A Whole-Farm Planning Decision Support System for Preventive Integrated Pest Management and Nonpoint Source Pollution Control

Jose Lopez-Collado

Dissertation submitted to the faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
In
Entomology

Nicholas D. Stone, Chair
Mary Leigh Wolfe
John W. Roach
Darrell J. Bosch
Timothy P. Mack

August 13, 1999
Blacksburg, Virginia

Keywords: Corn Rootworm Control, Constraint Satisfaction, Whole-Farm Planning, Nonpoint Source Pollution Control, Preventive IPM, Decision Support Systems

A Whole-Farm Planning Decision Support System for Preventive Integrated Pest Management and Nonpoint Source Pollution Control

Jose Lopez-Collado

(ABSTRACT)

A decision support system for preventive integrated pest management (IPM) and nonpoint source (NPS) pollution control was designed, implemented and evaluated. The objective of the system was to generate plans at the farm level to satisfy economic and production goals while limiting risks of insect pest outbreaks, nitrate and pesticide leaching and runoff, and soil erosion. The system is composed of a constraint satisfaction planner (CROPS-LT), a modified version of CROPS (Stone, 1995), a farm-level resource management system (FLAME), an NPS module, which includes a weather generator, CLIGEN (Nicks et al. 1995), and an NPS distributed-parameter model, ANSWERS (Bouraoui, 1994), databases, a database engine and utility programs. The performance of the system was analyzed and performance enhancing features were added to increase the planner's ability to find near-optimal plans within a limited planning time. Using heuristics to sort potential crop rotations based on profit generally improved the planner's performance, as did removal of fields that were not suitable for growing target crops. Not surprisingly, the planner was best able to find plans for crops that can be grown in a variety of rotational systems. Throughout, the ability to apply environmental constraints *selectively* to individual fields greatly improved the planner's ability to find acceptable plans. Preventive IPM (PIPM) heuristics to control corn rootworms CRW (*Diabrotica virgifera virgifera* and *D. barberi*) were added to the planner. The model was represented and solved as a constraint satisfaction problem. Results indicated that plans obtained using PIPM heuristics had less risk of CRW damage, reduced chemical control costs, higher profit and reduced soil erosion as compared to a control plan. Linking the planner to the NPS model in a feedback control loop improved the planner's ability to reduce soil losses while preserving economic and production goals.

I dedicate this research to my wife Josefina Solis Cruz, my son Egil Ivan, my daughter Jazmin, my mother Galdina Collado Garcia, and my father Francisco Lopez Sanchez

"Ut meus oblito pulvis amore vacet"

(Quevedo)

Acknowledgments

I would like to thank my adviser, Dr. Nicholas D. Stone who always listened to my ideas and suggested me how to improve my work in all aspects, and for his financial support. My committee members, Drs. Mary Leigh Wolfe, John W. Roach, Darrell J. Bosch, and Timothy P. Mack, who guided me with suggestions, directions, and corrections.

The Fullbright-CONACYT program, Colegio de Postgraduados, and Virginia Tech assisted me by providing financial support to complete my studies. Thank to the ISIS lab gang: Rebecca Scheckler, Clifton Wood, Benjy Cline, Jim Venuto, who taught me some skills to survive this computer-intensive research. Mrs. Claire Thomas kindly helped me with the ARC/INFO software. Special thanks to the faculty, staff, and graduate students of the department of Entomology; in particular to the graduate students of the ISIS lab: Peter Warren, Prabhakar Bhogaraju, and Jun Fan, whom I had the chance to talk about mutual interests in research and life. I also want to recognize my friends Rachel Gwin-Parker and Louis Gwin for all their help and support throughout these years. Jaime Acosta deserves a special mention, he is a friend who helped my family and me, and gave us so useful advice; my deepest thank goes to him. Hector Sosa was also supportive; he encouraged us to establish friendship with many people. Susan and Jun Fan, Lilia and Keith Kenyon, Rosario and Francisco Canales, Olimpia and Roberto Conte were always around, cheering and reminding me that research was not the only activity in this life. In Mexico, my brother Jorge managed many of my professional duties. Finally, I wish to give special thanks to my wife Josefina Solis-Cruz for her unconditional support. She took care of our children with the most patient attention I had ever seen and release me from some home-related daily activities I should had done.

