implications of the phrase. In reality, seldom does a system reach a level of competence that is deserving of the title expert. In this thesis, in accordance with most of the available literature, the name most often used will be expert system.

³ The equivalent to a knowledge engineer in a more conventional environment is the software engineer or system analyst.

should be aware of this fact and should exercise caution when deciding whether or not to accept a recommendation.

- 6) An expert system usually exhibits high-level expertise in one domain area. It can not solve problems outside of its domain of expertise. The robustness of an expert system measures the completeness of the system's knowledge in regards to solving a variety of problems in the domain. A robust expert system should follow a gradual degradation in performance when it is posed with problems which are at the edges of its scope of expertise. An expert system may require metaknowledge (i.e., knowledge about how to use the domain knowledge) in order to detect that a given problem is outside of its domain of expertise.
- 7) An expert system may be given the knowledge to balance the amount of effort expended on solving a particular problem with the quality of the solution. Again, as in point four, this feature depends on the quality of the knowledge transferred from the domain experts.

1.1.3 Conventional Computer Systems versus Expert Systems

Conventional systems are designed to perform algorithmic tasks on large volumes of data. The goal is to produce exact results efficiently. The results produced by conventional systems must be correct otherwise they are simply not acceptable. For example, an electronic card catalogue system in a library is designed to manage a database of book records. The operations are straightforward: add a book record, mark a book as being borrowed, print a list of books by a particular author, etc. The tasks to perform are not complex and require little expert-level intelligence.

Expert systems on the other hand are designed to solve problems by manipulating specialized knowledge in the problem domain. Like experts, the results produced may not be exact. The quality of the solutions depends on the quality of the problem-solving knowledge in the expert system. The recommendations they produce are generally neither correct nor incorrect, but only more or less plausible [Harmon, 1985]. This is due to the heuristic nature of their knowledge.

Common languages used for developing conventional systems are procedural languages such as COBOL, PL/I, Fortran, or Pascal. Procedural programs are expressed using

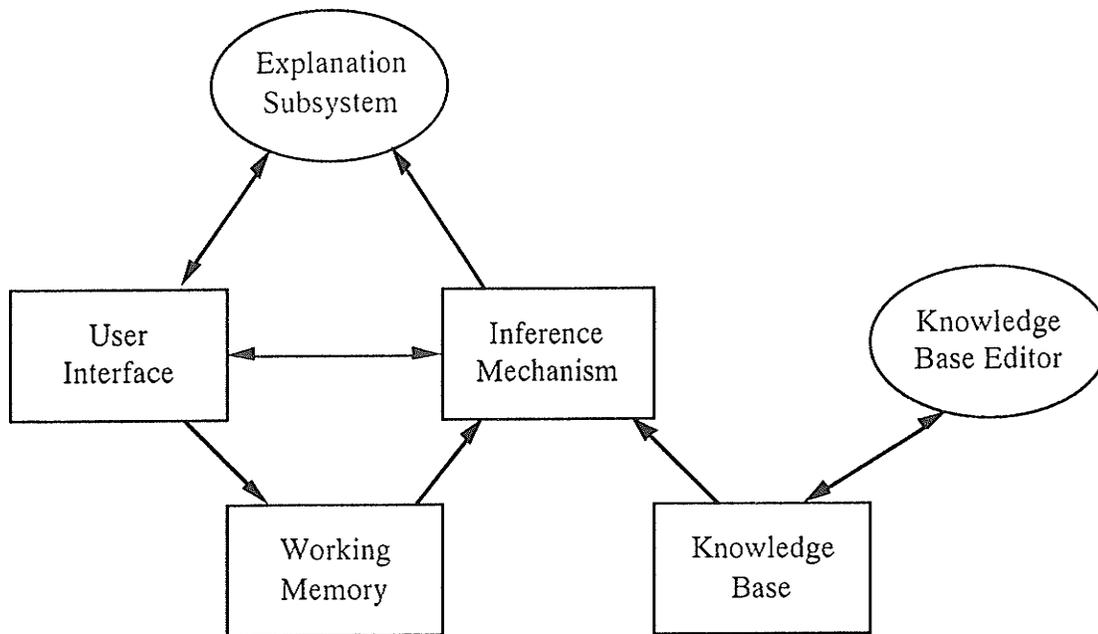


Figure 1.1: Conceptual expert system architecture
Adapted from [Evans et al., 1989]

control structures, logical expressions, variable assignments, etc. Although modular programming methodologies have been designed for building large complex conventional systems, the resulting code is ultimately difficult to understand and modify. Typical development of conventional systems is to specify the design plans before implementing the system. All problem solving methods must be specified in advance.

Expert systems are developed using a declarative programming style. Declarative systems separate the knowledge from the inference mechanism which decides in what order to apply the knowledge to solve a particular problem (refer to Figure 1.1 for the conceptual components of an expert system). The expert system builder creates the knowledge base of what is known about the problem and the system utilizes the knowledge to determine how to solve the problem. The execution of the control strategy for manipulating the knowledge base may be less efficient but the knowledge in the system is explicit, and is easy to understand and extend. It is desirable to have knowledge that is easy to change. This is

because expert systems are developed incrementally through successive extensions and modifications. At the start, the knowledge engineer does not know all the methods for solving a particular problem. More so, the knowledge engineer is not capable of understanding all of the problem solving knowledge at once. The expert system must be built up incrementally by implementing the knowledge as the knowledge engineer slowly begins to understand the problem and its solution methods. There are two ways in which the knowledge engineer (in cooperation with the domain experts) increases the expert system's level of performance: 1) the knowledge present in the system is gradually refined over time and 2) additional problem solving knowledge is put into the expert system. It is not possible to completely specify a problem and all of its solution methods in advance (Figure 1.2 depicts the expert system development process).

As can be seen by the above comparison, expert systems and conventional systems each have their strengths and weaknesses. It is most advantageous to employ expert systems technology in conjunction with conventional computing techniques to complement each

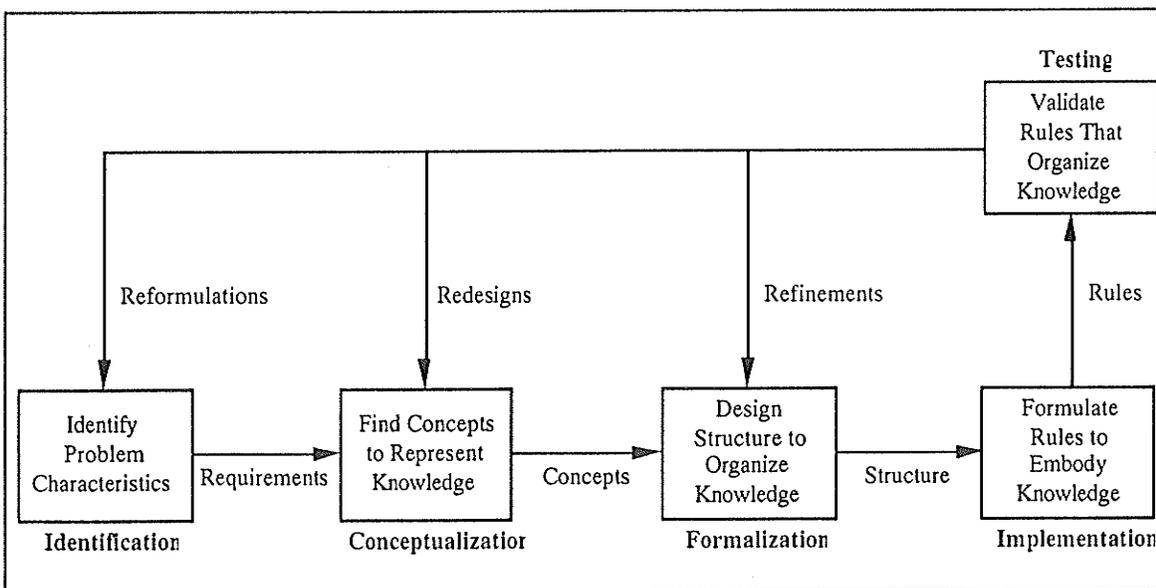


Figure 1.2: Stages of Expert Systems Development
Adapted from [Evans et al., 1989]

other's qualities. The Fertilizer Advisor is an example of a program which integrates both conventional and expert systems technology. This expert system is designed to solve what we call the fertilizer problem.

1.2 Fertilization Decisions

1.2.1 Overview of the Application Domain

During the early days of agriculture, farming was a tradition passed on from one generation to the next. It was necessary for it was a means of providing food for the community. Today, farming in the prairies has become a thriving business. Through modern transportation systems, the produce of a farm is made available to a "community" which is no longer restricted to nearby neighbours. Furthermore, the crops produced by a farm are not only a useful food source, but they can also be used for the production of other resources such as feed, fibre, pharmaceuticals, and fuel. The impact of science has turned crop production into a growing industry. Years of accumulated traditional farming experience and scientific research has changed crop production into a science. Several uncontrollable factors such as weather remain. For example, droughts and frosts can have an undesirable effect on that desirable target yield a farmer is planning for. But if the weather is favorable, it is quite possible for a farmer to reach a target yield provided that the right actions are taken.

Successful crop production begins with proper crop management practices. Choosing certified seed of a recommended variety will ensure that weed seeds and diseased kernels are removed, leaving only the high quality product to give good germination. Further treatment of the seeds with fungicides can lessen the likelihood of diseases. Before actual seeding, the seedbed must be well prepared. For example, major weed problems should be eliminated and the seedbed should be firm to allow shallow seeding because the moisture

condition near the surface is better. To properly seed, the rate, distribution, depth and method of seeding must be carefully chosen. Fertilizer application will provide the necessary nutrition to foster crop growth. An equally important step is pest control. This may include the prevention and/or removal of weeds, insect pests and plant diseases. Lastly, a good harvesting method will ensure an optimum recovery of the crop yield. The timeliness of the application of the above steps is also very important.

It is essential that the above crop management practices are followed in order to maximize crop production. These crop management steps are inter-related and each plays a very important part in improving crop growth. The necessary management practices to carry out vary from one farmer to another and may even vary from field to field. For example, one field may have soil which has a sufficiently high concentration of the required nutrients while another field with nutrient-deficient soil may require heavy fertilization. Each crop management step mentioned entails a different domain of knowledge. The focus of this thesis is on the fertilization step. The solutions to be presented on solving the problem of fertilization are realizable on the condition that the other management practices are correctly exercised. The expertise required in making recommendations for these other management practices are the subject of separate expert systems. It is envisioned that future research will explore expert systems which address each one of these management steps; resulting in a collection of expert systems which can be integrated together to make recommendations for a complete crop production plan. The remainder of this chapter will discuss the problem of making recommendations for fertilizer application.

1.2.2 Current System for Making Fertilizer Selection Decisions

At present, Manitoba Agriculture has a system in place for aiding Manitoba farmers and advisors in making decisions concerning fertilizers. Farmers can turn to Manitoba Agriculture for information on how to maximize their crop production potential. The main

component of this system is the Manitoba Provincial Soil Testing Laboratory. Farmers, agricultural representatives, and fertilizer dealers send soil samples to the Provincial Lab for analysis of nutrient concentrations in the soil. On the basis of this analysis, fertilizer recommendations are provided. For all crops, a *standard* soil test recommendation is given which suggests a *single application rate* for each nutrient required. These recommendations are said to provide a good economic return under reasonable Manitoba growing conditions.

This system has been criticized by many as being somewhat crude for a number of reasons. Firstly, the recommendations only suggest rates of each nutrient to apply. The farmer is free to use any combination of fertilizers, placement methods and timing to achieve these suggested nutrient levels. Selecting a fertilizer plan that is cost effective is a complex task. One farmer, for example, may be able to capitalize on low prices for a particular fertilizer whereas a neighbour may not because of a lack of suitable fertilizer application equipment. Secondly, the recommendations are standardized for one particular set of economic and environmental circumstances. In reality, each farmer deals with a different and dynamic set of crop and fertilizer prices. Growing conditions, being highly dependent on soil type and climatic factors, are also highly variable across many areas of the province as well as from year to year. Thirdly, the standardized recommendations are conservative. The low-risk nature of these recommendations also implies a less than optimum return on investment. Farmers who are willing to take a higher risk by increasing the fertilizer investment may realize a relatively higher crop yield. Therefore, for the best results, farmers should have the data from the soil test interpreted by an agronomist and tailored towards their specific conditions before fertilization begins.

1.2.3 Making a Fertilization Plan

The ideal recommendation would be a *fertilization plan* which would provide the farmer with information concerning the fertilizer compound(s) to apply, the amount of each

fertilizer compound(s) to apply, the method of application, and when the application should be performed (see Figure 1.3). Given a fertilization plan, additional information such as the per hectare cost of applying the fertilizer, amount of each nutrient applied and so on can be calculated. The recommended plan must also meet a number of goals:

- 1) achieve the maximum dollar return on the investment in crop fertilization
- 2) meet the requirements of the four nutrients: nitrogen (N), phosphorus (P), potassium (K), and sulphur (S), as dictated by the standard Manitoba Soil Testing Laboratory recommendations.
- 3) plan is customized to meet the farmer's needs
(e.g., equipment availability, limited selection of fertilizer products)

To achieve each of these goals, it is necessary to have knowledge about crop economics, soil science and the farmer's preferences. Crop economic knowledge such as the basic economics of field crop production, and soil and crop management knowledge concerning the yield potential of various crops, fertilizer toxicity behaviours and other factors, will shape the eventual fertilization plan into an expert level recommendation. Other factual data such as fertilizer compounds accessible to the farmer, their nutrient contents, prices, and toxicity behaviours, the forms in which they are available, the cost and relative efficiency of application methods, expected crop sale prices, and farmer's field conditions form the base of information from which fertilization plans can be constructed. Integrating the

<u>FERTILIZATION PLAN</u>	
<i>Source</i>	of fertilizer compound(s) to apply
<i>Rate</i>	of fertilizer compound(s) to apply
<i>Method</i>	of application
<i>Timing</i>	of application

Figure 1.3: Information offered by a fertilization plan

consideration for the farmer's preferences in the decision making process will help in building a fertilization plan suited to the farmer's needs. Working with the knowledge domains outlined above involves the manipulation of many variables.

1.2.4 The Fertilizer Problem

Ideally, formulating a recommendation concerning additional crop nutrients should take into consideration the following variables:

- concentration of nutrients already in the soil (*i.e.*, soil test results);
- soil type;
- soil moisture content at planting time;
- regional climatic conditions including growing season precipitation;
- the effect of method, cost, and time of application on *toxicity constraints* and the efficiency of added nutrients; (toxicity constraints will be discussed in detail later);
- different fertilizer's nutrient analysis, form, and price;
- the client's own preferences; a farmer may have certain equipment constraints, dollar constraints, risk aversions, or simply a desire to carry on traditional practices.

Each of the above variables can take on many values. For example, climatic conditions, soil type, and soil moisture content may take on a variety of different values from one field to another and from one year to another. Fertilizers can be applied either in the spring or fall using one or two combination of the following methods: banded prior to planting, banded with the seed, banded beside/under the seed, or broadcast. The common nutrients, nitrogen, phosphorus, potassium, and sulphur, can be combined in numerous ways to yield the various fertilizers in granular, gas, or liquid form. In addition, the preferences of the user should be addressed in order to provide a customized fertilizer recommendation. The difficulty in making effective fertilizer recommendations is in considering how to work with all the variables to arrive at the best recommendation—this is called the fertilizer problem.

Consider first the effects of a single nutrient, nitrogen. In Manitoba, nitrogen is the major nutrient required for successful crop growth. If it is determined that the supply of nitrogen in the soil prior to fertilization is not sufficient to yield an optimum return, the cost of adding more nitrogen to the soil is considered. The total cost of additional nitrogen is a combination of the cost of the fertilizer itself (*i.e.*, the cost per unit increment of nitrogen times the number of increments) plus the cost of application. Application costs are dependent on the method of application and may include such expenditures as fuel, repairs, time, and labour. The cost per unit increment of nitrogen supply is a function of fertilizer price, percentage of nitrogen in the fertilizer, and the efficiency by which the application method and time of application increases the nitrogen supply in the soil. Gross revenue, on the other hand, is a function of the projected crop price multiplied by the projected yield. Yield is, to a large degree, dependent on available moisture and total nitrogen supply (including added nitrogen). For a summary of these dependencies, see Figure 1.4.

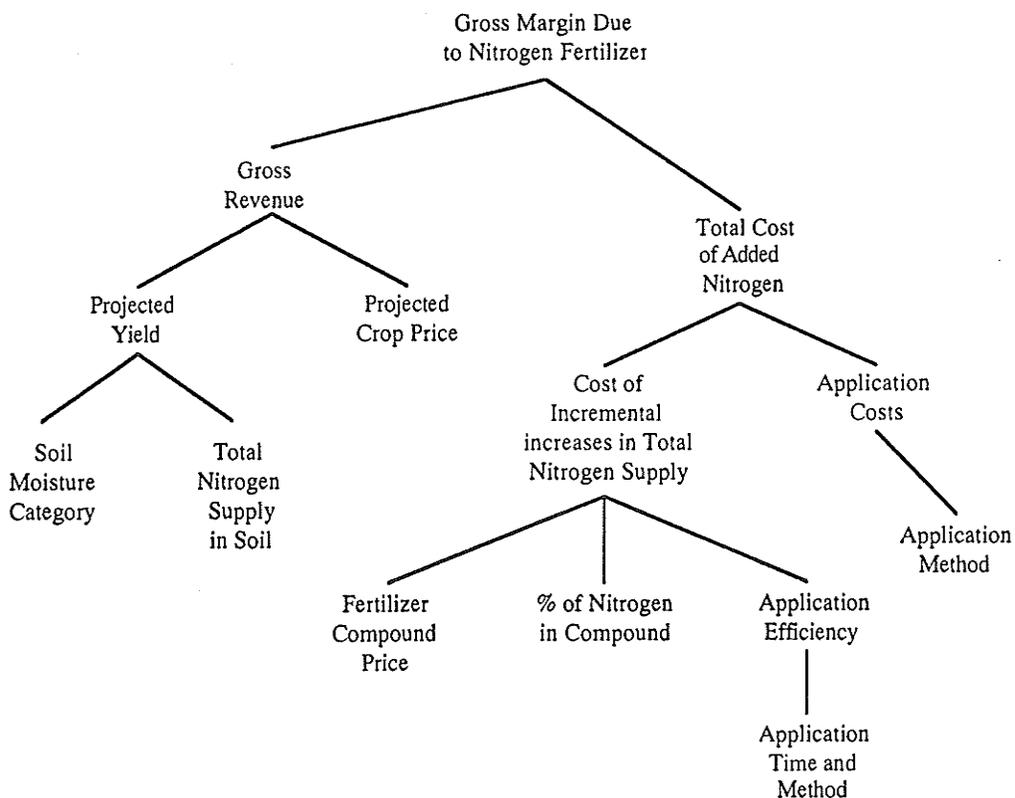


Figure 1.4: Task overview for a given level of existing nitrogen supply in soil.

In theory, to solve the fertilizer problem, one could consider all potential recommendations by generating all possible combinations of the variables and selecting the most cost effective alternatives (*i.e.*, an exhaustive search of the problem space). In fact, such an approach is not viable since the computational complexity of this problem is exponential in nature, implying that a brute force search of the problem space would not be cost effective. After consultation with various experts in crop management and soil sciences, we determined that agronomists do not use a formal algorithm for generating solutions to the fertilizer problem. Instead, agronomists rely on experience, intuition, and heuristics when making fertilizer decisions. The lack of a formal algorithm, combined with the recognition of the heuristic methods, suggested that the fertilizer problem was a good candidate for expert systems technology.

The FA expert system was built using an incremental development strategy. The initial step is to determine the feasibility of the project by solving a simplified version of the fertilizer problem. Chapter 2 will discuss the problem solving strategy in which the economic optimum fertilization plan⁴ is selected by examining the yield and cost curves representing the farmer's fertilization options. The simplified problem concerns the selection of the fertilization plan for applying the single nutrient nitrogen to obtain an optimum yield. Chapter 3 will present a blackboard framework in which the solution from chapter 2 can be scaled up to solve the fertilizer problem. That is, the four nutrient requirements N, P, K, and S must be met within the constraints imposed by the farmer's circumstances while achieving the optimum net return on fertilization investment. This chapter will show a solution method in which the solution is achieved incrementally by solving portions of the problem and integrating these partial solutions together to arrive at the final solution.

⁴The fertilization plans selected by the Fertilizer Advisor are expected to achieve an optimum economic return on fertilization investment provided that the information considered by the system are accurate. However, information such as weather, market price of the yielded crop, etc. can only be forecasted. In this thesis, the term optimum implies being very close to the optimum.

Chapter 4 will discuss the major stages of incremental development used to gradually build the solution presented in chapters 2 and 3. The discussion will include the objectives, the achievements, the problems encountered at each stage of development. Chapter 5 will be the conclusion to the thesis.

Chapter 2

The Nitrogen-Only Problem

The previous chapter discussed the large number of factors that must be considered in the fertilizer selection problem. In order to handle the complex interactions between these factors, it is best to place the initial focus on a small, manageable area of the problem domain, find a solution, and incrementally scale the scope of the system up to solve larger portions of the problem [Evans et al., 1990a]. This is a common technique used in building knowledge-based systems. The size and the complexity of the problem is reduced, thus allowing us to concentrate on details. This incremental development technique was employed in building the Fertilizer Advisor. The initial focus of the prototype development was on solving the *nitrogen-only problem*. This primarily involves the generation of fertilization plans which must fulfill some nitrogen requirement.