Table of Contents

List of Tables

List of Figures

Chapter 1 Introduction

In the field of artificial intelligence, planning is considered as the process of searching for a set of actions that lead to the achievement of certain goals (Wilkins, 1984; Georgeff, 1987). In agriculture, planning has a slightly different meaning. It is associated with scheduling agronomic activities (Buick et al. 1992). Whole-farm planning is a comprehensive approach to farm management decision-making. The objective of whole-farm planning is to help the farmer or grower reach his or her personal goals while protecting natural resources (Bridge, 1993).

The last definition of farm planning is compatible with and includes preventive integrated pest management (IPM) as one of its components. Preventive IPM is the use of a set of pest control methods to avoid pest outbreaks (Pedigo, 1989). Crop rotation, an important component in conservation farm planning (Stone et al. 1992), is also a preventive IPM technique. Crop rotation can also reduce soil erosion, improve soil fertility, and increase soil organic matter content and crop yields. In addition, insect pests, pathogens, and weeds cause fewer problems in rotated systems (Lazarus and White, 1984; Lee et al. 1988; Pimentel et al. 1993; Brust and King, 1994; Sumner, 1982; NRCS, 1996b).

Since the 1980's, the development of formal farm plans has emerged as a necessity for many farmers to fulfill participation requirements in federal and state cost-share and price-support programs. The requirements have included productivity targets, restrictions on land use, practice standards for environmental protection, and other federal regulations (Kay, 1986; Buick et al. 1992; Stone et al. 1992). For example, farmers must have approved cropping and management plans for environmentally sensitive areas, like highly erodible land (HEL) fields, in order to participate in share-cost programs (Feather and Cooper, 1995). Consequently, sustainable agricultural practices are often used in farm planning to fulfill the protection of natural resources while maintaining profit (Schaller, 1993). These practices include rotations, balanced fertilizer use, adequate land drainage, minimum tillage, stubble retention, low stocking rate, use of plant cover to minimize soil erosion, and efficient water and pesticide management (Smith and McDonald, 1998).

Different approaches for farm planning have been studied. The approaches include geographic information systems (GIS) (Sharifi and Van Keulen, 1994), expert systems (Nevo et al. 1994), simulation models (Tsai et al. 1987), and planners or schedulers (Buick et al. 1992; Stone et al. 1992; Rellier and Chedru, 1992). What is evident from these studies is the overwhelming complexity of solving the farm planning problem. Farms today are still losing topsoil, polluting streams with fertilizers, and often not making money - not because they must but because we don't have a way yet to find the whole farm plans that can meet all their goals while conserving resources.

Problem Statement

In short, farm planning is an important and complex activity because of the many requirements, goals, and regulations farmers must satisfy or consider. Farmers ought to use their limited capital resources like land and machinery in an optimal way, and the amount of information required to successfully run a farm business is growing (Doyle, 1990; Heinemann et al. 1992). Manual evaluation of all the possible combinations of factors that affect farm planning is impractical and prone to errors (Kay, 1986). Whole farm planning systems can help farmers find good plans by reducing or efficiently traversing the search space and eliminating unsatisfactory solutions (Stone et al. 1992). However, computer-based planning requires a model of the natural system (farm), upon which to operate. Whole farm planning systems use various methods to mimic the effects of alternative farm plans on farm income and the environment. The most realistic approach, using validated simulation models, has proven to be too expensive computationally to build into whole farm planners during search (Stone et al. 1992). On the other hand, use of simulation is an accepted technique for evaluating crop management practices in terms of pollution control (USEPA, 1997b).

This research set out to develop ways to merge the predictive power of simulation models with the search efficiency of artificial intelligence planning methods to improve our ability to find optimal or near optimal farm plans for crop/ livestock farms in the mid Atlantic region.

Objectives

Four broad objectives were proposed:

1. Develop a whole-farm planning decision support system to simultaneously consider some of the most important factors required to find satisfactory farm plans.
2. Evaluate some sorting heuristics to obtain high profit plans in short time.
3. Include preventive IPM heuristics in the whole-farm planning system to improve overall pest management on the farm, while maintaining profitability and resource objectives.
4. Add landscape simulation results as feedback to a whole-farm planner to improve the planner's ability to find plans that control nonpoint source (NPS) pollution economically.

Research Steps

The required steps to achieve the previous objectives were the following:

1. Design and implement a whole-farm planning decision support system (WFP) to include a preventive IPM control model and an NPS simulation model.
2. Perform an analysis of the behavior of the planning subsystem; specially evaluate some sorting strategies to find plans quickly.
3. Design, evaluate, and apply a preventive IPM control model for corn rootworms (CRWs) *(Diabrotica virgifera virgifera* and *D. barberi*).
4. Explore the value of using an NPS model to evaluate farm plans and to provide feedback to the planning subsystem.

The WFP system created under step one was the basis for the remaining work. The scope of the WFP system included searching for plans, managing data, and interfacing with an NPS simulation model. The plans generated were required to satisfy goals and preferences of farmers about target crop acreage, annual crop production levels, and economic profit. Plans also had to satisfy constraints limiting soil erosion, and nitrate and pesticide leaching.

The first step was broke down into three specific tasks: building a planning subsystem, integrating the planning subsystem with an NPS model, and building a database management system to manage the data required by the planner and the NPS model.

The first specific task was aimed at developing and modifying a planning engine to obtain whole farm plans using a constraint satisfaction approach, based on the CROPS system (Stone, 1995). The CROPS system is a whole-farm planning system that includes preventive IPM to address the problem of pesticide leaching and runoff, among others (Stone et al. 1992). The second task addressed the problem of sending information to and retrieving results from an NPS model. The third task was aimed at designing the databases required for the whole system, as well as including the capacity to visually represent the information and execute economic analysis.

Step two concerned the evaluation of the behavior of the planning system. Because whole-farm planning is such a complex process, determining how well a planner is working is very difficult. In this step, I explored the effect of some alternative system configurations on the planning process. In particular, I evaluated the inclusion of sorting heuristics and the effect of changing crop production targets on the performance of the planner.

The third step was aimed to assess the inclusion of a preventive IPM control model into the planning subsystem. Corn rootworms (*Diabrotica virgifera virgifera* and *D. barberi*) were used as a case study to demonstrate the feasibility of using crop rotation as a preventive control technique within a whole-farm planning system.

The fourth step explored the utility of using an NPS model to improve the planning process through its inclusion in a feedback control loop. In CROPS, nonpoint source pollution is evaluated by using the universal soil loss equation (USLE), an empirical model (Weischmeier and Smith, 1978). However, NPS models have been used in systems like FARMSCALE. The FARMSCALE system (Wolfe et al. 1995) is a farm-level decision support system that uses ANSWERS, an NPS simulation model, to evaluate farm plans in terms of soil and nutrient losses. However, plans in FARMSCALE have to be created manually. Therefore, in this step, the automatic planning capabilities of a computer-based planner were combined with the predictive power of a NPS simulation model.

Overall, the final objective was to build a decision support system for farm planning similar to FARMSCALE (Wolfe et al. 1995) but with automatic planning capabilities that can help farmers better select alternative plans while preserving income and satisfying production and economic goals.

Justification of the Research

Planning systems help farmers to better manage farms (Kay, 1986). The use of computer-based planning tools minimizes errors and reduces the time to process large amounts of information, providing the farmer with useful information to make better decisions (Harsh et al. 1981; Heinemann et al. 1992). Pimentel et al. (1993) suggested that it is possible to reduce pesticide usage by 50 percent without decreasing crop yields or changing cosmetic standards by using management practices like crop rotation. Thus, by including preventive IPM control practices like crop rotation, a farm plan could require less pesticide application. In addition, planning to reduce environmental impact while maintaining the economic benefits sought by farmers can contribute to reduced nutrient loads in streams and reservoirs and improve water quality. The IPM initiative, launched in 1994, declared as a goal that 75 percent of the crops in the U.S. would implement IPM by the year 2000 (USDA, 1994). A decision support system that helps include preventive IPM practices in farm plans could help to reach that goal. Ultimately, the goal of developing whole-farm planning systems is to help farmers to deal with the complexity of farm planning by minimizing errors and by simultaneously managing the most important factors that affect farm planning.