Nitrogen (N) fertilizer is effective in increasing the yield of crops. The effectiveness of fertilizer nitrogen is dependent on the supply of other nutrients, the moisture available to the crop, the level of available soil nitrogen, and losses of applied nitrogen, which are affected by the time and method of application. These are but some of the factors that will affect the

crop's response to nitrogen fertilizer. A proper balance of all nutrients essential for crop production must be maintained to obtain optimum yield increases from the application of nitrogen fertilizers [Saskatchewan, 1986]. In the nitrogen-only problem, the other nutrients phosphorus, potassium, and sulphur are assumed to be at optimum levels, allowing the effectiveness of the nitrogen to be maximized. The rest of this chapter will describe the solution to the nitrogen-only fertilizer selection problem. The next chapter will discuss the process of scaling the nitrogen-only solution up to generating recommendations for fulfilling all nutrient requirements.

2.1 Nitrogen Recommendations

The nitrogen-only problem can be categorized into two types of recommendations: 1) the *yield curve type*, and 2) the *grid type*. The type of recommendation to perform is dependent on the type of crop being fertilized. For example, crops which have tables predicting their potential crop production yield should generate the yield curve recommendations. The grid type of recommendations is a simplified case of the yield curve type of recommendations and will be dealt with after the yield curve type.

2.1.1 Yield Curve Type Recommendations

The main objective in formulating a recommendation for fertilizer application based on a crop production yield curve is to maximize dollar returns on the investment in fertilization. This type of recommendation takes advantage of the fact that for some crops, there are tables which describe the relationship between the soil nitrogen available to the crop and the potential crop yield under normal growing conditions. This relationship allows us to compare the economic benefits of various fertilization plans which achieve different soil nitrogen levels. In the context of solving the fertilizer problem, the economic benefit of a fertilization plan is defined as the net return on fertilization investment and the net return is the gross revenue acquired from the sale of the harvested produce minus the total

fertilization costs (*i.e.*, fertilizer costs and application costs). We are interested in the optimum yield, or more precisely, the yield that will provide the greatest margin of return over fertilization costs.

2.1.1.1 Yield Curves

The task of determining the optimum yield is complex. Generally, yield is described as a function of nitrogen supply as depicted in Figure 2.1. The curve describing the yield at each increment of nitrogen can be derived in one of two ways. First, a complex equation based on such variables as water supply, nitrogen supply, and growing degree days could be used to plot a continuous curve. However, to the best of our knowledge no equation exists that can be used to accurately predict the yield response of any crop. Second, a tabular approach based on observations of multiple trials under similar conditions can be used to estimate the response for set increments in nitrogen supply. The second alternative, a tabular approach, has been adopted for use in the FA system. This approach, when illustrated graphically, results in a segmented curve. However, for illustrative purposes, the curve has been smoothed.

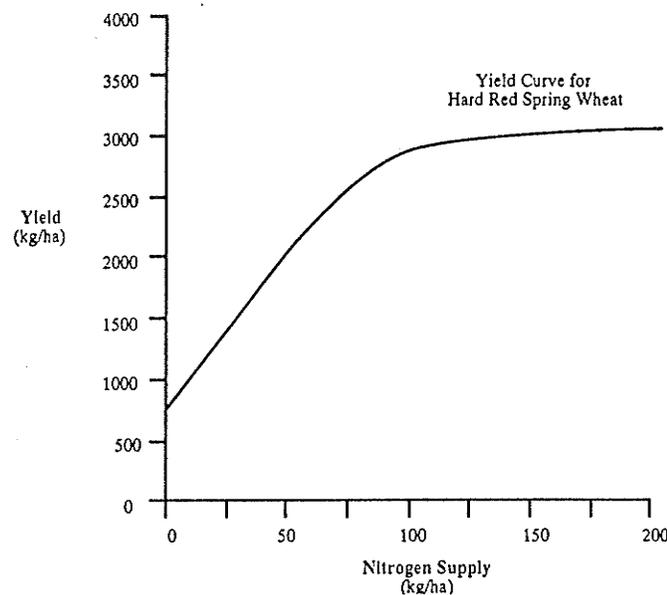


Figure 2.1: A yield curve for Hard Red Spring Wheat under moist growing conditions.

Nitrogen Supply (kg/ha)	Yield
0	800
10	1120
20	1440
30	1760
40	2080
50	2400
60	2680
70	2875
80	3005
90	3095
100	3150
110	3185
120	3200
130	3200

Table 2.1: Yield table for Hard Red Spring Wheat under moist growing conditions.

The specific values used for plotting a particular crop's yield response curve are adapted from the Yield/Nitrogen Supply Tables used by the Manitoba Provincial Soil Testing Laboratory (see Table 2.1). Different crops exhibit different yield responses and the same crop will exhibit different yield responses under different soil moisture conditions. For example, there are four separate tables available for Hard Red Spring Wheat; one for each of four moisture regimes (*i.e.*, arid, dry, moist, and ideal). Each table describes the *average expected yield* (kg/ha) for increments of total nitrogen supply in the soil. The FA prototype uses slight modifications to these tables where the curves are smoothed over finer increments of nitrogen supply.

2.1.1.2 Cost Curves

The second step in determining the optimum yield involves establishing a corresponding fertilization cost curve. The cost curve for the application of one fertilizer compound using one application method is linear (see Figure 2.2). A *simple* cost curve can be calculated by determining the cumulative costs of adding each increment of nitrogen using a specific application method and time, and a single fertilizer.

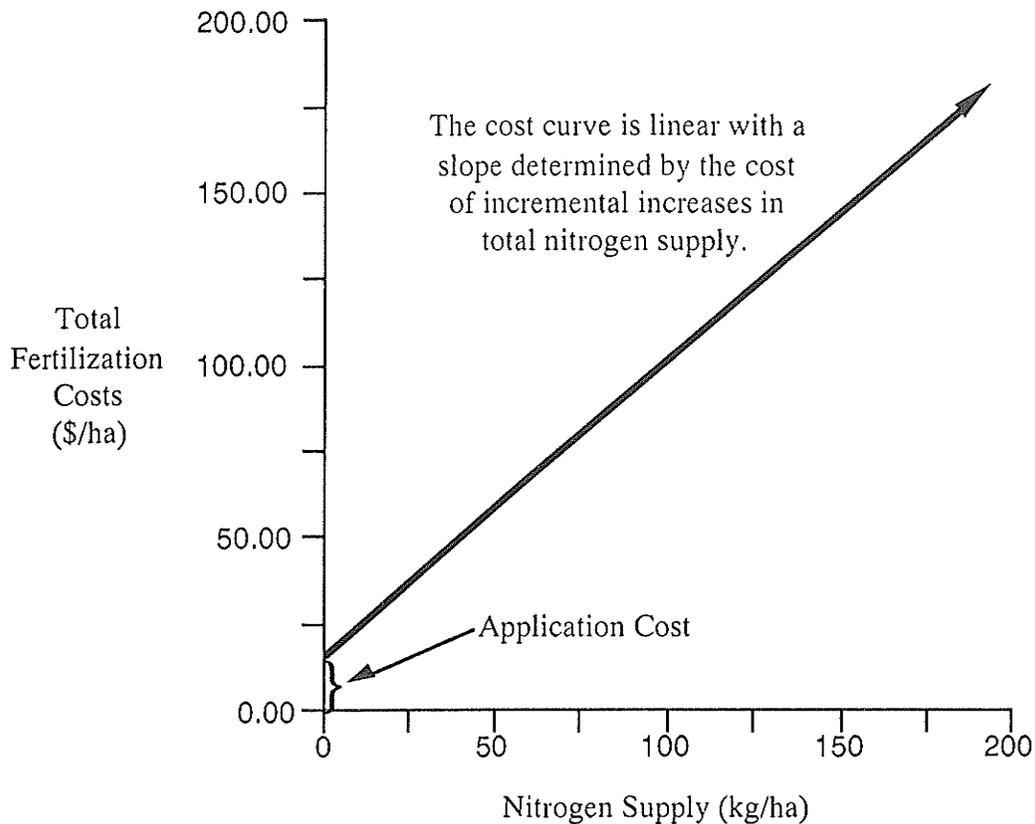


Figure 2.2: Cost curve for applying one fertilizer compound using one method.

2.1.1.3 Net Return on Fertilizer Investment

The third step is to superimpose the cost curve on the yield curve. By definition, return equals revenue generated from sales minus total cost of fertilization. This is represented by the distance between the two curves. The point at which the distance between the two curves is greatest marks the maximum net returns for that fertilizer and application method (see Figure 2.3). This is the optimum yield which we are interested in. To find this, a search is required along the nitrogen supply axis for the point of greatest distance between the two curves. Information can be extracted from the point of optimum return to formulate a complete fertilization plan. The method-fertilizer combination represented by the cost curve determines the time and method of application and the fertilizer compound to recommend. The point of optimum return gives the amount of nitrogen to supply. This is used to calculate the amount of the selected fertilizer compound to apply.

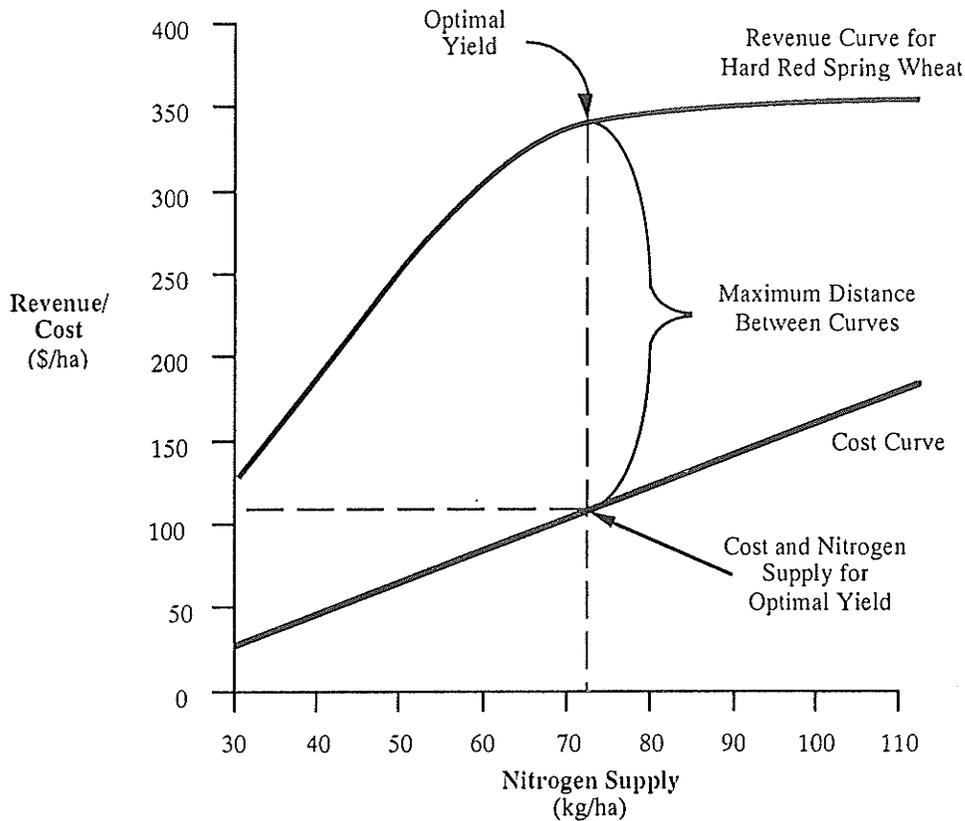


Figure 2.3: A fertilization cost curve superimposed on a revenue curve.

2.1.1.4 Complete and Partial Cost Curves

Recall that this is a simple scenario dealing with one method of applying one fertilizer compound. The system must also consider other possibilities of applying different fertilizer compounds using different application methods thus resulting in multiple curves. To further complicate the problem, a single cost curve may involve multiple application methods and possibly multiple fertilizer compounds; such cost curves will be referred to as *partial* cost curves because using a single application method with a fertilizer compound can only generate a part of the curve, requiring a second application method to complete the rest of the curve. For the purpose of definition, a curve is complete if the cost for supplying all increments of 5 kg/ha of nitrogen up to 175 kg/ha can be calculated (it is rare that an application of more than 175 kg/ha of nitrogen supply is necessary).

A partial cost curve is characteristic for the *seed-row* method of application. Here, the fertilizer is applied with the seed at planting time. In close proximity to the seed, the concentration of the fertilizer may be toxic, thereby injuring the germinating seedling. For each fertilizer, there is an upper limit on the amount that can be applied using a seed row method before the rate is deemed to be toxic. This upper bound is called the *toxicity limit* or *toxicity constraint*. Using a seed row method may be economically attractive because it is done concurrently with seeding and usually does not incur an extra application cost. But when the toxicity limit is reached, it is not viable to continue adding more of the same fertilizer with the seed. In order to contribute additional fertilizer nutrients, it is necessary to increase the rate of fertilization using a non-seed row method. There is an application cost involved in switching over to a non-seed row method because the non-seed row application has to be performed at a different time. The same fertilizer or a different fertilizer can be applied using the non-seed row method.

The cost curve illustrated in Figure 2.2 is an example of a *complete* cost curve; applying one fertilizer using one method. A complete cost curve is associated with a *non-seed row* method of application. Such an application method involves either broadcasting the fertilizer onto the soil surface before or after seeding, or applying the fertilizer in bands close to the seed bed (but separate from the seed) before or after seeding. Using a method of this type, the fertilizer is not concentrated around the seed and thus will not cause toxicity problems. The corresponding cost curve can be graphically represented by a linear curve. Note that this type of application method is performed before or after seeding, implying that an additional field operation is required to apply fertilizer. This extra operation is taken into account as an application cost. Figure 2.4 illustrates seed row and non-seed row application methods and their corresponding partial and complete cost curves.

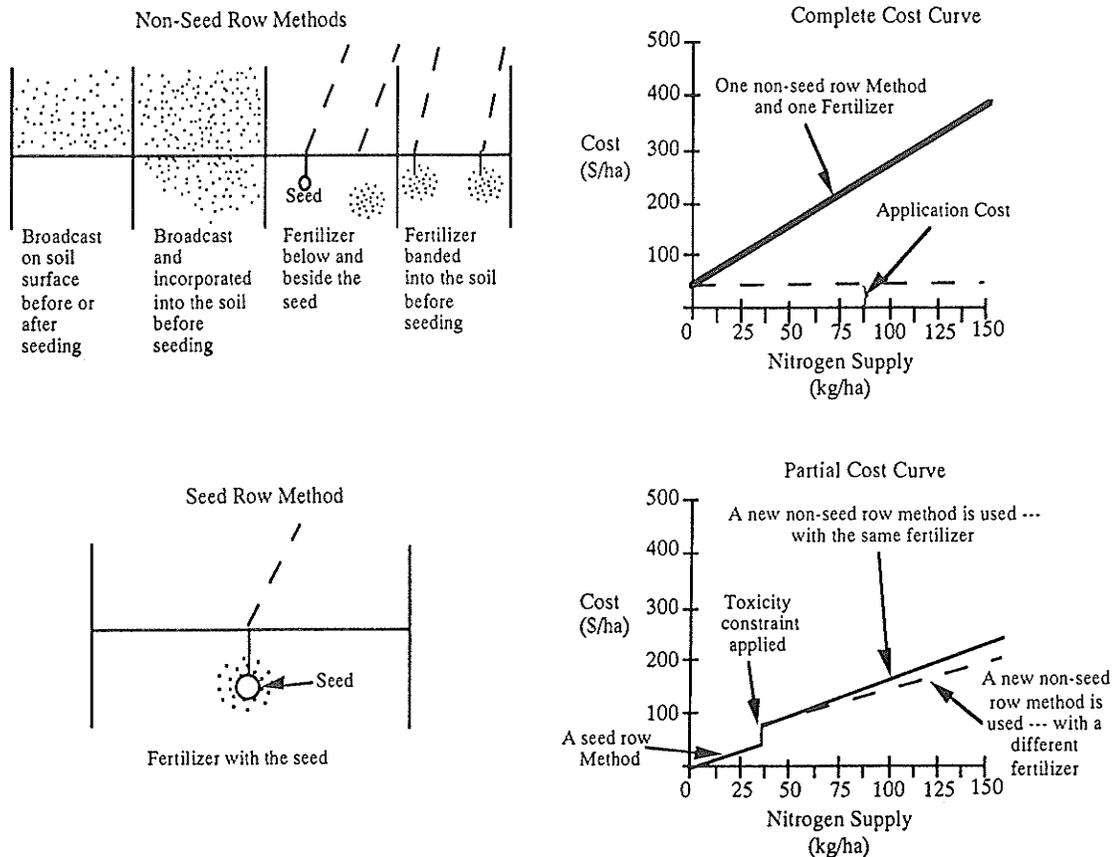


Figure 2.4: Examples of non-seed row methods, a complete cost curve characteristic of non-seed row application of fertilizers, seed row method, and a partial cost curve characteristic of seed row application of fertilizers

2.1.1.5 Multiple Cost Curves

The effect of calculating and comparing cost curves for all possible combinations of applying different fertilizers with different methods will result in a graph similar to that in Figure 2.5. This graph is a simple example depicting three cost curves: two partial cost curves and one complete cost curve. The lowest cost segments form the optimum cost curve. Graphically, to find the best return, we search for the greatest distance between the revenue curve and this optimum cost curve. Computationally, a search vertically between all cost curves to find the lowest cost curve is necessary at each increment along the nitrogen supply axis while searching for the point of optimum return.

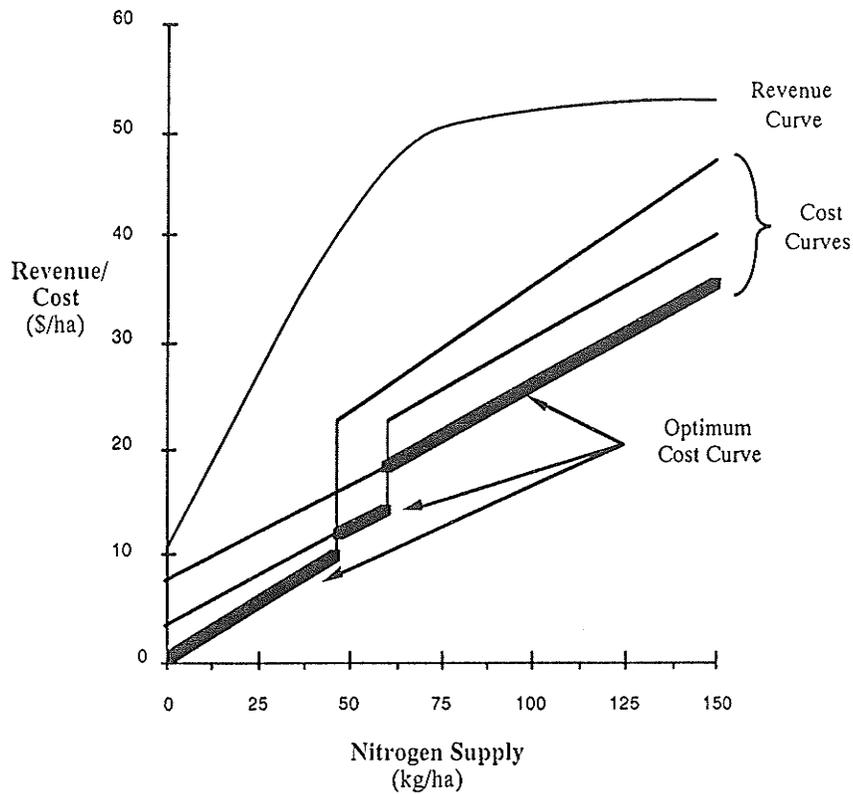


Figure 2.5: A simple example of three cost curves. The thickened segments form the optimum cost curve.

Consider the case where the system has in its database four fertilizer compounds and five different application methods (one seed row and four non-seed row) for applying these fertilizers. First determine all possible combinations of application method/fertilizer pairs, then generate a cost curve for each combination. A brute-force basic algorithm for the generation of these cost curves appears in Figure 2.6.

Four fertilizer compounds, where each can be applied in five different ways yield 20 potential complete cost curves and as many as 64 additional cost curves entailing combinations of seed row application and non-seed row application. A search through these different curves to find the point of optimum return is time-consuming. If more methods or fertilizers are added to the database, the number of combinations/cost curves to search through grows exponentially.

```

For each increment of nitrogen supply
  For each application method
    For each applicable fertilizer compound
      If the application method is form-compatible with the fertilizer compound
        If it is a seed row application method
          If toxicity constraint has been exceeded
            For each non-seed row application method
              For each applicable fertilizer compound
                Determine how much it will cost to achieve the current level
                of nitrogen supply using the current application method(s)
                and fertilizer compound(s)
              Endfor
            Endfor
          Else
            Determine how much it will cost to achieve the target nitrogen level
          Endif
        Else If it is a non-seed row application method
          Determine how much it will cost to achieve the target nitrogen level
        Endif
      Endif
    Endfor
  Endfor
Endfor

```

Figure 2.6: Pseudo-code for a brute-force method of generating all possible cost curves.

2.1.2 Grid Type Recommendations

The grid type recommendations are performed for crops which do not have tables which describe the potential crop yield based on a given nitrogen level. Instead, these recommendations make use of Standard Nitrogen Recommendation Tables available from the Manitoba Soil Testing Laboratory. A sample table for potatoes appears in Table 2.2. From the results of a soil sample test, the amount of nitrogen currently in the soil can be used as the search key to look up the corresponding recommendation in the table. The recommended value is the appropriate amount of nitrogen to supply through fertilization. The remainder of the task is to search for the lowest costing fertilization plan which will produce the recommended nitrogen supply.