Scope of the Research

The decision support system was developed to support dairy farm enterprises. The system builds plans from components taken from conventional and sustainable agricultural production systems (Stone et al. 1992; Buick et al. 1992). In this study, all biophysical data and practice data were from the state of Virginia; however, the system was designed to be able to suit production systems other than those in Virginia. Therefore, results from this research could be applicable with some modifications to different regions in the U.S. where farming is based on the use of crop rotations and environmental protection is a primary goal.

Contents

This dissertation is divided in 8 chapters. In Chapter 2, I present a review of literature relevant to preventive IPM, corn rootworms (CRWs) biology, life cycle, and control. Preventive IPM decision support systems for corn are also reviewed. The CROPS system is introduced and described as well. Also, the importance of NPS pollution and approaches to its control are described, and relevant NPS simulation models are introduced. Chapter 3 describes the design and implementation of the WFP system. The programming environment, system components, assumptions, and verification procedures are explained. Chapter 4 presents a numerical analysis of the behavior of the planning component. Strategies including domain ordering, varying target acreages and eliminating inconsequential fields are evaluated in terms of their effect on planning time and plan profit. Chapter 5 describes the design, implementation, evaluation, and application of a constraint-based preventive control model for CRWs. The model relies on crop rotations as the primary tool to control CRWs. Chapter 6 presents a reactive approach for NPS pollution control by using a simulation model to evaluate farm plans. Chapter 7 presents a summary and conclusions of the work. Finally, Chapter 8 makes recommendations for future research.

Chapter 2 Literature Review

Preventive IPM

Integrated pest management (IPM) can be defined as the combination of all the feasible methods to reduce and maintain pest populations below densities that cause economic damage while maintaining environmental quality. The control measures should be economically, ecologically, and socially acceptable (Stern et al. 1959; Pedigo, 1989; Zalom, 1993). There are two main IPM strategies: preventive and therapeutic. The therapeutic approach is the use of pesticides, microbial insecticides, and insect growth regulators to reduce pest populations once they have reached an economic threshold (Pedigo, 1989). Chemical control has proven to be effective as a therapeutic control; however, its continuous use over time has caused several problems such as soil and ground-water contamination, health hazards, secondary pest outbreaks, and loss of profit (Knight and Norton, 1989; Metcalf, 1980; Varshney et al. 1993).

Complementing the therapeutic approach, preventive IPM aims to prevent pest outbreaks. Preventive IPM techniques include the use of host-plant resistance, quarantines, cultural practices, and crop rotation (Pedigo, 1989). Crop rotation is a kind of environmental manipulation to reduce the pest's ability to survive or breed (Dent, 1991); and is most effective when the insects or arthropod pests to be controlled have restricted host ranges, short life cycles, slow reproduction, and low mobility (Metcalf and Metcalf, 1993; Dent, 1991; Bullock, 1992). The general effect of crop rotation on insect pests is to reduce populations by excluding the host (Bullock, 1992).

Crop rotation has been successfully used to control some insect pests. The Hessian fly, *Mayetiola destructor* (Say), and the wheat strawworm, *Harmolita grandis* (Riley), are controlled when wheat is rotated with other crops; also, rotations that include sorghum and legumes reduce the occurrence of the false wireworm, *Eleodes* spp. (Sumner, 1982). In corn, two important pests that can be controlled with crop rotation are the Western and Northern Corn Rootworms (CRW) (*Diabrotica virgifera virgifera* and *D. barberi*) (Gray and Luckman, 1994). Corn rootworms are important pests throughout the corn producing areas of the United States and Virginia (VCE, 1993; Gray and Luckman, 1994). Farmers use more insecticide to control this pest than any other row crop pest. Between 12.1 and 16.2 million hectares of corn are treated annually at a cost of about US$1 billion (Schroder, 1998).

The importance of corn cannot be overestimated. Corn is an important cash crop in the U.S. The production value in 1997 was estimated at US$24.4 billion with 29.8 million ha harvested (NASS, 1998). In the state of Virginia, corn grain ranked third in value of production with over US$81.5 million while corn silage had a value of US$64.5 million in 1997 (VASS, 1998). Also in Virginia, most field corn is grown continuously, and CRW control, therefore, relies mainly on the applications of granular insecticides. Youngman et al. (1993) estimated that 32 percent of continuous corn received applications of granular insecticides in the state of Virginia.