Finding the cheapest fertilization plan for a given nitrogen supply increment has already

Manitoba Provincial Soil Testing Laboratory Standard N Recommendations	
Soil Test N (Soil Nitrate Nitrogen, kg/ha)	Amount of Nitrogen to Add (N, kg/ha)
0	100
5	95
10	90
15	86
20	81
25	77
30	72
35	67
40	62
45	58
50	53
55	48
60	44
65	39
70	35
75	35
80	35

Table 2.2: Grid type recommendation table for Potatoes

been described in the previous section on yield curve type recommendations. Graphically, it is equivalent to vertically searching between cost curves at the given nitrogen supply increment to find the lowest costing fertilization plan.

In the process of searching for the most economical fertilization plan, other solutions which are less than optimum will be considered. These solutions form a list of alternative means for achieving the recommended rate of nitrogen application. The recommendations are kept for reference purposes and are listed in ascending order by cost. Each entry in the table is a valid fertilization plan for achieving the given nutrient supply; only one is lowest in cost.

2.2 Implementation of the Nitrogen-Only Module

For the FA prototype, we have taken a knowledge-based approach to constrain the exponential complexity of the fertilizer problem. Through knowledge engineering sessions with our experts, constraints are uncovered to reduce the amount of search that must be

performed in calculating the final cost curve. The constraints identified can be divided into two categories: *domain-specific constraints* and *farm-specific constraints*. Domain-specific constraints are those imposed by nature; conditions which hold within Manitoba regardless of the farm being considered. An example is the toxicity constraint of a fertilizer if applied in the seed row. Farm-specific constraints are those related to the particular farm being analyzed. For example, a farmer may constrain the choice of application methods because the farmer lacks the machinery to carry out a particular type of application.

The majority of work being performed by the FA prototype is in generating and searching through the cost curves [Evans et al., 1990b]. It is desirable to avoid generating cost curves for method-fertilizer combinations which are not potential solutions. For example, it is not necessary to generate cost curves which are known to be economically ineffective. To this end, a pre-processing stage which utilizes both domain-specific and farm-specific knowledge to prune out imprudent method-fertilizer combinations in the problem is inserted before the cost curves are generated. The nitrogen agent's processing is divided into four separate stages: rule base processing, pre-processing, main processing, and post-processing (see Figure 2.7).

2.2.1 Rule Base Processing

The rule base processing utilizes the farm-specific information entered by the user to modify the fertilizer and/or application method information. Some of the default constraint values such as toxicity constraints can be tightened or relaxed based on farm-specific information entered by the user (see Figure 2.8). The toxicity information is therefore tailored to the particular farm situation being considered rather than to all farms in one geographic area. This allows the system to work with case-specific information in the later stages of processing.

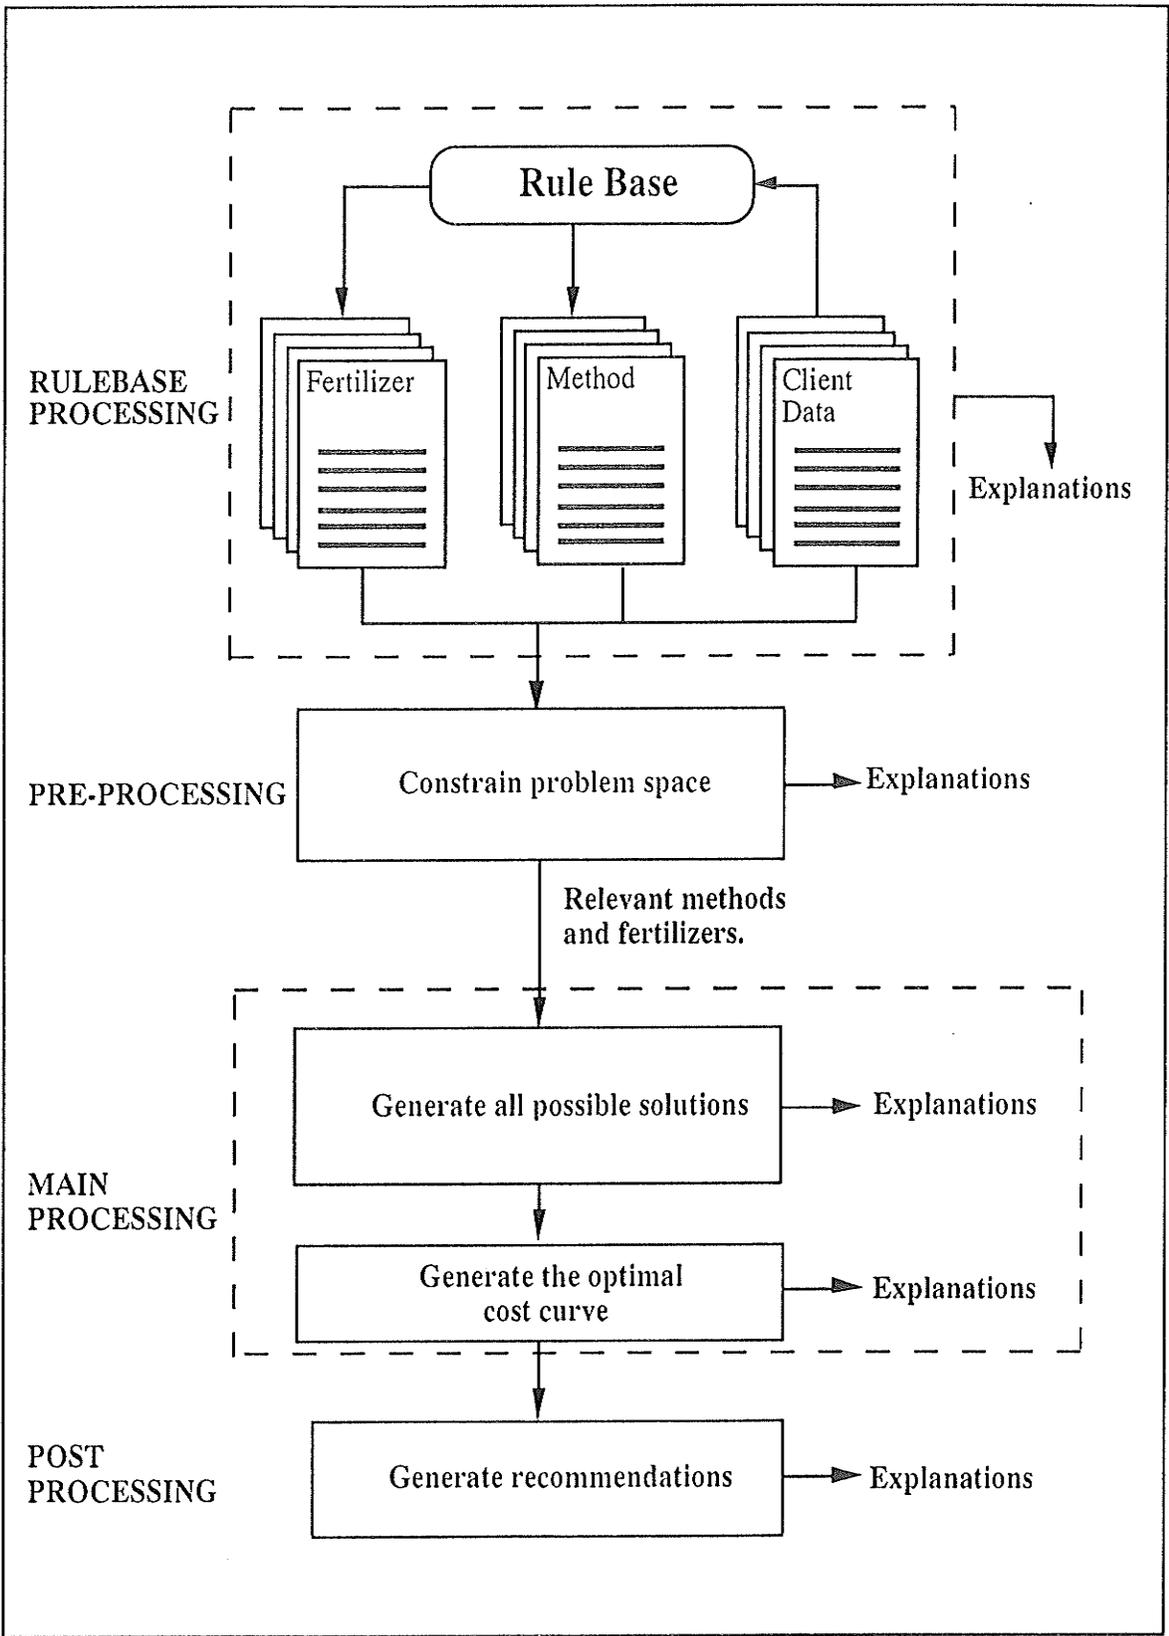


Figure 2.7: Processing component stages.

Rule 1: Tightening Constraints

IF ((fertilizer = Urea OR UAN Liquid)
AND
(soil moisture condition = Dry OR soil texture = Coarse))

THEN fertilizer's nitrogen banded toxicity limit is *lowered* to 0.0

BECAUSE such soil conditions prevent Urea or UAN Liquid from being applied in the seed row

Rule 2: Relaxing Constraint

IF ((fertilizer = Urea OR UAN Liquid OR Ammonium Nitrate)
AND
(soil moisture condition = Moist)
AND
(soil texture = Medium OR Fine)
AND
(seeding implement = Discer OR Scatter Airseeder
OR Sidebanding Drill))

THEN *double* fertilizer's nitrogen banded toxicity limit

BECAUSE such soil conditions make the fertilizer less toxic to seeds

Figure 2.8: Rule base processing: Use rulebase to check farm's soil conditions and adjust toxicity constraints of fertilizers accordingly by applying the rules to each fertilizer.

2.2.2 Pre-Processing

The pre-processing stage constrains the problem space to a manageable size. Initially, only *valid* application method and fertilizer combinations are formed. A valid combination is a match between the fertilizer's available forms (*i.e.*, granular, liquid, or gas) and methods which apply fertilizers in those forms. For example, a broadcast method can only apply granular or liquid fertilizers. This *form-compatibility* domain-specific constraint helps in reducing the number of method-fertilizer combinations significantly. At this point, additional heuristics are used to further trim the number of combinations. For example, if there is more than one fertilizer that can be applied using a particular broadcast method,

consider only the cheapest fertilizer under that method. Recall the linear cost curve associated with any fertilizer being applied using a non-seed row method (broadcasting is a non-seed row method); fertilizers with a higher per unit cost than the cheapest fertilizer will result in a steeper slope and thus will never be selected as an optimum cost curve. For seed-row methods, consider the fertilizer with the cheapest unit price and also other fertilizers (which may be more expensive) with higher tolerance to toxicity problems under the same seed-row method. Higher tolerance to toxicity means more of the fertilizer can be applied using the seed row method before having to switch over to a non-seed row method. We want to pursue the effect of adding as much of a fertilizer as possible using a seed row method because this type of application does not incur an extra application cost. Besides these built-in heuristics which prune out subsets of possibilities that incur high costs, the user may manually override the system; for example, to consider only a subset of the given application methods (possibly due to lack of available machinery) or to prune particular fertilizers because it is the user's traditional practice to avoid using those fertilizers. This pre-processing stage plays an important role in the success of the system because it utilizes domain knowledge to reduce the problem space which in turn reduces the search necessary in the next processing phase.

2.2.3 Main Processing

The main processing phase consists of three steps. The first step involves generating the (reduced) problem space of all possible potential solutions. For each valid method/fertilizer combination (or methods/fertilizer(s) combination in the case of partial solutions) calculate the cost of achieving each increment of nitrogen supply, up to the maximum nitrogen supply. The result is a solution matrix where each row i is method/fertilizer combination i and each column j in a row is the cost for using method/fertilizer combination i to supply j increments of nitrogen to the soil.

The second step sorts the columns of the solution matrix such that each increment of nitrogen supply will be a column containing solutions in ascending order by cost. An optimization technique is applied in this sorting step. It is noticed that by adding each increment of nitrogen supply, it is unlikely that the cost order of the solutions will change drastically. This is apparent in Figure 2.5 where the cheapest solution at the start continues to be the cheapest solution for a few increments until its toxicity constraint is exceeded. It is no longer the cheapest when a second method/fertilizer combination is incurred as part of its solution to supply the required level of nitrogen. Then another method/fertilizer solution becomes the cheapest and continues again for a few increments until its toxicity constraint is exceeded. There are two other factors, namely relative application method efficiencies and relative fertilizer compound costs which may cause changes in the cost order when additional increments of nitrogen supply is added. These effects are illustrated in Figure 2.9. In Figure 2.9(a), the solid cost curve represents solution *A*, the broadcasting of Urea and the dashed cost curve represents solution *B*, the banding of Urea. Although the cost of broadcasting Urea is cheaper than banding Urea in the first few increments, the less efficient method of broadcasting (relative to banding) becomes a less economical method for applying high levels of nitrogen to the soil. As the target nitrogen level increases, the amount of Urea required for achieving the same target is considerably more for broadcasting than is for banding. This is because broadcasting is a less efficient method of delivering the nitrogen content of the fertilizer to the seeds in the soil. In Figure 2.9(b), the solid cost curve represents solution *A*, the broadcasting of a high per unit cost fertilizer compound *A* and the dashed cost curve represents solution *B*, the banding of a low per unit cost fertilizer *B*. *A* and *B* are different fertilizer compounds. Again, the cost of broadcasting is cheaper than banding thus solution *A* is the cheaper solution in the first few increments. However, for higher amounts of nitrogen supply, the high cost of fertilizer *A* erodes the cost advantage of broadcasting over banding. Solution *B* then becomes the cheaper solution and remains the cheaper solution. Knowing that the cost of solutions will remain in order

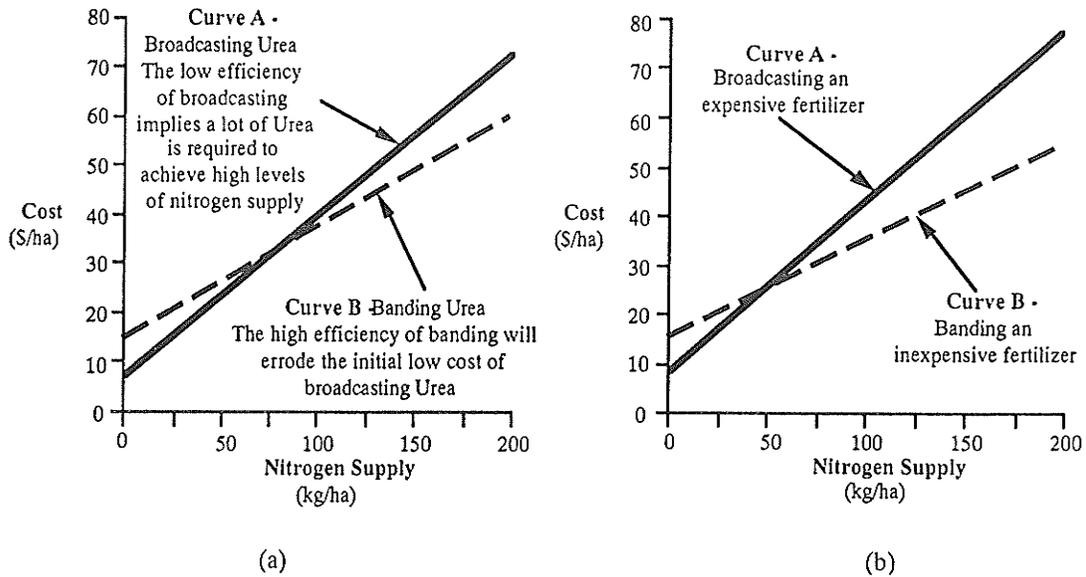


Figure 2.9: The effects of relative application efficiencies and relative fertilizer compound costs on the cost order of solution curves.

for intervals at a time suggested that a bubble sort is an efficient sorting algorithm for this step of processing. The bubble sort is efficient for sorting arrays which are almost sorted and is applied to the matrix by using the order of the solutions from the last increment as the starting order for the current solutions and bubble sorting the current solutions based on the costs of the current solutions. This sorting step prepares the solution matrix for the third step of main processing which performs the actual search of the problem space. Having the solutions sorted by cost will eliminate the need to search for the lowest costing solution at each increment because the most economical solutions are at the top of the matrix.

The third step simply searches along the increments of nitrogen supply to find the optimum return. This is done by taking the lowest costing method/fertilizer combination at each increment and subtracting it from the expected revenue generated by the potential yield associated with that nitrogen level. Then search for the increment with the highest net return on fertilizer investment.

The entire process of finding the optimum fertilization plan can be described as a *constraint-directed search* [Fox, 1987]. Constraints inherent in the domain constrain the search space and reduce the necessary search time [Newell et al., 1976]. An example of a constraint that is often encountered during the search process is the fertilizer toxicity limit. When considering the possibility of adding a fertilizer to the seed row, toxicity constraints govern the amount of fertilizer that can be added before damage to the seed occurs. The search algorithm must recognize these toxicity constraints and not surpass them when recommending the amount of fertilizer to apply. The process of performing a constraint-directed search has been split into the steps of initially generating the cost curves which ensured that constraints are not violated and then sorting them for each increment of nitrogen supply. Consequently, the system searches along the curves as illustrated in Figure 2.3, considering the lowest cost recommendation at each 5 kg/ha increment of N and also locating the point of greatest distance between the revenue and the cost curve (*i.e.*, the optimum solution).

2.2.4 Post-Processing

During post-processing, crop-specific details are looked after. For example, notes are attached to recommendations to warn the user of special procedures to be followed for the fertilization of such a crop.

The resulting recommendations (both optimum and less than optimum ones) form a summary list of potential solutions. The list is presented to the user with the most cost effective alternative being highlighted. The summary lists for each 5 kg/ha increment of N, the method (or combination of methods) that should be used, the fertilizer (or fertilizers) that should be applied under each method, the cost per hectare of applying the fertilizer, the additional revenue generated as a result of applying that fertilizer, and the forecasted profit margin.

The next step is to scale the system up to consider multiple nutrients. After all, nitrogen's effectiveness towards proper crop growth is maximized only when other nutrient requirements are fulfilled. An intuitive strategy would be to generate recommendations for each of the other nutrients also on the assumption that there are sufficient levels of the remaining nutrients available in the soil to the seed, and then bundling all four nutrients' recommendation into a final recommendation. In the next chapter, we will discover that there are interactions between the nutrients which will complicate the task of combining single nutrient recommendations.

Chapter 3

Multiple Nutrients

Common fertilizer compounds available in Manitoba contain some composition of the four nutrients Nitrogen (N), Phosphorus (P), Potassium (K), and Sulphur (S). Single nutrient recommendations will supplement a deficiency of any one nutrient in a farm's soil assuming that other nutrient requirements are already fulfilled. However, this type of recommendation is not sufficient for most farms. For fertilizer problems where the soil is deficient in more than one (or possibly all four) of the nutrients, an expert would have to analyze the soil's needs for each individual nutrient and arrive at some independent recommendation for each nutrient. Then the expert would resolve any conflicts between the individual recommendations and integrate the recommendations to arrive at the final recommendation. This result must meet the objective of having maximum cost efficiency, while maintaining a proper balance of all nutrients and also must satisfy any constraints imposed on the application of mixtures of fertilizer products. In this chapter, I will present a framework in which the fertilizer problem can be solved. The solution will ensure that the objectives mentioned above are met.

3.1 General Multiple Nutrient Fertilizer Problem

Consider an approach which takes advantage of the fact that cost curves and recommendations have been generated for the nitrogen-only problem, assuming sufficient levels of the other nutrients are already present in the soil. To truly complete these recommendations, we must make fertilization plans for the other nutrients to satisfy the above assumption. Suppose that there are four specialists (human or computational entities), each looking after a specific nutrient requirement. The general processing scheme is to first perform the rule base processing (as described in section 2.2.1) on the application methods and fertilizer compounds data to customize these to reflect farm-specific conditions. Then request the nitrogen specialist to make its recommendations. The processing performed by the nitrogen specialist will be identical to that described in the previous chapter. Then these recommendations are passed along to the phosphorus specialist, then the potassium specialist and then the sulphur specialist for each of their contributions. Each will contribute alternative single nutrient fertilization plans to meet the nutrient requirements necessary to complement the effectiveness of the nitrogen recommendations. In this scheme the global final recommendations are constructed incrementally. Let us examine the processing to be performed by the phosphorus, potassium, and sulphur specialists.

Each specialist is given a set of recommendations which is assumed to meet all other nutrient requirements except for the nutrient which the specialist is going to recommend. The contributions made by a specialist must not alter the other recommended nutrient levels. However, recall that fertilizers come in a variety of compositions of the four nutrients. Each fertilizer can be categorized by the major nutrient (*i.e.*, the nutrient which takes up the largest percentage of the overall composition) that it contains. Of the common fertilizer compounds currently being considered by the system, the common nitrogen fertilizer compounds have the characteristic that they contain only nitrogen and none of the

other nutrients. The phosphorus fertilizer compounds contain primarily phosphorus along with a small percentage of nitrogen. The sulphur fertilizer compounds have the same characteristic that they contain some nitrogen besides the major nutrient sulphur. Potassium fertilizer compounds contain only potassium. When a specialist focuses on one nutrient, the selected fertilizer compound(s) may contribute other nutrients as a by-product. When a specialist attempts to integrate these fertilizer compounds into the given set of recommendations, the levels of the other nutrients may be increased. This violates the initial restriction of a specialist not being allowed to alter the other recommended nutrient levels. If a specialist is to solve this problem involving other nutrients, then the specialist must have knowledge about other nutrients. This complicates the scope of each specialist and is not desirable.

In addition to this problem of a specialist having to handle fertilizer compounds which contribute other nutrients outside of its domain of expertise, each specialist must make sure that constraints *local* to its domain of expertise on the nutrient must not be violated. Concurrently, the decisions made must also adhere to the constraints *global* to the overall fertilizer problem. For example, a local constraint for the phosphorus specialist is that broadcasting a phosphorus fertilizer is inefficient and should not be considered. A global constraint for any specialist is that the total amount of all fertilizer product applied with the seed in the seed row should not exceed 195 kg/ha. Another global constraint is that the final recommendation for fulfilling all four nutrients' requirements should not involve more than two application methods. It is preferable to satisfy all nutrient requirements with as few application methods as possible. Having to deal with more than two different application methods is very inconvenient for any farmer. This is an example of a domain heuristic that our experts learned from field experience.

The goal of this chapter is to present an organization of specialists, each with a narrow

scope of the entire fertilizer problem. The specialists will incrementally make contributions, gradually leading to a complete solution. This organization is known as a **blackboard system**. The rest of this chapter will describe the general blackboard architecture and organization, the blackboard system in FA, and compare and contrast the general blackboard characteristics to those in FA. The discussion will include the solutions to those problems mentioned in the above paragraphs and also other problems and their solutions in the multiple nutrients domain.

3.2 Overview of General Blackboard Systems

Conceptually, a general blackboard system consists of three major components: the knowledge sources, the blackboard, and the control. The **Knowledge Sources** are self-contained units, each having specialized problem solving knowledge for partially solving the problem. There is no restriction on the method by which a knowledge source solves a problem. Each knowledge source may apply a paradigm completely different from the others. For example, one knowledge source may use some nested procedures to solve its assigned problem while another knowledge source may apply a set of rules to produce its results. The **Blackboard** is a global structure on which knowledge sources contribute information to incrementally solve the problem. The **Control** component coordinates the activities of the knowledge sources and opportunistically executes appropriate knowledge sources to provide partial solutions to the problem. This will eventually lead to the problem being solved.

The problem solving mechanism lies in the sharing of information between the knowledge sources via the blackboard. Since each knowledge source is a specialist in a section of the problem domain, the pooling of their individual efforts on solving portions of the problem will eventually lead to a solution to the overall problem. There is no direct communication between the individual knowledge sources. They communicate through the blackboard.

Their activities on the blackboard are overseen by the control component. The control component can be described as having two parts: the control procedure and the control data. The control procedure takes action in determining which knowledge source(s) are most appropriate for execution and opportunistically activates them to make their contributions. The control procedure's decisions regarding the selection of knowledge source(s) for execution is based on information kept in the control data. This information may be in the form of an agenda which is updated constantly to reflect current blackboard status and pending knowledge sources to be executed. There may also be domain-specific heuristics kept in the control data which are useful in the problem solving process. Centralizing the control into one component of the system ensures that the coordinated efforts of the knowledge sources are being directed towards the desired solution.

Having individual and independent knowledge sources implies a potential for parallel execution of knowledge sources as long as there is sufficient information on the blackboard for each knowledge source to successfully perform its work. In some cases, more than one knowledge source can work in parallel on solving different parts of the problem. At other times, a knowledge source must wait for the availability of specific information on the blackboard before it is ready to make its contributions. The required information may be the results produced by another knowledge source. This implies that some knowledge sources execute in serial order.

The information contributed to the blackboard by various knowledge sources can usually be arranged into a hierarchy. That is, solutions may be made to lower level subproblems, and these lower level solutions can be used in deriving solutions to subproblems at a higher level of abstraction. For example, in the HEARSAY-II Speech Understanding System [Erman et al., 1980], syllables are interpreted by one knowledge source at the syllabic level and are grouped together by another knowledge source at the lexical level to form words

which are in turn used by a phrasal level knowledge source for interpretation as phrases. "Generally, a knowledge source uses information on one level as its input and produces output information on another level." [Nii, 1986a].

Consider an arrangement of the blackboard into levels of abstraction where data produced at the different levels form an hierarchy. The hierarchy has raw input data at the lowest level and through gradual refinement up the hierarchy, one or more solutions to the problem are produced at the top of the hierarchy (see Figure 3.1). The knowledge sources operate between data abstraction levels and are responsible for refining data from one level and producing data for a higher level. This blackboard architecture has a number of advantages. The notion of having knowledge sources which independently solve parts of the entire problem at various levels of abstraction lends well to modular design. The entire problem and the knowledge to solve the problem is divided into subproblems which can be packaged into knowledge sources. Each knowledge source is in essence a self-contained module, having the knowledge it requires for doing its job. The fact that the knowledge sources are independent of one another allows for incremental system development between multiple system builders. The entire problem can be broken down into smaller subproblems and knowledge sources can be built to solve each subproblem; analogous to modular program design. In addition to being advantageous in reducing the complexity of the design, this divide-and-conquer approach is also effective in constraining the amount of search required in problems which have large problem spaces. The modular structuring also makes the system knowledge easier to maintain.

3.3 The Fertilizer Advisor Blackboard

The organizational structure of FA follows a blackboard model as depicted in Figure 3.2. The specialists (another name for knowledge sources) in FA are procedures and functions (or nested procedures and/or functions) each designed to solve a subproblem in the

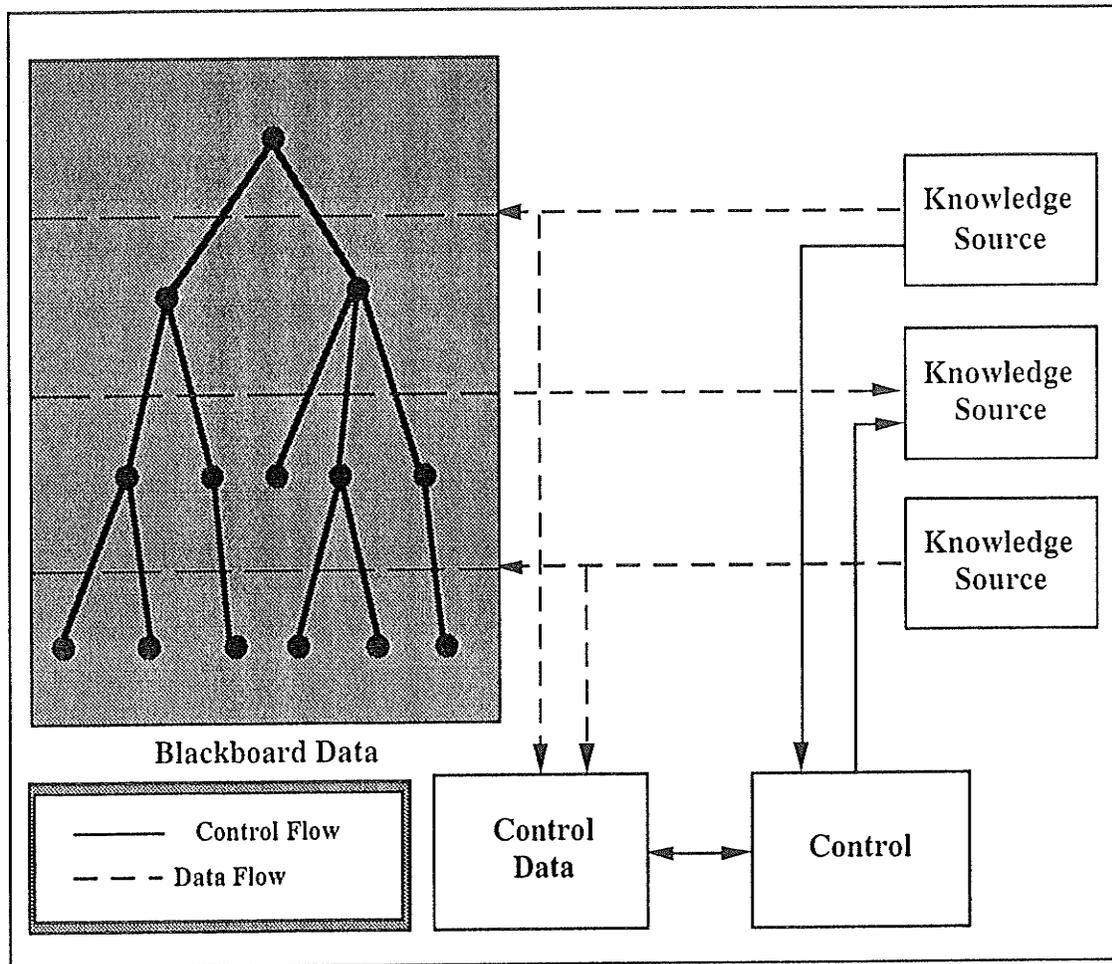


Figure 3.1: Blackboard Architecture. Adapted from [Nii, 1986a]

domain. The controller in FA is the program mainline which contains the agenda of specialists to call. A data structure called SITUATION is used to represent the blackboard (see Figure 3.3). Specialists take information from this blackboard and also make contributions to it. The information about different application methods is stored in an array of METHOD structures (see Figure 3.4). Likewise, the information about different fertilizer compounds is stored in an array of FERTILIZER structures (see Figure 3.5). Information regarding the client's field being fertilized is stored in a CLIENT_DATA structure (see Figure 3.6). Nitrogen yield tables are stored as columns of a numeric matrix and phosphorus, potassium, and sulphur recommendation tables are stored in numeric arrays. The FERTILIZER, METHOD, and CLIENT_DATA, YIELD_TABLE,

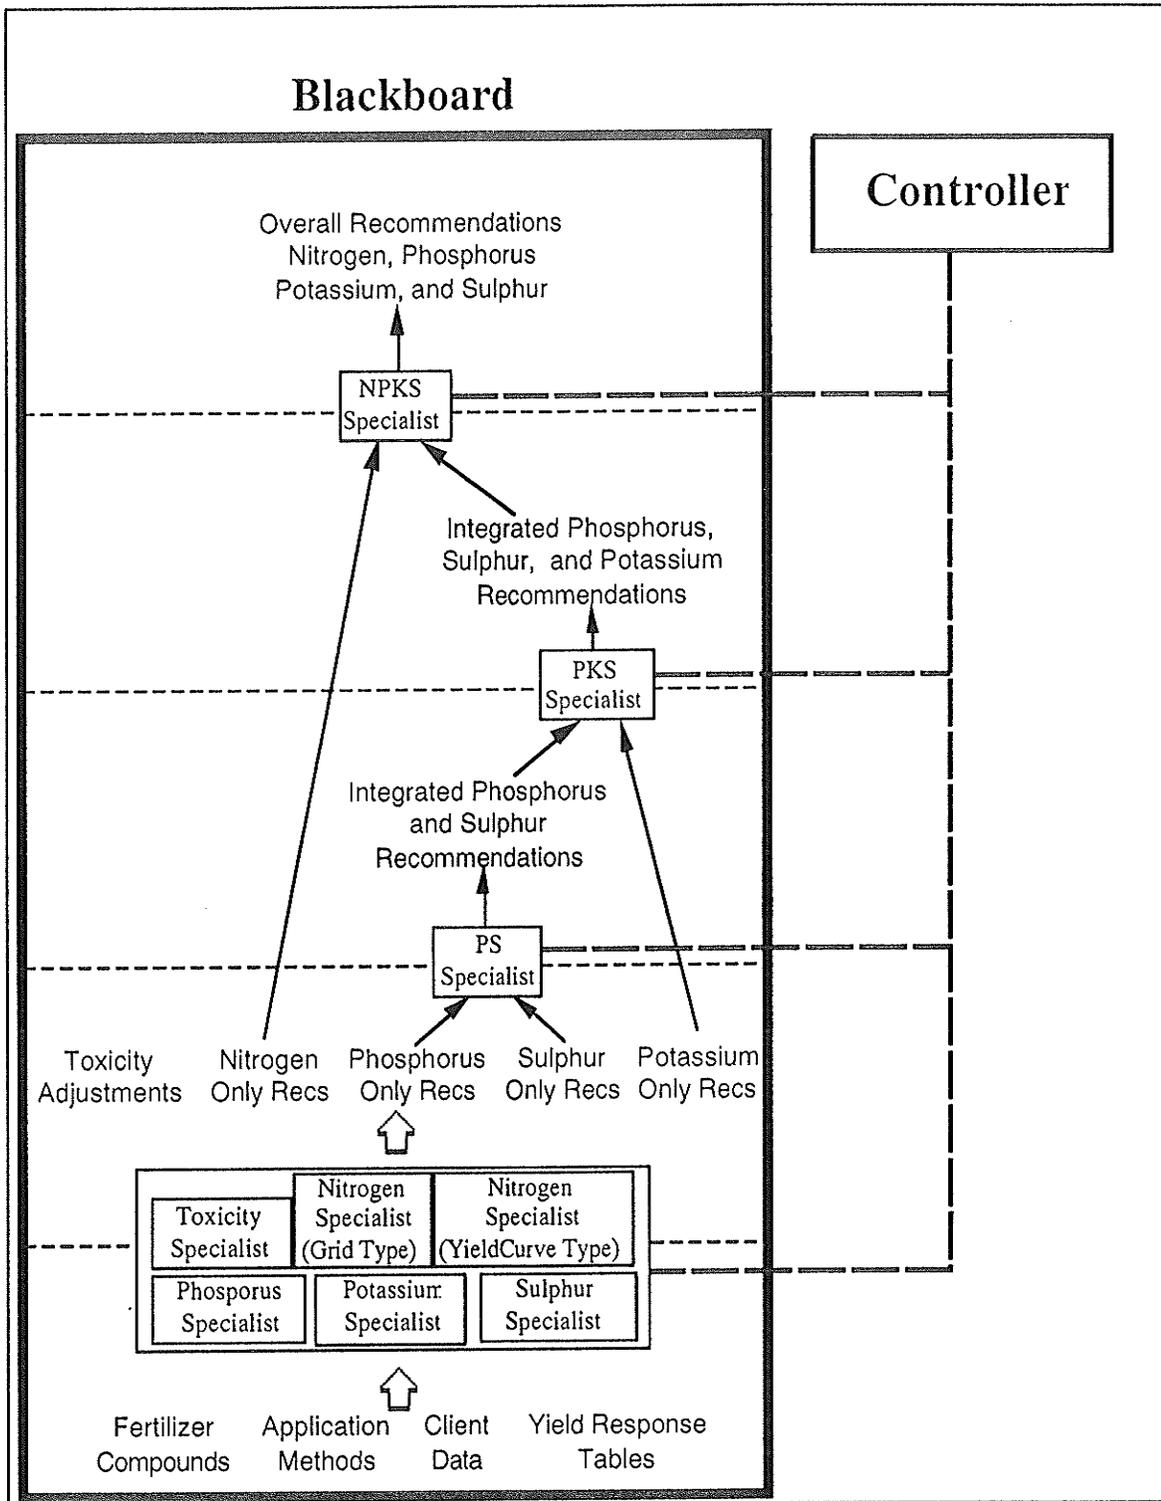


Figure 3.2: The Blackboard Architecture used in the Fertilizer Advisor. This shows one instance where the PS specialist is called upon to integrate the phosphorus and sulphur recommendations. There may be situations under which the requirements of the potassium must be given priority. In such a case, a PK specialist or a SK specialist will be executed depending on the priority between the phosphorus and the sulphur requirements.

```

struct SITUATION {
    struct CLIENT_DATA    client;
    struct FERTILIZER     ferts[MAX_FERTILIZERS];
    struct METHOD          methods[N_METHODS];
    int    num_solns;
    struct SOLUTION solns[MAX_NUM_SOLNS];          /* Cost Curves */
    int    *sorted_solns;
    int    num_sr_n;
    int    sorted_sr[MAX_N_FERTILIZERS];          /* Sorted SR only */
    int    num_nonsr_n;
    int    *sorted_non_sr; /* Sorted Non-SR for each increment */
    struct N_RECOMMEND    n_rec[MAX_N];
    struct P_RECOMMEND    p_rec[MAX_P];
    struct S_RECOMMEND    s_rec[MAX_S];
    struct K_RECOMMEND    k_rec[MAX_K];
    int    pks_size;          /* Size of PKS_RECOMMEND list */
    struct PKS_RECOMMEND  *pks_list;
    struct NPKS_RECOMMEND npks_rec[MAX_NUM_INCR + 1];
    char    *client_expl; /* Explanation of client's situation */
};

```

Figure 3.3: SITUATION structure

```

#define METHOD_NAMESIZE 30
struct METHOD { /* application methods */
    int    id;          /* method id */
    char    name[METHOD_NAMESIZE]; /* name of the method */
    double cost;          /* application cost per ha */
    double efficiency;    /* w.r.t. spring broadcast */
    struct FERT_FORMS    forms; /* applicable forms of fert */
};

```

Figure 3.4: The METHOD structure for storing data about one application method.

P_RESPONSE, K_RESPONSE, and S_RESPONSE structures are all storage areas for input data. There are specialists which display appropriate input screens to request values for some data fields in each of these structures. Other data fields in the structure are areas for sharing information between specialists.

3.3.1 Rule Base Processing Specialist

The rule base processing described in the previous chapter is the responsibility assigned to the toxicity specialist. The crop field conditions, as described by the data stored in

```

struct FERT_FORMS { /* forms of fertilizer available */
    BOOLEAN gas;          /* available in gas form? */
    BOOLEAN liquid;      /* available in liquid form? */
    BOOLEAN granular;    /* available in granular form? */
};

#define FERT_NAMESIZE 20

struct FERTILIZER { /* Information about a fertilizer */
    int    id;           /* fertilizer id */
    char   name[FERT_NAMESIZE]; /* a name for the fertilizer */
    double s_price_per_kg; /* price per kilogram */
    double s_price_per_ton; /* price per 1000 kilo's */
    double s_price_per_kg_of_n; /* price of a kg of nitro */
    double f_price_per_kg; /* price per kilogram */
    double f_price_per_ton; /* price per 1000 kilo's */
    double f_price_per_kg_of_n; /* price of a kg of nitro */
    double nitro;        /* nitrogen content */
    double phos;         /* phosphorus content */
    double sulf;         /* sulfur content */
    double potas;        /* potassium content */
    struct FERT_FORMS forms; /* forms available */
    double banded_toxicity; /* maximum kg/ha of N which can */
                          /* safely be applied with the seed */
};

```

Figure 3.5: The FERTILIZER structure for storing data about one fertilizer compound.

```

struct CLIENT_DATA { /* Client's farm information */
    int    client_id;    /* client id */
    int    field_id;    /* field id */
    int    sit_id;      /* situation id */
    char   client_name[CLIENT_NAMESIZE];
    char   field_name[FIELD_NAMESIZE];
    char   sit_name[SIT_NAMESIZE];
    char   legal_desc[LEGAL_DESCSIZE];
    int    moist_cond;  /* field moisture condition */
    int    yldtbl_choice; /* yield table to use */
    int    texture;     /* soil texture */
    int    ph;          /* pH level of the soil */
    int    soil_n_level; /* level of nitrate in soil */
    int    soil_p_level; /* level of phosphate in soil */
    int    soil_k_level; /* level of potassium in soil */
    int    soil_s_level; /* level of sulphur in soil */
    int    seed_impl;   /* seeding implement */
    int    seed_impl_qual; /* a qualification of seed. imp. */
    int    crop_grown;  /* the crop to be grown */
    double crop_price;  /* potential sale price of crop */
    BOOLEAN adjustments_made; /* adjs made to the fert data? */
};

```

Figure 3.6: The CLIENT_DATA structure for storing data about a client's field.

CLIENT_DATA, can have an effect on the toxicity behaviour of the fertilizers being analyzed; the FERTILIZER data must be customized to reflect this difference in behaviour. The toxicity specialist has the knowledge about fertilizer toxicity behaviours expressed as a set of rules. Upon data being entered into the CLIENT_DATA, this specialist applies the rules in its rule base against the data and modifies the toxicities in the FERTILIZER accordingly. A rule has a premise and an action corresponding to the IF clause and the THEN clause of a rule like that shown in Figure 2.8. The rule base processing is performed in a forward chaining fashion by cycling through the toxicity adjustment rules to check if any of the rule have its premise(s) match the information in the CLIENT_DATA. If so, that rule is fired. Firing a rule will result in the rule action being performed. In this case, the action is the adjustment of one or more fertilizers' toxicity limit values. The rule base processing stops when no more rules can be fired.

3.3.2 Single Nutrient Recommendation Specialists

To generate the optimum recommendation for the given information, the first specialists to be executed are the single nutrient specialists. Based on the information given in the input data, each single nutrient specialist produces a list of alternative single nutrient recommendations. Generally, to ensure that the crop is properly fertilized, specific concentrations of each nutrient must be made available in the soil to the crop. There are tables distributed by the Manitoba Soil Testing Laboratory which recommend, for any single nutrient, the appropriate amount of that nutrient to add based on the level of that nutrient currently present in the soil. These tables have been compiled by experts and are based on research and statistical analysis of historic data. Note that each single nutrient recommendation is given based on the assumption that either there are satisfactory levels of the other nutrients already in the soil or the concentrations of the other nutrients will be brought to the appropriate levels. Thus each specialist focuses on the nutrient it is responsible for and does not have to concern itself with the requirements of the other

nutrients. Basically, a single nutrient specialist is given the level of the nutrient presently in the soil (results from a soil test analysis) and performs a table look up for the recommended amount of nutrient to apply. The relevant tables are statically stored in memory and are part of a specialist's knowledge about the particular nutrient. Each table simply indicates the rate of the nutrient to apply and does not indicate what fertilizers to apply and how to apply the fertilizers to achieve the recommended levels. Each specialist extends the function of these tables by proposing a list of alternative means for achieving the recommended rate of nutrient application. Each entry in the list is a valid fertilization plan composed of a fertilizer source along with the rate, method, timing, and cost of application for achieving the recommended nutrient supply.

Generally, to derive each entry, a fertilizer source, and the method and timing of application must be chosen. The compatibility between a fertilizer source and its method of application is constrained by the form in which the fertilizer is available. For example, fertilizers in the gas form can not be broadcasted. Combinations which are incompatible in terms of fertilizer form are discarded. For each form-compatible recommendation, the rate and cost of application (including fertilizer product cost and cost of application) is calculated. The recommendations are listed in ascending order by cost. This is the basic process in generating the lists of single nutrient recommendations. Other nutrient specific constraints may apply and are discussed below.