Western and Northern Corn Rootworms (*Diabrotica virgifera virgifera* and *D. barberi*)

Biology and Life Cycle

Levine and Oloumi-Sadeghi (1991) and Chiang (1973) have reviewed the biology and damage of corn rootworms (CRW). Here, a summary of the basic aspects of their biology, as related to their control, is presented. Both species are univoltine and have a narrow host range (Woodson and Jackson, 1996; Elliot et al. 1990; Boetel et al. 1992). Adult females lay eggs in the roots of corn in late summer. The eggs remain buried in the soil during the winter and hatch the next spring (Gray and Tollefson, 1987). The larvae produce most of the damage by feeding on the roots, causing lodging and reducing yield (Gray and Tollefson, 1987; Spike and Tollefson, 1991; Kahler et al. 1985).

Prolonged diapause is a trait that allows the eggs of CRW to survive for more than one year. Through the years, populations of the Northern Corn Rootworm (*D. barberi*) have appeared to increase the percentage of eggs with prolonged diapause possibly in response to the use of crop rotation. Prolonged diapause was reported to occur in less than two percent of the eggs in the 60s (Levine and Oleoumi-Sadeghi, 1991; Levine et al. 1992a), but this rate has increased to as much as 51 percent in east central Illinois in the 80s (Levine et al. 1992a). Prolonged diapause in the

Western Corn Rootworm has also been reported but at rates less than 1 percent (Levine et al. 1992b).

Control of CRW

Because of its univoltine nature, host-specificity to corn and low dispersion rate, CRW is a good candidate for control by crop rotation (Riedell et al. 1991; Roth et al. 1995). Alternating corn with non-host crops like soybeans, prevents the larvae feeding and causing damage (Levine and Oloumi-Sadeghi, 1991; Steffey et al. 1992). Crop rotation has been a common practice to control CRW since corn started to be grown in the U.S. (Levine and Oloumi-Sadeghi 1991). However, for economic reasons, continuous corn is often preferred by farmers (Allee and Davis, 1996; Meinke et al. 1998). When continuous corn is planted, soil insecticides applied at planting are recommend (Foster et al. 1986; Meinke et al. 1998; Riedell et al. 1991).

IPM Systems for CRW and Other Corn Pests

Decision support systems have come to play an important role in farm management due to the increased need of farmers to fulfill different goals and comply with regulations. The information needed to effectively drive a farm business has increased, and decision support systems can help to process that information (Plant and Stone, 1991; Heinemann et al. 1992). Corn is a major crop in the U.S. and several decision support systems have been developed to address some of the entomological problems that corn producers face. This section describes some of these systems. Phenological models like degree-days models are not included because they are used more for short-term forecasting (Jackson and Elliot, 1988; Elliot et al. 1990; Schaafsma et al. 1991; Levine et al. 1992b; Woodson and Jackson, 1996) than for planning over one or more years.

MAIZE

MAIZE was developed to assist farmers, private consultants, and county extension agents in making recommendations for within-season crop production decisions and developing crop production plans (Heinemann et al. 1992). MAIZE is a set of expert systems that offers recommendations to manage insect, diseases, and weeds. MAIZE gives recommendations for the current season until drying and storage of the product. A preseason module assists in identification of potential pest problems and suggests control tactics. For example, for corn rootworms, the system includes crop rotation as a suggested preventive measure.

VICE-Corn

The Virginia Insect Control Expert for Field Corn, VICE-Corn, is a field-level expert system designed to predict pest outbreaks requiring control for corn pests in Virginia (Buick et al. 1993). VICE-corn has three main components: a database to store information about corn pests, a graphical database to display images of the pests and their life cycles, and a set of rules used by the system to provide advice (Buick et al. 1993). Control recommendations are based on current management practices and environmental conditions entered by the user. The VICE-Corn system can recommend sampling, chemical, cultural, and biological control. Crop rotation for CRW control is included as a recommendation but no specific crop rotations are listed.