3.3.2.1 Nitrogen-Only Specialists

Recall from chapter three that nitrogen requirement is handled differently depending on the crop type. The grid-type crops each have recommendation tables which indicate the appropriate level of nitrogen to apply given the level tested to be currently in the soil. As for crops which have corresponding yield tables describing their yield potential relative to the level of nitrogen supplied to the soil, we can perform a more in-depth analysis of the

economic feasibility of various recommendations. There are two nitrogen specialists implemented in FA, each managing a different type of recommendation. Basically, the yield curve type nitrogen specialist handles the crops which there is a corresponding expected crop yield for a particular amount of nitrogen applied. The grid type nitrogen specialist handles the crops for which there is no way to predict the yield potential. The controller activates the appropriate one depending on the type of crop being grown by the client as indicated by the `crop_grown` data field in `CLIENT_DATA`. The processes for making these two types of recommendations have been described in the previous chapter. Both specialists begin by generating a list of all valid combinations of method/fertilizer pairs and their corresponding cost curves. The procedure for generating the cost curves ensures that the local constraint of not exceeding individual seed row fertilizer toxicity limits and the global constraint of not applying more than 195 kg/ha of fertilizer product in the seed row are not violated. The list is stored in an array of `SOLUTION` structures (see Figure 3.7) within `SITUATION`. Each `SOLUTION` structure in the array represents the fertilizer(s) and method(s) components of a fertilization plan. The *partial* field is set to `TRUE` if the solution being represented is partial; it is set to `FALSE` for complete solutions. The cost array within a `SOLUTION` structure contains the costs for achieving each 5 kg/ha increment of nitrogen supply using that solution. The next step is to sort the solutions in ascending order by cost for each increment of nitrogen supply. The sorting can not be performed on the array of `SOLUTION` structures because each `SOLUTION` has its own cost array embedded within the structure. Instead, we want to create matrix where each column represents a particular increment of nitrogen supply (*i.e.*, the first column represents supplying 5kg/ha of nitrogen, the second column represents supplying 10kg/ha of nitrogen and so on). Each row in the column is a solution for achieving the represented nitrogen supply. In a given nitrogen supply column, the rows of solutions are sorted according to cost in ascending order. To conserve memory usage, this matrix is coded as an integer matrix where the value I stored at position (R,C) in the matrix represents that the Rth lowest costing solution

```

struct SOLUTION {
    BOOLEAN partial;          /*Solution partial or complete?*/
    int    meth_id1;         /*Method id of 1st component*/
    int    fert_id1;        /*Fertilizer id of 1st component*/
    int    meth_id2;        /*Method id of 2nd component*/
    int    fert_id2;        /*Fertilizer id of 2nd component*/
    int    start_of_comp2;   /*Sub-script of start of 2nd*/
                                /*component if partial solution*/
    double cost[MAX_NUM_INCR]; /*Array of costs for applying*/
                                /*each 5 kg/ha increment of N*/
};

```

Figure 3.7: The SOLUTION structure. In FA, there is an array of this structure each representing a unique nitrogen fertilizer/method combination.

for achieving $((C+1)*5)$ kg/ha of nitrogen supply is the fertilizer(s) and method(s) given in the I th element of the SOLUTION structures array¹. Furthermore, the C th element of the cost array in that I th SOLUTION structure contains the cost of performing the fertilization. This matrix is stored in the block of memory pointed to by the **sorted_solns* pointer field in SITUATION.

The processing steps leading up to the generation of the SOLUTION list are common to both specialists, although only one specialist will be activated depending on the type of crop being grown. The SOLUTION list is the only essential nitrogen-only contribution which will be subsequently used by the NPKS recommendation specialist for the production of the final overall recommendations. However, the currently active nitrogen specialist continues on to generate nitrogen-only recommendations. The processing for each nitrogen specialist is as follows. The yield curve type nitrogen specialist must form a recommendations list consisting of the fertilization plan with the best net return at each increment of nitrogen supply. Recall that $net\ return = revenue$ generated from crop yield - fertilization *cost*. The revenue generated is constant for a given nitrogen supply. This

¹Note that in the C language, the first element of an array has the subscript zero. Thus to represent that the first column supplies 5 kg/ha of nitrogen, we must add one to the subscript before multiplying by five.

```

/* This N_RECOMMEND structure is used to hold recommendations for both
yield curve-type recommendations and grid-type recommendations. The
structure contains fields which are used in either or both types of
recommendations. The fields are grouped together according to whether
they are used in both types of recommendations or just specific to
one type of recommendation. */
struct N_RECOMMEND {
/* The following fields are commonly used by both type of recs. */
int      meth_id1; /* method id of 1st component */
int      fert_id1; /* fertilizer id of 1st component */
double   fert_amt1; /* amount of 1st component fert to apply*/
int      meth_id2; /* method id of 2nd component */
int      fert_id2; /* fertilizer id of 2nd component */
double   fert_amt2; /* amount of 2nd component fert to apply*/
double   total_cost; /* total cost of this solution */

/* The following fields are used by yield curve type recs only. */
int      n_level; /* current nitrogen level */
int      yield; /* current yield (kg/ha) */
double   gross; /* gross return */
struct   SOLUTION *soln_ptr; /* Pointer to the solution */
/* which has the lowest cost at the */
/* current n_level. */
/* The following fields are used by grid type recs only. */
char     *explain; /* The reason we picked this combination*/
};

```

Figure 3.8: The N_RECOMMEND structure.

implies that at any given nitrogen supply, the net return and the fertilization cost are inversely related. It follows that the cheapest fertilization plan is also the plan which will achieve the best net return at that nitrogen supply. Thus, the yield curve type specialist simply takes the first fertilization plan (*i.e.*, the cheapest costing plan or the best return) represented by the first row of the list pointed to by *sorted_soln to form the required list and identifies the nitrogen supply which gives the overall optimum net return. The grid type nitrogen specialist determines the amount of nitrogen to supply based on the Manitoba Soil Testing Laboratory Nitrogen Recommendation Tables and takes from the matrix pointed to by *sorted_soln the column representing the list of alternative fertilization plans for achieving the recommended increment of nitrogen supply. The list of alternative nitrogen only recommendations (regardless of type) is stored in the array of N_RECOMMEND structures (see Figure 3.8) within SITUATION. A summary of the processing for generating the two types of nitrogen-only recommendations is described in Figure 3.9.

Form all valid fertilizer source and application method combination

For each combination, generate the cost curves for applying each 5kg/ha increment of N

Sort the solutions for each increment of nitrogen by cost

For *yield type* recommendations, display the top solution for each increment of nitrogen nitrogen level

For *grid type* recommendations, determine the recommended level of nitrogen to apply as suggested by tables from the Manitoba Soil Testing Laboratory, and display the top solution for the recommended nitrogen level

Figure 3.9: Generating N-only recommendations.

3.3.2.2 Phosphorus-Only Specialist

Phosphorus-only recommendations are derived in the same way as the grid type nitrogen-only recommendations. The Manitoba Soil Testing Laboratory provides phosphorus recommendation tables for a large variety of crops grown in Manitoba. The soil phosphate content as determined from a soil test analysis is used to look up the appropriate level of phosphorus to apply. The phosphorus-only specialist builds a list of alternative fertilization plans where each plan meets this phosphorus requirement. These fertilization plans all meet the local and global constraints imposed. For example, the form in which the suggested fertilizer compound is available must be applicable using the corresponding recommended application method in each fertilization plan. The total amount of fertilizer product to apply in the seed row to meet the phosphorus requirement must not exceed 195 kg/ha. These are global constraints which the phosphorus-only specialist must conform to when generating its recommendations. There are two local constraints for the phosphorus-only specialist. Phosphorus should not be broadcast since the amount of broadcast phosphorus used during the first number of years has to be increased two or fourfold to equal yield increases from placement near the seed. Therefore, broadcast application of phosphorus may be uneconomical on many soils. Also, the amount of nitrate nitrogen given off as a by-product of the phosphorus fertilizer must not exceed 45 kg/ha when applied in the seed row. The

```

struct P_RECOMMEND {
    int    meth_id;    /* Id of the method being used */
    int    fert_id;    /* Id of the fertilizer being used */
    double fert_amount; /* Amount of fertilizer to apply */
    double n_amount;   /* Amount of N by-product produced */
    double total_cost; /* The cost of applying the P */
    char   *explain;   /* Reason we picked this combination */
    BOOLEAN keepit;    /* Indicates if we should keep this */
                /* combination or forget it due to */
                /* violation of toxicity constraints */
};

```

Figure 3.10: The P_RECOMMEND structure.

Determine the recommended level of phosphorus to apply using phosphorus recommendation tables from the Manitoba Soil Testing Laboratory

Form all valid phosphorus fertilizer source and application method combinations

For each combination, determine the amount of fertilizer to apply to achieve the recommended phosphorus level and determine the cost of fertilization

Display all phosphorus recommendations in ascending order sorted by cost

Figure 3.11: Generating P-only recommendations.

recommendations produced by the phosphorus-only specialist are stored in the array of P_RECOMMEND structures (see Figure 3.10) within SITUATION. Also, a summary of the phosphorus recommendations generation process is described in Figure 3.11.

3.3.2.3 Potassium-Only Specialist and Sulphur-Only Specialist

Potassium-only recommendations and sulphur-only recommendations follows the same procedure as the building of the phosphorus-only recommendations. The Manitoba Soil Testing Laboratory has tables which give recommendations for these two nutrients based on their concentration in the soil as indicated by a soil test analysis. These specialists must enforce the same global constraints as the phosphorus-only specialist in building its fertilization plans. It is the local constraints which differ. For potassium, if a fertilization

plan recommends using a broadcast application method, then the suggested rate of fertilizer to apply must be twice as much as that if using a non-broadcast application method to meet the same potassium requirement. For sulphur, special constraints are in effect for its use in elemental form. It is inefficient to apply elemental sulphur because it is not in a form which is immediately available to plants. This results in twice the rate of application for applying elemental sulphur to meet the required sulphur level. The recommendations produced by the potassium-only and sulphur-only specialists are stored in the array of K_RECOMMEND structures (see Figure 3.12) and the array of S_RECOMMEND structures (see Figure 3.14) respectively within SITUATION. Summaries for the generation of potassium-only and sulphur-only recommendations are described in Figures 3.13 and 3.15 respectively.

```

struct K_RECOMMEND {
    int     meth_id;      /* Id of the method being used */
    int     fert_id;     /* Id of the fertilizer being used */
    double  fert_amount; /* Amount of fertilizer to apply */
    double  total_cost;  /* The cost of applying the K */
    char    *explain;    /* Reason we picked this combination */
    BOOLEAN keepit;     /* Indicates if we should keep this */
                /* combination or forget it due to */
                /* violation of toxicity constraints */
};

```

Figure 3.12: The K_RECOMMEND structure.

Determine the recommended level of potassium to apply using potassium recommendation tables from the Manitoba Soil Testing Laboratory

Form all valid potassium fertilizer source and application method combinations

For each combination, determine the amount of fertilizer to apply to achieve the recommended potassium level and determine the cost of fertilization

Display all potassium recommendations in ascending order sorted by cost

Figure 3.13: Generating K-only recommendations.

```

struct S_RECOMMEND {
    int     meth_id;      /* Id of the method being used */
    int     fert_id;     /* Id of the fertilizer being used */
    double  fert_amount; /* Amount of fertilizer to apply */
    double  n_amount;    /* Amount of N by-product produced */
    double  total_cost;  /* The cost of applying the S */
    char    *explain;    /* Reason we picked this combination */
    BOOLEAN keepit;      /* Indicates if we should keep this */
                /* combination or forget it due to */
                /* violation of toxicity constraints */
};

```

Figure 3.14: The S_RECOMMEND structure.

Determine the recommended level of sulphur to apply using sulphur recommendation tables from the Manitoba Soil Testing Laboratory

Form all valid sulphur fertilizer source and application method combinations

For each combination, determine the amount of fertilizer to apply to achieve the recommended sulphur level and determine the cost of fertilization

Display all sulphur recommendations in ascending order sorted by cost

Figure 3.15: Generating S-only recommendations.

3.3.3 NPKS Solution Method

After the single nutrient specialists have made their recommendations to the blackboard, the specialists at higher levels of abstraction can begin to integrate these single nutrient recommendations together to form NPKS recommendations for fulfilling all four nutrient requirements. The organization of the higher level specialists is influenced by the problem solution method. Recall that the central objective of the fertilizer problem is to find the economic optimum fertilization plan. In order to retain the yield potential characteristic of each fertilization plan for the purpose of comparing the economic consequences, it is best to take a nitrogen-only recommendation, integrate the phosphorus-only, potassium-only, and sulphur-only recommendations into it and then resolve any conflicts between the components with the goal of maintaining the rate of nitrogen application proposed by the

original nitrogen-only recommendation. Note that although the rate of nitrogen application does not change in the resulting NPKS recommendation, the cost of application will change due to additional fertilizers being added to the fertilization plan. This will alter the economic return of the fertilization plan and may make it more or less attractive relative to other fertilization plans. Thus we must construct all possible alternative plans which offer the same level of nitrogen supply to find the one with the best net return. A brute-force approach would be to first build all NPKS recommendations by taking all combinations of the nitrogen-only, phosphorus-only, potassium-only, and sulphur-only recommendations, and then comparing their relative returns on fertilizer investment. The problem with this approach is that there are too many combinations to generate-and-test. For example, consider the situation where there are five phosphorus-only recommendations, five potassium-only recommendations, five sulphur-only recommendations and five alternative nitrogen-only solutions for achieving each 5 kg/ha increments of nitrogen supply between 5 kg/ha and 130 kg/ha of nitrogen supply inclusive. Then there can potentially be 625 NPKS recommendations to generate-and-test to find the optimum fertilization plan at each 5 kg/ha increment of nitrogen supply and 26 increments to compare to find the overall best fertilization plan. That is a total of 16,250 NPKS recommendations to generate-and-test! A large number of these solutions are invalid solutions which break one or more constraints in the problem. Like the solution method in the nitrogen-only problem, we need to rely on constraints in the problem to direct the search through only the NPKS recommendations which have a high potential of being the optimum fertilization plan. Again, to manage the complex set of constraints that must be considered, the NPKS combination problem is broken down into abstracted levels of smaller problems where only a subset of the constraints have to be considered in each subproblem. Each subproblem will yield intermediate solutions that can eventually be integrated into the final solution to the fertilizer problem.

3.3.4 Two-Nutrients Specialists and PKS Specialist

The first task to perform after the generation of single nutrient recommendations is to combine the phosphorus-only, potassium-only, and sulphur-only (PKS) recommendations. These nutrient recommendations are combined first before attending to nitrogen for two reasons. 1) The fertilizers supplying these nutrients may contribute small amounts of nitrogen. Their total nitrogen contribution can be accounted for at once by integrating these single nutrient recommendations. 2) The number of resulting PKS combinations will be small. This is because the requirement for each of these nutrient is a single target nutrient level (in contrast to nitrogen where any target nitrogen level may be selected along a yield curve), thus the number of recommendations to consider for any one nutrient is the same as the number of valid fertilizer source and application method combination that can be used to apply that nutrient. The task of combining the PKS recommendations with the nitrogen-only recommendations at a later stage will be simplified if the number of PKS recommendations is reduced. The PKS recommendations are generated in two stages.

The first stage is to combine two of the three nutrients. That is, form phosphorus and sulphur (PS) recommendations, or phosphorus and potassium (PK) recommendations or sulphur and potassium (SK) recommendations. There is a corresponding specialist for handling each of these two-nutrients recommendations. The type of two-nutrients recommendations to form is determined by the crop being fertilized. Each crop has different priorities in terms of the nutrient requirements in the seed row. For example, it is more important to apply phosphorus and potassium in the seed row than sulphur and nitrogen for corn. Therefore, the PK recommendations should be formed first. The appropriate specialist is called upon opportunistically by the controller to combine the two most important PKS nutrients first before adding the least important nutrient. The paragraphs below will discuss the processing performed by first calling the PS specialist and then

completing the PKS recommendations by calling the PKS specialist to combine the PS recommendations with the potassium-only recommendations. The processing will be similar for other priorities of combining the PKS nutrients. For the various priorities, the only difference will be that potassium fertilizers do not contribute any nitrogen, thus the total PKS nitrogen contributions will be accounted for at different stages of the PKS processing. The appropriate specialists will handle this difference accordingly.

The responsibility of the PS specialist is to solve the subproblem of integrating the phosphorus and sulphur recommendations together to generate a list of all valid phosphorus-and-sulphur-only recommendations. The strategy is to check all combinations of phosphorus-only and sulphur-only recommendations for their validity. A valid PS recommendation must meet a number of constraints governing the combination of these fertilizer compounds. PS combinations which violate one or more of these constraints are discarded. Generally, these are global constraints which are applicable to any combination of fertilizer compounds. The first constraint is one that is a heuristic learned from experience in the field by our expert. It has been noted that farmers generally do not like to use more than one or two types of fertilizer application methods for convenience. Thus other than applying fertilizer while seeding, the farmer is not likely to perform more than one additional fertilizer application operation. For example, it is inconvenient to first incorporate some fertilizer into the soil before seeding, then plant the seeds (with or without any seed placed fertilizers) and later broadcast more fertilizer over the soil. Thus, PS combinations which suggest more than one non-seed row application method are not considered. Another constraint is that with average fertilization equipment, it is not possible to seed place or band fertilizers which are of different forms. The average equipment simply can not be used to apply more than one form of fertilizer compounds. Besides these crop management types of constraints, there are two soil science types of constraints which have been applied earlier in the generation of single nutrient recommendations and now

carry over into multiple nutrient recommendations. The first is the seed row nitrogen toxicity constraint of a fertilizer compound. Phosphorus and sulphur fertilizer compounds contribute a small percentage of nitrogen besides their major nutrient. As it turns out, when there is more than one source of nitrogen as in a phosphorus and sulphur combination, the upper limit on the amount of nitrogen which can be supplied through fertilization is 45 kg/ha of nitrogen or 20 kg/ha of nitrogen in the urea form. The last constraint on the PS recommendations is that no more than 195 kg/ha of total fertilizer product can be applied in a seed row application. That is, if a PS recommendation suggests seed row application for both the phosphorus and sulphur fertilizer compounds, then those combinations which use more than 195 kg/ha of total fertilizer product to achieve the required nutrient levels are discarded.

The list of PS recommendations which pass the above constraint tests are stored in the array of PKS_RECOMMEND structures (see Figure 3.16) pointed to by *pks_list on the SITUATION blackboard. The potassium (i.e., the K_RECOMMEND) component of the PKS_RECOMMEND will be completed by the specialist to be described in the next section. The above constraint tests will be applied again when other nutrient recommendations are integrated into the PS recommendations. The results of the tests at this stage are summarized and stored with each recommendation to simplify the same constraint tests to be applied later. The number and types of application methods used, the amount of nitrogen supplied in the seed row, the total amount of fertilizer product applied in the seed row, and the total cost of PS fertilization are summary information about a recommendation which will be useful when integrating other single nutrient recommendations with the PS recommendation.

In Manitoba, most soils are deficient in phosphorus. Therefore phosphorus fertilizer application will be required to give improved yield responses. However, a satisfactory level

```

struct PKS_RECOMMEND { /* A PKS fertilization plan */
    struct P_RECOMMEND *p_soln_ptr; /* Pointer to P component */
    struct K_RECOMMEND *k_soln_ptr; /* Pointer to K component */
    struct S_RECOMMEND *s_soln_ptr; /* Pointer to S component */
    double ps_sr_n_amount;          /* Total N supplied to */
                                    /* the seed row by PKS */
    double pks_sr_fert_amount;      /* Total fert product */
                                    /* in the seed row */
};

```

Figure 3.16: The PKS_RECOMMEND structure.

of sulphur is present in most soils and these soils will not require sulphur fertilizer application. Sulphur-only recommendations are, therefore, not needed. In this case, the PS recommendations list is simply a carbon copy of the phosphorus-only recommendations.

The second stage of generating the PKS recommendations is to add to the potassium to the PS recommendations. The subproblem of integrating the potassium-only recommendations into the PS recommendations is solved by the PKS specialist. The basic solution method is straightforward; build all valid combinations of potassium-only and phosphorus-and-sulphur-only recommendations. The constraint checks to determine the validity of a PKS combination are the same as those applied by the PS specialist. A PKS recommendation inherits the information from a PS recommendation (*i.e.*, a PKS_RECOMMEND structure) and the PKS specialist completes the potassium component of the PKS_RECOMMEND structure. Some information may have to be modified to reflect the effects of adding the potassium recommendations.

Based on the crop being fertilized, determine which two of the PKS nutrients should be given priority for application in the seed row. For this example, let the PS nutrients be given seed row application priority

Form all valid PS combinations

Combine each PS combination with each compatible potassium-only recommendation to form a list of valid PKS recommendations

Figure 3.17: Generating PKS recommendations.

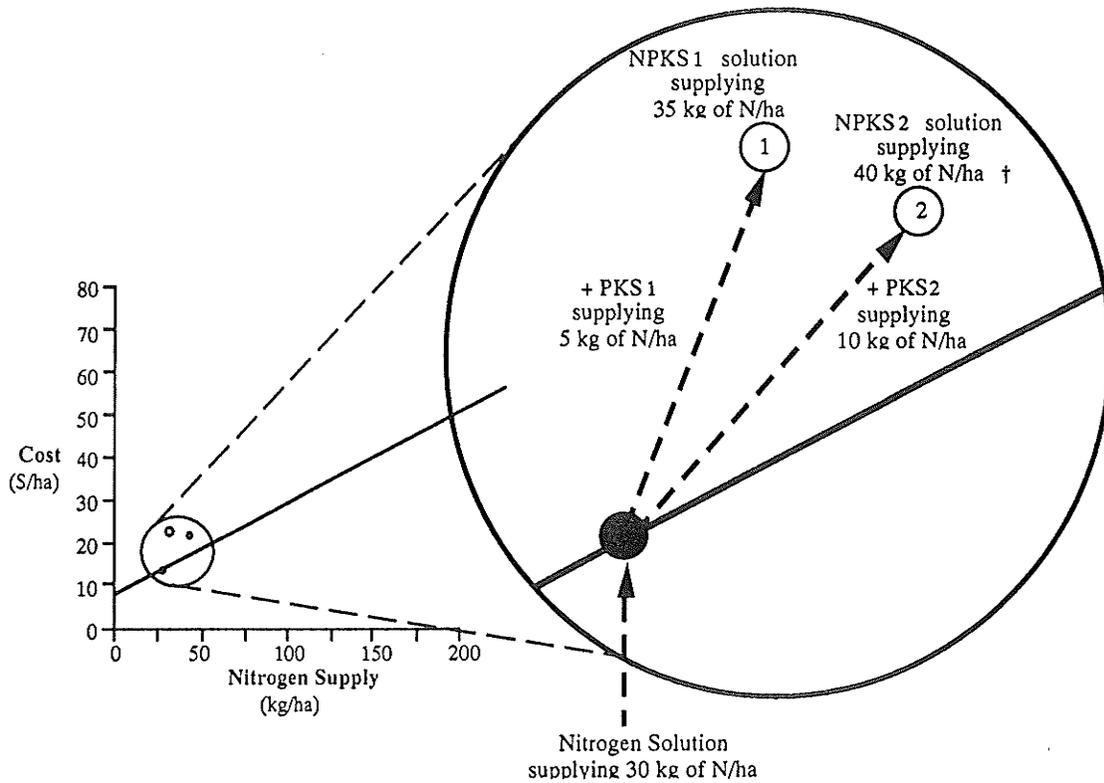
Because potassium fertilizer compounds do not contain any nitrogen, the amount of nitrogen supplied by a PKS recommendation and the effect on the toxicity constraint in the seed row does not change as a result of adding potassium to the base PS recommendation. The summary information that needs to be updated are the total cost of fertilization and total fertilizer product applied in the seed row (if the potassium fertilizer is seed placed). The resulting list of valid PKS recommendations is stored in the array of PKS_RECOMMEND structures pointed to by the *pks_list on the SITUATION blackboard. The array of PS recommendations formerly pointed to by the *pks_list is disposed of and the *pks_list is set to point at the newly completed array of PKS recommendations. Like sulphur, potassium may not be required in some fields. In this case, the PKS recommendations are simply the PS recommendations. Since each PS recommendation may be combined with zero or more potassium-only recommendations, the size of the PKS recommendations array pointed to by *pks_list may be greater than or equal to the size of the original PS recommendations array pointed to by *pks_list. A summary of the generation of the PKS recommendations is described in Figure 3.17.

3.3.5 Yield Curve Type NPKS Specialist

The specialist which is responsible for producing the list of the final NPKS recommendations (including the optimum NPKS fertilization plan) is the NPKS specialist. Since an NPKS recommendation will have a nitrogen component, the amount of nitrogen supplied can be used to determine the associated net return based on the nitrogen-to-yield relationship described in the previous chapter. The NPKS recommendations can now be compared in terms of their economic benefits. The NPKS recommendations are constructed by integrating the PKS recommendations into the nitrogen-only cost curves. Consider the simple approach of generating all combinations of nitrogen-only solutions with PKS recommendations. Note that each PKS recommendation supplies some amount of nitrogen.

The additional nitrogen will increase the resulting nitrogen supply when a PKS recommendation is combined with a nitrogen-only recommendation. For example, given that a PKS recommendation supplies 10 kg/ha of nitrogen, combining it with a nitrogen-only recommendation which supplies 30 kg/ha of nitrogen will increase the nitrogen supply to 40 kg/ha. The resulting fertilization plan will have a different yield potential than that associated with the original nitrogen-only component. Unfortunately, not all PKS recommendations will contribute the same nitrogen supply. Thus combining the same nitrogen-only recommendation with a different PKS recommendation which supplies 5 kg/ha of nitrogen will result in 35 kg/ha of nitrogen supply, and a different potential yield. The problem with this simple solution method can be illustrated graphically. Working from the nitrogen-only cost curves, giving a PKS component to each nitrogen fertilization plan will change its position on the graph because both its cost and its yield potential (revenue) is affected by the nitrogen from the PKS component. Completely new cost curves are formed and a search of the entire graph must be repeated to find the best return at each increment of nitrogen supply and the overall optimum return. This graph transformation is illustrated in Figure 3.18. Also, solving the NPKS problem by simply taking all combinations of PKS recommendations and nitrogen-only cost curves is inefficient because the number of solutions will increase by a multiple equal to the number of PKS recommendations. Furthermore, toxicity problems may also be introduced into an NPKS combination. Consider a partial solution at the higher levels of nitrogen supply. If such a solution requires the use of a non-seed row application to fulfill the target level of nitrogen because the seed row component has already supplied up to the toxicity limit, and the PKS component suggests seed row application of phosphorus and sulphur, the additional nitrogen supplied by the PKS combination may break the toxicity constraint and make the NPKS combination invalid. Discarding such a solution is undesirable because seed row application of phosphorus is the most efficient method of applying phosphorus. It is preferable to reduce the amount of nitrogen applied in the seed row to allow the seed-row

application of phosphorus and sulphur and make up the reduced nitrogen supply using more of the non-seed row component of the original nitrogen solution. The simple approach of combining a PKS recommendation with a pre-fabricated nitrogen solution does not allow for this adjustment of the nitrogen seed row component to nitrogen non-seed row component. Another limitation which may require the same type of adjustment is the total amount of fertilizer product applied in the seed row. Let us look at a more efficient method of combining the PKS recommendations with the nitrogen solutions. Keep in mind that the list of global constraints enforced during the construction of the PS recommendations and the PKS recommendations is still applicable in this stage of recommendation.



†Note: NPKS2 may be cheaper than NPKS1 even though it contributes more N.

Figure 3.18: The effect of combining an N solution with different PKS solutions.

Instead of building an NPKS fertilization plan and determining its total nitrogen supply as in the previous approach, it is more efficient to fix on a target nitrogen supply, determine the amount supplied by a PKS recommendation, and find the lowest costing nitrogen fertilization plan to complete the nitrogen supply to meet the target. The resulting NPKS recommendation must not violate any constraints. Since lowest cost is the main constraint, the search for the lowest costing and compatible nitrogen fertilization plan to combine with a given PKS recommendation can be optimized if the search through nitrogen fertilization plans is performed in ascending order by cost. The components needed to facilitate this optimization are two structures. One is the matrix pointed to by `*sorted_non_sr` in `SITUATION`. This matrix is identical to the one pointed to by `*sorted_solns` in `SITUATION` except that the new matrix only contains the sorted cost curves for the non-seed row applications. That is, the columns in the matrix represent ascending multiples of nitrogen supply and the rows in a nitrogen supply column represent the various nitrogen fertilization plans for achieving the given nitrogen supply in ascending order by cost. It is necessary to have columns showing the cost order between the application of different fertilizer compounds because the efficiency of the application method being used may affect the cost order as more fertilizer is applied (recall the effect of relative efficiencies on competing cost curves as presented in Figure 2.9). The second structure is an array representing seed row applications of the various nitrogen fertilizer compounds sorted by ascending cost of the fertilizers. This array is stored in `SITUATION` as the `sorted_sr` array. In contrast to the non-seed row applications, the cost order for the seed row application of different fertilizer compounds is only affected by the cost of the fertilizer compounds and thus cost order will not change as more fertilizer is applied. The NPKS recommendations can now be constructed using these two structures and the PKS recommendations.

For each 5 kg/ha multiple of nitrogen supply, FA finds the lowest costing NPKS combination which will produce the desired nitrogen supply. Given a target nitrogen level,

FA then completes each PKS recommendation with an appropriate amount of the cheapest compatible source of nitrogen to meet the target. There will be at most the same number of NPKS combinations as there are PKS recommendations. The cheapest NPKS combination is also the fertilization plan which will give the best net return at that increment of nitrogen supply. To complete a PKS recommendation, it is necessary to consider both seed row application of nitrogen fertilizer and non-seed row application of nitrogen. For brevity, a nitrogen fertilizer being placed with the seed in the seed row will be abbreviated as SrN and the nitrogen fertilizer being applied using a method other than placing it with the seed row will be abbreviated as NonSrN. Although the seed row application of nitrogen is generally more economical because there is no application cost incurred, using a NonSrN with a PKS combination costs less than using a SrN with the same PKS combination. This could be due to a number of reasons: 1) the NonSrN fertilizer compound has a much lower cost than the SrN fertilizer compound, 2) the SrN application has reached the seed row toxicity and must incur a second non-seed row application to complete the nitrogen supply up to the target amount, and/or 3) the PKS recommendation being completed has absorbed the application cost of the NonSrN because the same non-seed row application method is already being used to apply one or more components of the PKS combination. Completing a PKS recommendation with a SrN makes use of the sorted array of seed row applied nitrogen fertilizer compounds. Given that the PKS may have some fertilizer compounds already applied in the seed row, FA looks through the sorted array by ascending cost order and finds the first nitrogen fertilizer compound that can be applied in the seed row and is compatible with the other fertilizers (*i.e.*, seed row applied fertilizers must be of the same form). When a compatible SrN fertilizer is found, FA checks if it violates any seed row constraints in combination with the PKS recommendations. Given an NPKS combination where the nitrogen component is applied in the seed row, a check must be performed to determine if the combined nitrogen toxicity limit and/or the total seed row fertilizer product limit have been exceeded. If so, the SrN is applied up to the maximum amount allowed

without violating any constraints and the remaining amount is to be filled out by a NonSrN up to the target nitrogen supply.

The procedure for selecting the cheapest compatible NonSrN to finish off what the SrN can not complete is exactly identical to the procedure for selecting a single NonSrN (also the cheapest compatible one) to complete a PKS recommendation. Given the desired amount of nitrogen to supply, FA searches down the column which supplies that nitrogen amount in the sorted matrix of non-seed row applications of nitrogen fertilizer compounds. The first compatible NonSrN is used to complete the PKS combination. A NonSrN is compatible with the PKS combination if 1) there is no more than one non-seed row application involved in the resulting NPKS combination, and 2) if a banding method is used, the fertilizer compounds being banded must be of the same form. The process of combining the PKS recommendations with the nitrogen-only solutions is described in Figure 3.19.

The resulting list of NPKS recommendations contains the NPKS fertilization plans for achieving the highest net return at each 5 kg/ha multiple of nitrogen supply. This information is stored in the array of NPKS_RECOMMEND structures (see Figure 3.20) in SITUATION. For display purposes, the NPKS recommendations are sorted in descending order by net return. A sample display of the top three recommendations is shown in Figure 3.21.

Form list of alternative ways for applying nitrogen in the seed row (SrN's) sorted by cost
Form list of alternative non-seed row applications of nitrogen (NonSrN's) for each N increment
For each increment of nitrogen supply, find the lowest cost (and compatible) NPKS combination. This is done by combining each of the PKS recommendations with the cheapest compatible SrN and with the cheapest compatible NonSrN. Search through all pairs of PKS+SrN and PKS+NonSrN combinations to find the cheapest NPKS.

Figure 3.19: Generating NPKS recommendations.

```

struct NPKS_RECOMMEND { /* A NPKS fertilization plan */
    int    current_n_level; /* N level reached by this plan */
    struct N_SOLUTION  n_soln; /* N component */
    struct P_RECOMMEND *p_soln_ptr; /* Pointer to P component */
    struct K_RECOMMEND *k_soln_ptr; /* Pointer to K component */
    struct S_RECOMMEND *s_soln_ptr; /* Pointer to S component */
    double gross_revenue; /* Gross revenue generated */
    double total_cost; /* Cost of this plan */
    double net_return;
    BOOLEAN valid; /* Does plan violate any constraints? */
};

```

Figure 3.20: The NPKS_RECOMMEND structure.

Recommendations Summary	
Scenario	Fertilizer Plan
1	N - 108 kg/ha N, 82-0-0 (Anhydrous Ammonia), Spring Banded P - 35 kg/ha P2O5, 12-51-0 (Monoammonium Phos.), Seed Placed K - None Recommended S - None Recommended Potential Yld: 3150 kg/ha Cost: \$87.31 /ha Return: \$385.19 /ha
2	N - 100 kg/ha N, 82-0-0 (Anhydrous Ammonia), Spring Banded P - 35 kg/ha P2O5, 12-51-0 (Monoammonium Phos.), Seed Placed K - None Recommended S - None Recommended Potential Yld: 3125 kg/ha Cost: \$83.66 /ha Return: \$385.09 /ha
3	N - 117 kg/ha N, 82-0-0 (Anhydrous Ammonia), Spring Banded P - 35 kg/ha P2O5, 12-51-0 (Monoammonium Phos.), Seed Placed K - None Recommended S - None Recommended Potential Yld: 3170 kg/ha Cost: \$90.97 /ha Return: \$384.53 /ha
----- <F5>: Graphical Representation; <F7>: Options Menu; <ESC>,<F10>: Main Menu	

Figure 3.21: Example of an NPKS recommendations display.

3.3.6 Grid-Type NPKS Specialist

Recall the grid-type nitrogen-only recommendations for some crops where there is no predictable nitrogen-to-yield relationship. A straight nitrogen application rate is recommended based on the amount of nitrogen present in the soil. These recommendations are given by the Manitoba Soil Testing Laboratory Nitrogen Recommendation Tables. To

form grid-type NPKS recommendations, FA sets the target nitrogen level at the nitrogen supply level suggested by the tables and complete each PKS recommendation with the best compatible nitrogen-only solution to meet this target. The solution method to this problem is identical to the method applied over each 5 kg/ha increment of nitrogen supply in the yield curve-type recommendations. The result is a list of NPKS recommendations all of which produce the required nitrogen supply. Since there is no yield predicted for a grid-type NPKS recommendation, the lowest cost fertilization plan is the optimum recommendation.

3.3.7 Other Specialists in the FA System

Besides the specialists which contribute to the building of the NPKS combined recommendations, there are other specialists which offer specialized features to enhance the expert system. The following is a brief description of the features and the specialists which implement these features.

3.3.7.1 Explanation Facilities

Explanation is provided for the reasoning applied at decision points throughout the expert system. For example, when a rule is fired resulting in the adjustment of the default toxicities, explanations are provided to warn the user of the actions taken. Figure 3.22 shows an example of explanations displayed for some toxicity adjustments performed. The

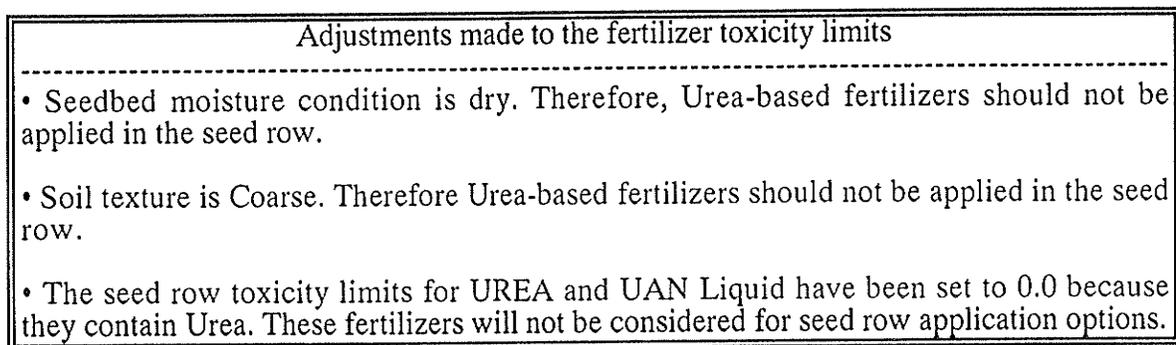


Figure 3.22: Example of explanations displayed for toxicity adjustments.

explanation specialist is responsible for setting up textual or graphical explanations (and also summaries) to present the information to the user.

3.3.7.2 Comparison Facility

In a recommendation system, it is beneficial to have the capability of performing what-if analysis by adjusting factors in the problem and comparing the impact of the differences in the recommendations. For example, compare the effects of different crop prices on net returns to fertilization. The comparison specialist uses two blackboard panels to generate recommendations for different scenarios. Each panel starts with the same set of input except for the difference in the factors being changed. The blackboard panels have identical organizations allowing the same recommendation specialists to work independently on the separate panels. When the recommendations are completed on both panels, the results can be compared either in a tabular format or graphically as curves. A sample comparison table for different sale price of Hard Red Spring Wheat under conditions which are otherwise identical is presented in Figure 3.23.

Comparison Summary						
Situation 1: Lower Wheat Price						
Situation 2: Higher Wheat Price						
N in Soil(kg/ha)	N Applied(kg/ha)		Total Cost		Return(\$/ha)	
	Sit1	Sit2	Sit1	Sit2	Sit1	Sit2
70	67	67	\$65.36	\$65.36	\$365.89	\$509.64
75	75	75	\$69.02	\$69.02	\$372.73	\$519.98
80	83	83	\$72.68	\$72.68	\$378.07	\$528.32
85	92	92	\$76.34	\$76.34	\$381.91	\$534.66
90	100	100	\$80.00	\$80.00	\$384.25	\$539.00
95	108	108	\$83.66	\$83.66	\$385.09	\$541.34
100	117	117	\$87.31	\$87.31	\$385.19*	\$542.69
105	125	125	\$90.97	\$90.97	\$384.53	\$543.03*
110	133	133	\$94.63	\$94.63	\$383.12	\$542.37
115	142	142	\$98.29	\$98.29	\$380.96	\$540.71
120	150	150	\$101.95	\$101.95	\$378.05	\$538.05

* - Best Net Return for Situation

<F6>: Details for Situation 1; <F7>: Details for Situation 2
 <F5>: Graphical Representation; <ESC>, <F10>: Return to Main Menu

Figure 3.23: Comparison table showing the effects of different crop sale prices.

3.3.7.3 Trace Facility

The constraint-directed search performed in FA by-passes many invalid or less than optimum fertilization plans. The fertilization plan preferred by a farmer for achieving a particular target yield may be a plan not recommended by the system. The trace facility allows a user to construct a fertilization plan and query the outcome. While the list of alternative recommendations shows the economic optimum fertilization plan for each increment of nitrogen supply selected by FA, the basic purpose of the trace facility is to allow the user to focus on any one increment of nitrogen supply and explore alternative (less than optimum) fertilization plans. The different paths of exploration offered by the list of alternative recommendations and the basic trace facility are illustrated in Figure 3.24.

The trace facility can be thought of as another blackboard panel where the user takes part as an active specialist in building NPKS recommendations. An example of a screen which represents this conceptual blackboard panel is shown in Figure 3.25(a). The user initiates the trace facility by selecting one of the recommendations from the list of alternatives. The trace specialist takes a snapshot of the selected recommendation onto the trace blackboard.

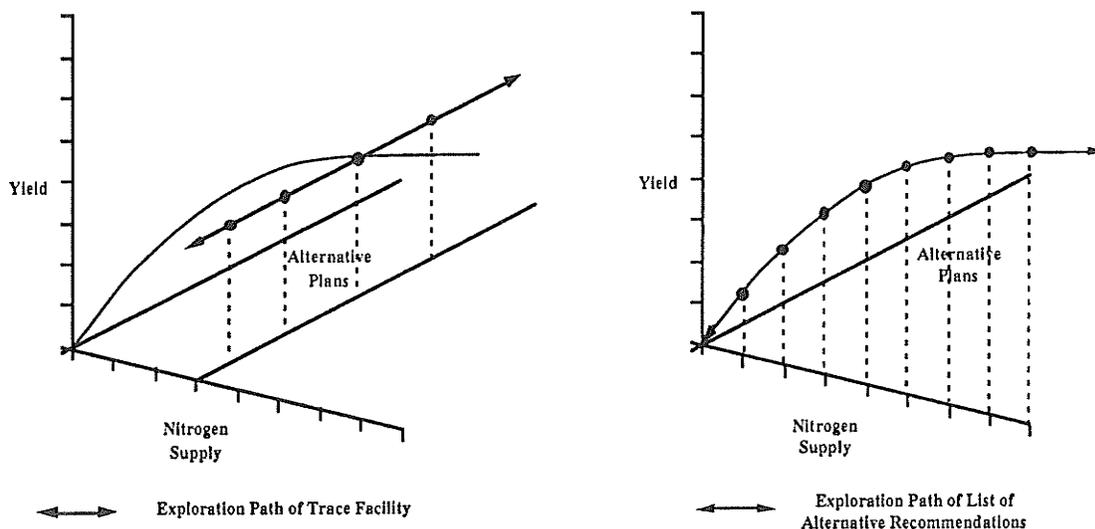


Figure 3.24: Exploration paths of Trace Facility and List of Alternative Recommendations.

The selected recommendation's nutrient levels become the target nutrient levels. The user is allowed to change the fertilizer source and application methods for any of the four nutrients. The trace specialist assumes that the user is going to apply as much of the selected fertilizer compounds as possible using the application methods to achieve the target nutrient supplies. The rates of each fertilizer to apply and the cost of their application are determined by letting the single nutrient specialists and the PS, PKS and NPKS specialists from the FA blackboard provide their expertise in solving the NPKS problem dynamically. The specialists also critique the user's recommendations and provide explanation for the status of the plan. For example, it may incur a higher cost than the system's best recommendation, or it may inadvertently violate toxicity constraints. Figure 3.25(b) is an example of the system's explanation of a constraint being broken by the user selected plan shown in Figure 3.25(a). In this example, although the user's plan appears to achieve a better dollar return, it is in reality not attainable because such a plan of fertilization will damage the seedlings and will result in far lower yield than expected. The critiquing ability of the trace component is designed to advise the user of such violations.

Examine Alternative Recommendations						
<u>System's Best Rec.</u>			<u>Your Recommendation</u>			
Nutrient	60 Target	Fertilizer N	60 Target	Fertilizer N		
Supplies	34 P	0 K	0 S	34 P	0 K	0 S
in kg/ha						
	<u>Method</u>	<u>Fertilizer</u>	<u>Method</u>	<u>Fertilizer</u>		
N:	(1) Spring Banded	82-0-0(0)	Seed Placed	46-0-0(0)		
	(2) None	None	None	None		
P:	(1) Seed Placed	12-51-0(0)	Seed Placed	12-51-0(0)		
K:	(1) None	None	None	None		
S:	(1) None	None	None	None		
	Yield	2758 kg/ha		2758 kg/ha		
x	Crop Price	x \$150.00 /T	x	\$150.00 /T		
=	Revenue	= \$413.70 /ha	=	\$413.70 /ha		
-	Cost	- \$80.00 /ha	-	\$74.58 /ha		
=	Return	= \$333.70 /ha	=	\$339.12 /ha		

<F7>: Detail Explanation <F10>: Explanation Regarding Problems

Figure 3.25(a): Trace Facility Screen.

Press <F10> function key to see explanation as shown in (b).

EXPLANATION

You are attempting to place 91.76 kg/ha of Nitrogen in the seed row using Urea. This amount of Urea will be too toxic if placed with the seed under the current soil conditions. The maximum that you can safely apply in the seed row using Urea is 20.00 kg/ha of Nitrogen.

Figure 3.25(b): System's critique of the plan in (a).

When the trace facility is initiated by the user selecting one of the recommendations from the list of alternatives, the comparison of the user's recommendation to FA's recommendation is for a particular nitrogen level. The user can change the nitrogen level and see the effects of using the same or a different fertilization plan to supply a higher or lower level of nitrogen. Graphically, this is equivalent to moving the exploration path shown in the left graph of Figure 3.24 to one of the adjacent tick marks along the nitrogen supply axis. For a given fertilizer and application method combination, the user can see the agronomic effects of using the same combination for supplying different amounts of nitrogen. For example, the user can choose to place Urea with the seed to achieve a 20 kg/ha of nitrogen supply. However, if the user decides to increase the amount of nitrogen supply using this same fertilizer and application method combination, the amount of fertilizer required will result in the fertilization plan exceeding the toxicity limit of Urea. The critiquing ability of the trace facility will warn the user of such a violation. This feature of allowing the user to step through different nitrogen supplies can be useful in explaining the reasons which cause FA to switch from using one fertilizer and application method combination to another as the amount of nitrogen supply changes.

In summary, a user can perform three different types of explorations using the trace facility. 1) For a given nitrogen supply, compare the options of using different fertilizer and application method combinations. 2) Using the same fertilizer and application method combination, see the economic and agronomic effects of that fertilization plan against FA's

recommendations for different nitrogen supply levels. 3) Change both the amount of nitrogen supply and the fertilizer and application method combination in comparison to FA's recommendations for the various nitrogen supply levels.

3.3.7.4 Multiple Crops

The crops specialist manages the retrieval of appropriate tables (*e.g.*, nitrogen yield response tables, phosphorus, potassium, sulphur recommendation tables, etc.) pertaining to the selected crop. These tables must be set up for use by the other specialists.

3.3.7.5 Database Routines

The users' data is organized and stored in a hierarchical fashion. Each client can have a number of crop fields. Each crop field can have a number of situations representing different scenarios for the field. Recommendations are performed for any one client's field situation. Situations can be compared using the comparison facility described above. The database specialist manages the organization and the low level saving and retrieving of user data to disk. The volume of data being stored can increase dramatically, especially when situations are being cloned and modified for comparison purposes. The database specialist applies an efficient data saving scheme. Each client has a default set of data describing an arbitrary field. This entire data set is saved on disk. When fields and situations are created, they inherit the default data. Only changes made to be different from the default data are stored to reduce the amount of disk storage required. Upon selecting a situation, the default set of data is retrieved and also the changes stored for that situation are retrieved from disk. The changes are applied to the default set to form the set of data describing the selected situation. This is analogous to an object-oriented environment where subclasses inherit the attributes and attribute values of ancestors unless explicit attributes and/or attribute values are specified for overriding the inheritance.

3.3.7.6 Imperial/Metric Conversions

In Manitoba, the agricultural community generally works with an imperial system of measure. However, for academic purposes, the FA system has been developed based on a metric system of measure. It is useful to have the option of selecting whether the system is to work in imperial or metric. A conversion specialist is built into FA to perform conversions between the two systems of measure. The data stored and the calculations performed are in metric. If the imperial system is chosen, then inputs (in imperial) must be converted to metric for storage, and the calculated results (in metric) must be converted to imperial for display.

3.4 The Fertilizer Advisor Blackboard vs. General Blackboards

In FA, there is an *a priori* order in which the knowledge sources are activated. This suits well to a serially executing computer program. In [Nii, 1986a], a blackboard framework appropriate for a serial-computing environment is described. Although FA's blackboard model evolved over development without reference to this model, there is a strong resemblance between the two models. What follows is a description of that blackboard framework as presented in the above mentioned article and how it correlates to FA's blackboard model. Keep in mind that FA is a program developed entirely in C, a procedural language.

The Knowledge Sources

The knowledge sources are represented as procedures, sets of rules or logic assertions.

Each knowledge source in FA is either a function, a group of nested functions, or a rule base processing routine which executes some rules. The toxicity specialist is a rule-based component which customizes the information to be used by the other processing specialists. The other specialists are generally written in procedural code. They utilize knowledge about the domain to solve the fertilizer problem. All of these knowledge sources are represented in the C language.

The objective of each knowledge source is to contribute information that will lead to a solution to the problem.

Each knowledge source is responsible for solving a subproblem of the overall fertilizer problem. The knowledge sources are designed to take the output of lower level knowledge

sources and apply their individual specialized knowledge to update the incoming information for knowledge sources at a higher levels. This action incrementally builds recommendations based on recommendations given by lower level. Adding one nutrient at a time to attend to individual nutrient requirements will eventually lead to the final NPKS recommendations.

The knowledge sources modify only the blackboard or control data structures (that also might be on the blackboard), and only the knowledge sources modify the blackboard.

The blackboard structure, SITUATION is accessible to each and every knowledge source. Knowledge sources use data from the blackboard and output data to it.

Each knowledge source is responsible for knowing the conditions under which it can contribute to a solution.

Due to the serial execution order of the knowledge sources, a knowledge source is given sufficient information to make its contributions. The lower level single nutrient knowledge sources have knowledge in the form of tables and/or logical assertions to decide whether or not they need to make a contribution. For example, if the sulphur-only knowledge source refers to the Manitoba Soil Testing Laboratory Sulphur Recommendation Table and finds that the level of sulphur present in the soil is above the required threshold, then it will indicate on the blackboard that a sulphur recommendation is not necessary. Higher level knowledge sources wait until the input data they need are available before they attempt to contribute to a solution.

The Blackboard Data Structure

The purpose of the blackboard is to hold computational and solution-state data needed by and produced by the knowledge sources.

The data produced by knowledge sources at one level is stored on the blackboard and is used as input for knowledge sources at a higher level. The knowledge sources interact with each other indirectly via the blackboard.

The blackboard consists of objects from the solution space.

The FA blackboard stores the list of alternative NPKS recommendations. These recommendations are a subset of the entire solution space. Intermediate recommendations (e.g., PS, PKS recommendations) are stored on the blackboard as well.

The objects on the blackboard are hierarchically organized into levels of analysis.

The data produced by knowledge sources at one level is stored on the blackboard and is used as input for knowledge sources at a higher level.

The objects and their properties define the vocabulary of the solution space.

The knowledge sources at each level produce data which is expressed by a different C structure. The type of information represented by the C structures varies from one level to another.

The blackboard can have multiple blackboard panels.

In the case of FA, multiple blackboard panels serve the purpose of generating two sets of recommendations, each illustrating different factors affecting a crop field. Recommendations can be compared in terms of economic return on fertilizer investment. Knowledge sources manipulate these blackboard panels one at a time and do not use the information from one blackboard panel to make contributions to another blackboard panel.

Control

There is a set of control modules that monitor the changes on the blackboard and decide

what actions to take next.

There is a main control procedure in FA which handles the execution of appropriate knowledge sources to generate the NPKS recommendations.

Various kinds of information are made globally available to the control modules.

Some knowledge sources indicate to the control module via the blackboard the action that should be taken next. Other knowledge sources are just one of a stream of knowledge sources to be executed in a fixed order. The control module is always aware of the focus of attention.

The focus of attention indicates the next thing to be processed.

The focus of attention is always on which knowledge source to execute next.

The solution is built one step at a time.

A model driven reasoning scheme is applied. The NPKS recommendations are built in steps starting with the single-nutrient recommendations. The P-only and the S-only recommendations are integrated into a list of valid PS recommendations. The K-only recommendations are added to the PS recommendations to form the list of valid PKS recommendations. Finally, the PKS recommendations are combined with the N-only solutions to produce NPKS recommendations. Appropriate knowledge sources are called upon opportunistically depending on the data made available to the blackboard. For example, if the type of crop being grown does not have tables or formulae describing the nitrogen-to-yield characteristics, then the grid type recommendation knowledge sources are called rather than the yield curve-type recommendation knowledge sources. Also, the priority for combining phosphorus, potassium, and sulphur recommendations will be opportunistically determined by the crop being fertilized.

Pieces of problem-solving activities occur in a prescribed iterative sequence.

The iterative sequence is defined *a priori* for solving the NPKS problem.

Criteria are provided to determine when to terminate the process.

The NPKS specialist is the last knowledge source to execute before the process terminates. It generates a list of alternative solutions to the NPKS problem.

In FA, the opportunistic problem solving characteristic is its ability to apply knowledge about the problem constraints to improve problem solving speed.

In this chapter, I have examined the problem of meeting the requirements of multiple nutrients. The implementation of its solution follows a blackboard framework. The use of specialists within a blackboard framework to solve small, manageable parts of the problem allows the problem solving components of the expert system to be incrementally developed. The next chapter will discuss the incremental development strategy used in building FA.

Chapter 4

Evolution of Fertilizer Advisor

4.1 Incremental Development

The distinctive quality of an expert system is the use of expert knowledge about the problem domain to efficiently solve the problem at hand. The problem solving knowledge could be represented in an unlimited number of different ways. Some common representation schemes include rules, frames, and object-oriented representations. The process of extracting an expert's knowledge and representing it in a suitable form for use by an expert system is called *knowledge engineering*. A knowledge engineer developing an expert system follows a development life cycle similar to that of software engineering. The cycle begins with the *definition* of the problem to solve. In sessions with domain experts, the knowledge engineer attempts to uncover the knowledge typically used to solve problems in the domain. The next step is to *conceptualize* the knowledge, concepts and relations uncovered in the problem definition stage. The concepts developed are then *formalized* into representation structures before the *implementation* of the expert system. The formalization stage includes the selection of an appropriate tool for development of the expert system and the specification of knowledge representation structures in the framework of the chosen tool. Implementation is the encoding of the formal specifications

developed in the previous stage into an expert system. The last but not the final step in the development life cycle is the *testing* of the implemented system. Testing usually leads to the discovery of details which must be refined and further problems to address. The cycle starts all over with the definition, conceptualization, formalization, implementation, and testing of solutions to these discoveries (refer to Figure 1.2). Expert systems are developed incrementally through successive extensions and modifications. The objective of this chapter is to present the iterative process involved in the development of the fertilizer selection expert system prototype.

Development of the fertilizer selection prototype began in 1989. The current prototype has evolved over two years of development. Work on the development of the expert system has involved the efforts of graduate students from the Department of Computer Science and staff members from the Department of Computer Science, the Faculty of Agriculture and the Solomon Sinclair Farm Management Institute at the University of Manitoba. My role as a research assistant in this project involved performing knowledge engineering and program development. The tasks of knowledge engineering and program development are shared amongst the members of the development team.

4.2 MacIntosh Prototype

In early 1989, the expert system was originally prototyped on a MacIntosh using a version of Common Lisp. The prototype was designed to provide nitrogen-only recommendations on the assumptions that all other nutrient requirements are satisfied. It demonstrated that using expert systems technology is suitable for the automation of the fertilizer selection task. The MacIntosh was chosen as the prototype platform because its mouse-driven facilities and the extensive graphics capabilities provided friendly interfaces to a user. Also, Common Lisp was a predominant expert systems development language at the time. Unfortunately, there were drawbacks to the MacIntosh environment. The application

developed using the version of Common Lisp was too slow for performing the amount of search required in the fertilizer problem. More importantly, the prevalent platform on which the system is intended to be delivered is the IBM MS-DOS environment. For these reasons, the prototype was ported over to the IBM PC environment at the beginning of year two.

4.3 IBM Prototype

4.3.1 Nitrogen-Only Version

Re-implementation of the MacIntosh prototype in the IBM PC environment using Microsoft C 5.0 began in the fall of 1989. The initial goal was to build a prototype that would have the same capabilities as the MacIntosh version (*i.e.*, to solve the nitrogen-only problem). A new user interface was required because the new environment lacked the mouse-driven facilities of the MacIntosh.

4.3.1.1 Hardware and software requirements

The hardware requirement was largely dictated by the existing machines currently in the agricultural community. A majority of the Manitoba Agriculture regional offices have IBM PC's with 8088 microprocessors and dual floppy disk drives. This fact shaped the initial specifications of the target delivery machine. Anticipating that speed of execution will be a major factor in the success of the expert system, the C language was chosen as the development language. Another key factor in the selection of C is that a comprehensive set of user interface tools written in C was already being used by other graduate students in the Department of Computer Science and the applications developed using these tools were well received.

4.3.1.2 User-interface, rule base component and explanations

From a design perspective, it was necessary to reorganize the user's interface to the program. The intent was to streamline the amount of information requested from the user

so as to reduce the complexity in using the program. The information relevant to the problem was analyzed and organized into three categories of information describing a farmer's scenario: the fertilizer compounds that may be used by the farmer, the methods by which these compounds can be applied, and the information describing the conditions of the field to be fertilized. A separate input screen was set up for collecting the information in each category. We found that structuring the information into categories helped the users focus their attention on individual categories of input. The organization of these categories into individual screens reduced the apparent complexity and volume of input required. To further reduce the amount of input required from a user, a prepared list of common fertilizer compounds (see Figure 4.1) and application methods were provided for the user on entry to the system. Thus upon entering a session, it was only necessary to modify the default fertilizer, method and field characteristics information before proceeding on to generate a recommendation.

Analysis Keep?	Analysis					Available Form	Fertilizer Products			Current Tox (kg/ha)
	N	P	K	S	Fertilizer Name		Price per tonne Spring	Fall		
Yes	82	0	0	(0)	Gas	Anhydrous Ammonia	\$360.00	\$310.00	0.00	
Yes	28	0	0	(0)	Liquid	UAN Liquid	\$160.00	\$150.00	30.00	
Yes	34	0	0	(0)	Granular	Ammonium Nitrate	\$225.00	\$210.00	45.00	
Yes	10-34	0	0	(0)	Liquid	Amm. Polyphosphate	\$300.00	\$270.00	45.00	
No	16-20	0	0	(14)	Granular	Amm. Phos. Sulf.	\$280.00	\$252.00	45.00	
Yes	12-51	0	0	(0)	Granular	Monoammonium Phos.	\$360.00	\$324.00	45.00	
Yes	0	0	60	(0)	Granular	Potash	\$160.00	\$160.00	45.00	
No	0	0	62	(0)	Granular	Potassium ?	\$170.00	\$153.00	45.00	
Yes	20	0	0	(24)	Granular	Amm. Sulphate I	\$215.00	\$193.50	45.00	
No	21	0	0	(24)	Granular	Amm. Sulphate II	\$215.00	\$193.50	45.00	
Yes	12	0	0	(26)	Liquid	Amm. Thiosulphate	\$193.50	\$174.15	45.00	
No	0	0	0	(90)	Granular	Elemental Sulphur	\$375.00	\$337.50	45.00	

<F8>: edit selected fertilizer in full screen mode
 <ESC>: exit without save <F10>: exit with save

Figure 4.1: List of pre-defined fertilizer compounds. Each row in the Keep? column may be toggled between yes or no to logically include or exclude the fertilizer represented by the row. Pressing <F8> allows the user to modify the current row's fertilizer information in a full screen mode.

The information which described the user's field included such details as: the soil test results, the crop to be grown and its expected sale price, the field's soil and moisture conditions, and the equipment to be used for fertilization (see Figure 4.2). For the prototype version, average default values were provided for these items upon start of the input session and modifications could be made to customize the information to reflect the farmer's current situation. Also some items could only take on one of a limited set of values. For example, the soil moisture level could only be described as one of dry, arid, moist, or ideal. The user interface tool implemented an entry of this type as an entry field which provided a list of the selectable options. This simplified the input process by allowing the user to select one of the options. Upon entering the field descriptions, a recommendation could readily be generated. As mentioned previously, the field condition may have an effect on the behaviour of the fertilizers. For example, if the soil moisture content is dry or the soil texture is coarse, it is not advisable to put nitrogen in urea form

Edit Client's Field Data			
Name:	Colin Farmer		
Field Name:	Test Field (Wheat)	Legal Desc.:	
<u>Soil Test Results:</u> kg/ha			<u>Rating</u>
Soil Nitrate - N Level:	30 NO3-N, 2ft.		Medium
Soil Phosphate - P Level:	20 Sod. Bicarb. P		Medium
Soil Potassium - K Level:	400 Amm. Acetate K		Very High
Soil Sulphate - S Level:	50 SO4-S, 2ft.		Very High
Crop to be Grown: Hard Red Spring Wheat			
Expected Crop Price: \$150.00/T			
Soil Texture: Fine (Clay)			
Seedbed Moisture Conditions: Moist			
Soil Moisture Curve to be Used: Moist			
Seeding Implement: Airseeder			
Seed and Fertilizer Placement: Separated or scattered			

<F5>: Options List; <ESC>: Don't Save; <F10>: Exit With Save			

Figure 4.2: The client's field data entry screen. A data entry field such as Seedbed Moisture Content has a fixed number of selectable values and can be scrolled through using the arrow keys, or displayed in a menu format by pressing <F5> when the cursor is on that entry.

with the seed at seeding time because doing so would likely damage the seed. To reflect this toxicity constraint, the nitrogen fertilizers in the urea form would have their toxicity limits lowered to zero. A rule base process then applies the field characteristic dependent adjustments to the fertilizer data.

The rule base processing component is part of a C-based rule base management package developed by Computer Science graduate students at the University of Manitoba [Miller, 1991]. The rule base management package has two components: a front-end knowledge acquisition component and a rule base processing driver component. The front-end knowledge acquisition has a rule base editor for entering the rules into a rule base. The rule base can also be tested through the knowledge acquisition component. When the rule base has been tested for validity, it can be exported to a C code version of the rule base. The C code rule base and the rule base processing driver are then linked into the program and a call can be made to the rule base processor to execute the rule base. To continue the example from the previous paragraph, the rule base processor is called after the completion all field description entries. The appropriate conditions will cause the rule:

IF Soil Moisture Content = Dry OR Soil Texture = Coarse THEN Urea Toxicity = 0.0
to be fired, resulting in Urea's toxicity being set to zero. There are other rules similar to this which take the information for a field and conditionally modify the current set of fertilizer compound and application method characteristics. Representing the knowledge about fertilizer behaviour under special field conditions as rules is advantageous in that the knowledge is explicit and is highly maintainable.

The recommendations produced by this version of the prototype are exactly as described in chapter three. The result is a list consisting of the best nitrogen-only recommendation for achieving each 5 kg/ha increment of nitrogen supply. All other nutrient requirements are assumed to have been satisfied by other means.

One aspect of expert systems which was explored during the development of this prototype was the capability to provide explanations for the actions taken by the system. Experts exhibit the characteristic that they can explain the reasoning they used in arriving at solutions. An expert system should have a similar capability in order to gain the user's confidence in the system's problem solving strategies. This attempt was modestly successful. Explicit knowledge such as the toxicity adjustment rules were straight forward translations from the rules to textual explanations. Other results such as the final nitrogen-only recommendations were found to be difficult to develop explanations for. The initial suggestion from our expert was to explain a selected recommendation by showing the details of that recommendation. Details included the fertilizer(s) to apply, the method of application, the rate, the amount of nitrogen applied, the cost of the fertilizer product, the cost of application, the total cost, the potential yield, the revenue generated and the net return. This detailed explanation has continued to be a part of the system even in the current prototype. It is useful in that a lot of the details about a recommendation are hidden until the user requests to see them. However, this type of explanation does not explain why the recommendation was selected over other fertilization plans which are capable of supplying the same amount of nitrogen. As will be seen in the following sections, the ability for the system to explain its reasoning continued to be one of the important issues in which development effort was invested.

4.3.1.3 Objectives for the Next Phase of Development

The IBM-implementation which solved the nitrogen-only problem was completed at the end of 1989. It became a stable platform on which additional knowledge could be added to make the system more robust at solving the fertilizer problem.

4.3.2 Nitrogen-and-Phosphorus Version

4.3.2.1 Knowledge Acquisition and Knowledge Engineering

At the beginning of 1990, the objectives for the FA project expanded to address the problem of integrating another nutrient, namely phosphorus, into the recommendations. This stage of development involved many knowledge acquisition sessions with domain experts to formalize the knowledge necessary to solve the phosphorus-only problem and also address the problem of integrating single nutrient recommendations to form combined recommendations. "*Knowledge acquisition* is the transfer and transformation of problem-solving expertise from some knowledge source to a program" (*or knowledge base*) [Buchanan et al., 1983]. The main sources of knowledge for FA were the domain experts. Other knowledge sources included text books, and data bases. Through a lengthy set of interviews, knowledge acquired from the domain experts was formalized into knowledge for representation in a knowledge base.

Knowledge acquisition tends to be the bottleneck in the construction of many expert systems [Buchanan et al., 1983]. Extracting knowledge from an expert involves extensive communication between the expert and the knowledge engineer. The exchanges are often awkward and clumsy in the initial stages when the knowledge engineer has limited familiarity with the domain. This problem is eventually overcome as the knowledge engineer acquires experience in the problem domain. Another problem which hinders the knowledge acquisition process is that "human experts are notoriously unreliable in explaining exactly what goes on in solving a complex problem. Often they forget to mention steps that have become obvious or even automatic to them after years of work in their field" [Luger et al., 1989]. This is actually not a fault, but rather a sign of expertise. The expert has mastered the problem solving skills and thus can solve a complex problem by going from point A to point C without explicitly going through point B. However, each intermediate step in solving the problem must be made explicit to the knowledge engineer

because the knowledge engineer is not an expert in the domain. The experts we work with are aware of this fact and make special efforts in their attempts to explain the problem solving techniques to us. Nonetheless, communication problems that occur between the domain expert and the knowledge engineer often impede the development process.

A typical session with our experts involved a review of the current system, the steps that have been taken to address the objectives proposed during the last session, and the determination of new objectives to be achieved for the next session. The objectives proposed included fixing program errors, improving the current system, and adding new features to the system. We have found that it was beneficial to demonstrate the most current version of the program at each session. The experts often played the role of the users and critique the usefulness of the system.

4.3.2.2 System Enhancements

During this phase of development, one important objective was to enhance the system's usefulness from the user's perspective. It was undesirable to have a clumsy system which would be difficult to use. Much work was put into improving the user interface. Effort was invested in two areas: adding a database management facility for storing and retrieving data for users, and experimenting with different arrangements of the system's components to increase the system's ease of use. With a database management facility, users could better manage the input data, reducing the amount of data entry in subsequent sessions. To further enhance the system's usefulness, we experimented with different screen arrangements. Subsequently, we have decided on an arrangement where the input screens requesting data that must be modified for each session (e.g., data describing the farm characteristics) were placed in the path towards the generation of recommendations. Screens which contained default data that would infrequently be modified (e.g., methods and fertilizers data) could be displayed when optionally selected by the user. This

arrangement allowed the novice user to quickly input the essential data and receive the corresponding recommendations, and yet allowed an experienced user to experiment with different sets of data to perform what-if analysis.

Another objective which was passed on from the previous (nitrogen-only) stage of development is the goal of improving the system's ability to explain its processing. In an attempt to explain the reason for selecting the optimum recommendation over other, less than optimum recommendations, FA was modified to display the competing recommendations. Again, such extraneous information was not normally displayed unless requested by the user. The system allowed the user to select an optimum recommendation and display all other, less than optimum recommendations which supplied the same amount of nitrogen and phosphorus as the optimum recommendation. It was apparent from this list that the optimum recommendation was chosen due to its high net return. However, such a list might not include a particular fertilization plan that the user is expecting to be economical.

To allow the user to see the effects of a specific fertilization plan, a *critiquing component* was added to FA. The critiquing component allowed the user to fill in, on screen, a fertilization plan. The user input the fertilizer products to apply, the methods of application and the amount of each product to apply. Existing modules which were already implemented for providing recommendations were called on by the critiquing component to critique the user's fertilization plan. If the plan was valid (i.e., does not violate any constraints), then calculation modules would be called on to determine details such as costs and nutrient supplies. The critiquing component could also determine if the plan violated constraints and warn the user of those violations.

The manner in which the critiquing component called on existing modules to critique the user's fertilization plan was made possible by the modular design of the system. Incremental development had lead us to implement the system one component at a time. The fact that the system had components which independently solved parts of the problem allowed the critiquing component to share these other component's problem solving knowledge. The value of the modular design was further supported by the work put into meeting two other objectives in this stage of development. The first objective, transforming the nitrogen-only recommendations into nitrogen-and-phosphorus recommendations, was accomplished by adding the knowledge about phosphorus recommendations to the existing nitrogen-only system. The integration of the processing for making phosphorus recommendations was achieved without extensively modifying the existing nitrogen component of the system. To meet the second objective, improving the user interface of the system, several different arrangements of the system's user interface components (and their associated processing components) were evaluated. In a relatively short period of time, several versions of the system were tested for their ease of use. We found that the system components could easily be configured in a number of different ways to achieve the same results. The differences between the configurations were in the presentation of information, screens and processing flow. This experiment illustrated that modular design methodology is valuable for further development of the program.

4.3.2.3 Objectives for the Next Phase of Development

At the end of summer 1990, the system was capable of providing nitrogen and phosphorus recommendations. The project team had a firm understanding of the integration of multiple nutrients and began to explore the integration of potassium and sulphur recommendations into the system. In preparation for more knowledge acquisition sessions, two major areas of focus were established for the next stage of development. 1) The addition of potassium and sulphur fertilizers to be considered by the system would dramatically expand the scope

of the problem. With each individual nutrient source being applicable using one of five application methods (and the nitrogen fertilizer may require a two method application), the total number of combinations could grow exponentially. There was a need to perform some type of guided search through only the relevant solutions. 2) There was still the need for a mechanism of providing comprehensive explanations for the reasoning applied by the system.

4.3.3 NPKS version

Development of the system continued in the fall of 1990 with the integration of potassium and sulphur recommendations into the existing nitrogen and phosphorus recommendation system. Throughout the fall of 1990, efforts were concentrated on developing a scheme for integrating the four nutrient recommendations into one combined optimum recommendation.

4.3.3.1 Controlling the Complexity of the Problem

After acquiring the knowledge necessary for implementing the potassium and sulphur components, the initial attempt at the integration of all four nutrients into combined recommendations was to exhaustively form all valid combinations and search for the optimum solution. The solution method was straight forward; the pseudo code representing the solution method is given in Figure 4.3. However, given the number of nested loops in the algorithm, the exhaustive method is time consuming. For an average test case (i.e., average initial soil test nitrogen level, only common fertilizers being considered), the generation of the recommendations took approximately ten seconds using a 10 MHz. 80286 processor. The worst case recommendations (i.e., low initial soil test nitrogen level, all fertilizers are considered) took 25 seconds to generate. The length of the response time was unacceptable, especially if the target end-user machines would have 8088 processors. The response time needed to be reduced. The strategy was to exploit some of the

knowledge we have about the interaction of the four nutrients to constrain the search for the optimum solution.

To reduce the complexity involved in dealing with each nutrient and the interactions between them, we incrementally built the problem solving components of the system such that the components themselves incrementally build the solution to the problem. The modular design methodology employed resulted in a system where the organization of the problem solving components was similar to a blackboard model. The problem solving components and the ways they contributed to solving the entire problem in a blackboard framework were described in the preceding chapter. The current system, tested under the same conditions as the initial system, took only three seconds to process the average test case and up to five seconds to process the worst case. The increase in the potential number of combinations to search through did not dramatically increase the execution time as in the exhaustive case.

```
For each 5 kg/ha increment of nitrogen supply
  Set current increment of nitrogen supply as the target nitrogen level
  For each phosphorus-only recommendation
    For each potassium-only recommendation
      For each sulphur-only recommendation
        If the PKS combination does not violate any constraints
          Form the PKS combination
          Determine nitrogen level supplied by the PKS combination
          Determine the amount of nitrogen required to meet the target
          For each nitrogen-only recommendations supplying the required amount
            If the NPKS combination does not violate any constraints
              Form the NPKS combination
              Determine the total cost of applying the NPKS combination
              Keep track of lowest costing combination at current nitrogen supply
            End If
          End For
        End For
      End For
    End For
  End For
  Determine net return for lowest costing NPKS at current nitrogen supply
  Keep track of the combination with the highest overall net return
End For
```

Figure 4.3: Pseudo-code for the exhaustive method of generating all valid solutions.

4.3.3.2 More System Enhancements

The organization of the problem solving components in a blackboard framework satisfied the first concern we had of the NPKS combinations growing exponentially as the number of factors to considered increased. The second concern regarding the improvement of the explanation facility was still going through experimental phases. Two features have been added to improve the system's ability to provide explanations. These are the *trace facility* and the *comparison facility* as described in the previous chapter.

In the NP version of the system, we implemented an explanation facility for showing alternative recommendations which achieved the same nutrient supplies as a selected optimum recommendation. Displaying all the recommendation alternatives excessively cluttered up the display screen. Furthermore, it was difficult to focus in on a specific type of fertilization plan a user was interested in to see how the plan of interest compared to the optimum recommendation. The development of the trace facility was motivated by the need to focus in on particular recommendations and compare their effects to the optimum recommendation given by the system. The trace facility replaced the critiquing component developed in the NP version of the system. The critiquing component was used to focus in on one particular fertilization plan for analysis of its economic effects. The trace facility uses the same interface as the critiquing component, allowing the user to select the fertilizer products and methods of application for each nutrient. However, instead of allowing the user to specify rates of fertilizer application, the system determined the appropriate rates of application for each fertilizer compound selected based on preset target nutrients supplies. These targets were set to be equivalent to the nutrients levels supplied by the optimum recommendation being compared by the user. This allowed the user to see the economic effects of the selected optimum recommendation and the recommendation proposed by the

user. Like the critiquing component, the trace facility warned the user if the proposed fertilization plan violated any domain constraints.

The second attempt at improving the explanation facilities was to extend the recommendation system to have the ability to compare recommendations made for different scenarios of the same field. That is, perform what-if analysis on the same field given slightly different circumstances. For example, what would be the economic effects and the recommendations given by the system, if the same field was fertilized with lower priced fertilizers. Such a facility is useful for illustrating the effects of changing factors on the economics of fertilization.

Finally, to prepare the system for acceptance in the agricultural community in Manitoba, facilities were implemented for handling multiple crops and operating in metric and imperial units of measure. The common unit of measure is imperial except for the price of fertilizers. To suit all users' preferences, either metric or imperial or a mixture of both units of measure could be chosen by the user. The selected units of measure were also saved for subsequent sessions.

Chapter 5

Conclusion

The ideal approach to making sound fertilizer application decisions is to consider all the factors that will affect the fertilization and devise a fertilization plan which will be customized to meet all the constraints imposed by these factors and yet will successfully yield the optimum return on investment. This is a complex task because there are many factors involved. First and foremost, the availability of the various fertilizers, the forms in which they are available, their nutrient compositions and costs must be determined. Next the rate, timing and method with which the fertilizers should be applied, along with the cost and the efficiency of the application methods must be taken into consideration. To maximize the effectiveness of the fertilizer application, many other factors must also be considered. To name a few, some of these factors may include: the soil type, texture, seedbed moisture content, concentration of nutrients presently in the soil, fertilizer toxicity problems, regional climatic rainfall, and farmer's preferences such as equipment availability, safety preferences, availability of financial resources, and crop management practices.

It is difficult to keep track of all of the above factors. This is further complicated by the substantial number of options available for choosing fertilizer products and suitable application methods. This combined with the goal of arriving at the best recommendation forms what we call the *fertilizer problem*. An agronomist would apply his/her knowledge of the agricultural domain and the knowledge of the farmer's specific circumstances to solve the fertilizer problem. There is no formal algorithm for carrying out this process. Instead, agronomists rely on facts known about the problem, their previous experience, intuition, and heuristics when making fertilizer application decisions. These problem characteristics suggested that the fertilizer problem was a good candidate for expert systems technology.

The result is the development of a prototype expert system called the *Fertilizer Advisor*. It takes a knowledge-based approach at solving the fertilizer problem. The many factors involved are considered in the problem solving process. Knowledge about the problem characteristics are applied to perform a constraint-directed search of the problem space. This search screens out irrelevant information, thereby reducing the potential complexity of the problem.

The system was developed using an incremental development strategy common to building expert systems. The process involved working in cooperation with experts to determine the knowledge required for solving a narrow scope of the problem and then implementing the solutions in the expert system. Gradually the scope of the problem was expanded and the process reiterated with the extraction of more knowledge from our experts about solving the enlarged problem, followed by the implementation of the solutions to incrementally complete the system's expertise.

The FA system was designed with a number of fundamental objectives in mind. Initial response from demonstrations of the system to groups in the agricultural industry have confirmed that these fundamental objectives were successfully achieved. These objectives were:

- 1) **Flexibility:** The system is capable of rendering recommendations customized to a farmer's circumstances. For example, different moisture regimes will result in the use of yield curves under different moisture categories. This allows the farmer to more accurately predict the yield potential. Unlike the manual system where recommendations have fixed yield goals, the FA system takes advantage of using the yield curves to more precisely choose an economic optimum yield goal. Furthermore, the selection of this yield goal is sensitive to the changing factors in the fertilizer problem.
- 2) **Constraints:** The system takes into consideration a number of constraints, but only when conditions require that those constraints be considered. For example, certain soil conditions will increase the toxicity hazard for seed placed fertilizer. These toxicity are adjusted accordingly to reflect the variety of field conditions which may tighten or relax these constraints.
- 3) **Speed:** The system is efficient in the generation of the recommendations. FA delivers recommendations within seconds due to its ability to apply knowledge about the problem characteristics to ignore irrelevant information in its search for the optimum recommendation.
- 4) **Ease of Use:** The system provides quick recommendations based on many default values stored in the knowledge base. The user may modify these if necessary. Furthermore, modified values may be saved for use in future sessions. Packaged with

the system are other facilities (*e.g.*, explanations, comparison facility, list of alternative recommendations, trace facility) which offer the user different recommendation needs. The user may be interested in knowing about the details of a particular recommendation, or why a fertilizer's toxicity was adjusted. These queries can be posed to the explanation facility. A user can explore alternative recommendations by two means: see the best alternative at each increment of nitrogen supply by displaying the list of alternative recommendations, or explore the alternatives at one particular increment of nitrogen supply using the trace facility. An interested user may also perform what-if analysis using the comparison facility.

Response from the demonstrations to agricultural groups have also identified some areas where improvements can be made. The major concern was that the system was lacking in expertise in some areas. For example, the system did not have data for predicting the yield potential of canola and flax, two major crops grown in Manitoba. This is in part due to lack of recent research data to formulate reliable yield curves. The need to update the data used by the system to reflect current crop production conditions has been identified as an objective to be tackled by agricultural researchers. Furthermore, other gaps in the system's knowledge base are to be filled in through more knowledge acquisition sessions with experts. For example, the system currently does not take into consideration variable soil testing depths. It only assumes that the soil test was performed at the depth prescribed by the Manitoba Soil Testing Laboratory. The ability to manage soil tests performed at other depths will require further consultation with our experts to determine the methods to handle the differences. Some areas were intentionally left for future development due to the need to restrict the functionality of the system during this research in order to avoid trying to solve too many problems at the same time. This is a common strategy used in the incremental development of expert systems.

The prototype system is expected to undergo thorough testing by users from the agricultural industry in the spring of 1991. The purpose of this testing phase is to uncover the strengths and weaknesses of the system in making real world recommendations. The users' responses will be used to determine the project's directives for reiteration of the incremental development cycle. To date, the system has been under significant scrutiny and we are confident that the underlying mechanism for solving the complex fertilizer problem will scale up gracefully to solve the fertilizer problem as the scope of the problem enlarges. New objectives have already been established as the system progresses into the next phase of development.

The first objective is to continue exploring ways to explicitly explain the knowledge applied by the system in solving the fertilizer problem. A user will be more confident with the recommendations suggested by the system if the system's reasoning is made explicit to user.

Another objective is to work with our experts in producing yield curves which have yields associated with small increments of nitrogen levels. The current yield tables show the corresponding yield for each 5 kg/ha increment of nitrogen supply. Having tables with finer increments (*e.g.*, yield for every 1 kg/ha increment of nitrogen supply) would allow the system to more accurately predict yield. The need to have finer increments in the yield table has been motivated by a demand for the system to perform target yield recommendations. That is, given a yield goal or a set of yield goals, what is the economic optimum fertilization plan for achieving the goal(s). Quite often, farmers prefer to meet specific yield goals due to market demands on their crop production. The following is what was conceptualized as the solution for solving this problem. The system currently can make recommendations for a given target nitrogen supply level (*i.e.*, grid type recommendations). Performing target yield recommendations is an extension of performing

target nitrogen supply recommendations because each nitrogen supply level is associated with a particular yield potential as dictated by the yield curve tables (Figure 2.2). Given a target yield, simply look up the nitrogen supply which generates that yield and perform a target nitrogen supply recommendation. The only drawback is that the nitrogen supply levels are in 5 kg/ha increments of nitrogen supply and their corresponding yields are in increments of 100 kg/ha or more. Given a target yield, it is likely that it is not in the yield table. Having finer increments will give more flexibility and accuracy in performing both target nitrogen supply and target yield recommendations.

Other objectives planned are features to cater to a variety of recommendation needs. Some farmers may want recommendations which meet a cost constraint. For example, the farmer has ten fields to fertilize and is willing to put \$10,000 into fertilization. Each field may have different field conditions and may be growing different crops. The system should perform an analysis for each field and determine the overall optimum allocation of the farmer's funds to maximize the economic return on fertilizer investment.

Another need is for a facility to provide recommendations on improving farm management practices to maximize the effectiveness of the fertilization. The application of the recommended fertilization plan to achieve the proposed target yield is only realizable given that proper management practices are exercised. It may be necessary to assess the history of a farmer's farm management practices and previously achieved crop yields before recommendations can be made regarding improvements in management practices.

Lastly, a long term goal is to make the expert system available for use by the various provincial soil testing laboratories and fertilizer dealers across Western Canada. The knowledge currently in the system is specific to Manitoba and must be replaced by the knowledge about crop fertilization practices in the province where the system is to be used.

The above problems have been identified as objectives to address in the next phase of development. Methods for solving these problems will be discussed in future consultations with our experts.

References

- [Buchanan et al., 1983] B. G. Buchanan, D. Barstow, R. Bechtal, J. Bennett, W. Clancey, C. Kulikowski, T. Mitchell, and D. A. Waterman, "Constructing an Expert System", in *Building Expert Systems*, Addison-Wesley, Reading, Massachusetts, 1983.
- [Erman et al., 1980] L. D. Erman, F. Hayes-Roth, V. R. Lesser, and D. R. Reddy, "The Hearsay-II Speech-Understanding System: Integrating Knowledge to Resolve Uncertainty", *Computing Surveys*, Vol. 12, No.2, June 1980.
- [Evans et al., 1989] M. Evans, R. Mondor, and D. Flaten, "Expert Systems and Farm Management", *Canadian Journal of Agricultural Economics*, 37, 1989, pp. 639-666.
- [Evans et al., 1990a] M. Evans, C. Li, P. Duddridge, M. Wilgosh, R. Mondor, D. Flaten, "Applying Expert Systems Technology to Farm Management", *Proceedings of the CIPS CONGRESS 90*, Ottawa, May, 1990.
- [Evans et al., 1990b] M. Evans, R. Mondor, D. Flaten, "Using Expert Systems to Generate Fertilizer Recommendations", *Journal of AI Applications in Natural Resource Management*, Spring, 1990.
- [Firebaugh, 1988] M. W. Firebaugh, *Artificial Intelligence: A Knowledge-Based Approach*, Boyd & Fraser, Boston, 1988.
- [Fox, 1987] M. Fox, *Constraint-Directed Search: A Case Study of Job-Shop Scheduling*, Morgan-Kaufmann Publishers, Los Altos, California, 1987.
- [Harmon, 1985] P. Harmon and D. King, *Expert Systems*, John Wiley & Sons, New York, 1985.
- [Luger et al., 1989] G. F. Luger and W. A. Stubblefield, *Artificial Intelligence and the Design of Expert Systems*, Benjamin/Cummings, California, 1989.
- [Manitoba 1988] Manitoba Agriculture, *Field Crop Production Guide for Manitoba 1988-1990*, Winnipeg, 1988.
- [McNurlin, 1988] B. C. McNurlin (Ed.), "Trends in Artificial Intelligence", *I/S Analyzer*, No. 2, 26(1988).
- [Miller, 1986] A. Miller, *Expert Systems: The Structure, history, and future of successful AI applications*, IEEE Potentials, (October 1986).
- [Miller, 1991] D. Miller, *Integrating AI and Conventional Computing*, Master's Thesis, University of Manitoba, 1991.
- [Newell et al., 1976] A. Newell and H. A. Simon, "Computer Science as Empirical Inquiry: Symbols and Search", *Communications of The ACM*, No. 3, 19(1976).

- [Nii, 1986a] H. P. Nii, "Blackboard Systems: The Blackboard Model of Problem Solving and the Evolution of Blackboard Architectures", *The AI Magazine*, Summer, 1986.
- [Nii, 1986b] H. P. Nii, "Blackboard Systems: Blackboard Application Systems, Blackboard Systems from a Knowledge Engineering Perspective", *The AI Magazine*, August, 1986.
- [Saskatchewan, 1986] Saskatchewan Agriculture, *Nitrogen Fertilization of Spring Annual Crops*, Agdex 54I, January 1986.
- [Waterman, 1986] D. A. Waterman, *A Guide to Expert Systems*, Addison-Wesley, Reading, Massachusetts, 1986.