Pandora

Pandora (Bhogaraju, 1996) uses a case-based reasoning approach to evaluate the risk of pest outbreaks based on a crop rotation as input. The system is composed of a case-base library and methods for indexing, similarity assessment, and retrieval. The system retrieves cases from its library that are similar to the one under assessment. Based on the known occurrences of pest outbreaks in the similar cases, Pandora estimates pest outbreak risks for the new case. The case library consists of rotations typical in southwestern and south-central Virginia (Bhogaraju, 1996).

The Comprehensive Resource Planning System (CROPS)

The Comprehensive Resource Planning System is a whole-farm planning tool that focuses on crop production. The expected users of CROPS are farmers who want to move toward more sustainable agricultural practices, or those who face multiple environmental restrictions (Buick et al. 1992; Stone et al. 1992). A farm plan can be defined as a set of rotations (including tillage and management practices) assigned to the fields in a farm. CROPS considers plans as acceptable if they satisfy production goals and economic and environmental constraints. The plans must satisfy acreage and production targets for one or more selected crops across all the planning years and also limit pesticide leaching and runoff, nitrate leaching, and soil erosion control to acceptable levels, based on regulations and the preferences of the farmer (Buick et al. 1992).

The CROPS system uses a planning engine that considers the generation of farm plans as a constraint satisfaction problem (CSP) (Stone et al. 1992; Buick et al. 1992). The plan search is conducted in a three-level hierarchical mode. The first step is to satisfy field-level constraints. Rotations and management practices that result in unacceptable levels of soil erosion, pesticide leaching and runoff, or nitrate leaching are eliminated at this stage on a per-field basis. CROPS uses the universal soil loss equation (USLE) (Wischmeier and Smith, 1978) to assess soil erosion for the fields in the farm. The second step is to plan at the farm level, one year at a time. At this stage, CROPS searches for annual assignment of crops to each field of the farm that combined satisfy the requested acreage or yield for the target crops. It also performs an economic analysis to estimate the likelihood of successfully meeting the annual economic requirements. Finally, in the third step, backtracking plus partial arc consistency techniques adapted from Nadel (1989) are used to assign rotations to fields and field combinations to years to satisfy target crop acreages and to find satisfactory plans (Stone, 1995). A final economic evaluation determines if the plan is acceptable.

Nonpoint Source Pollution

Definition and Importance of Nonpoint Source Pollution

Pollution is defined by Novotny and Olem (1994, p. 13) as, “change in the physical, chemical, or biological quality of the resource (air, land or water) caused by man or due to man’s activities that is injurious to existent, intended, or potential uses of the resource.” In general terms, the sources of pollution can be classified as *point* and *nonpoint*. Pollutants from point sources enter the transport systems at discrete, identifiable places. Pollutants from nonpoint sources originate from wide areas such as agricultural fields, and enter waterways at intervals linked to meteorological events (Novotny and Olem, 1994).

The most common NPS pollutants are sediment, pesticides, heavy metals, microorganisms, and nutrients (Frere, 1982; Novotny and Olem, 1994). Each pollutant causes different types of damage. Soil erosion reduces on-site soil productivity while off-site effects include impairment of the quality of streams, water bodies, and rivers (NRCS, 1996a, 1998). Annual economic losses caused by soil erosion are close to US$27 billion from reduced productivity and additional losses of US$17 billion for off-site costs (NRCS, 1998). Pesticides, heavy metals, and sediments detrimentally affect fisheries, wildlife resources, and recreational activities (Crutchfield et al. 1993). Nitrogen from manure and fertilizers may negatively affect human health when percolated to wells as nitrates (Bouchard et al. 1992; Crutchfield et al. 1993). Nonpoint source pollution is the most significant source of pollutants to waterways, streams, and water bodies (Osmond et al. 1997). For example, Davenport (1994) mentioned that about 55 percent of the degradation of estuaries and coastal areas are caused by nutrients that come from nonpoint sources.

Nonpoint pollutants are mainly produced by agricultural activities. For example, the concentration of nitrogen in watershed discharges increases as the ratio of cropland increases (Jordan et al. 1997). The Chesapeake Bay's productivity has been deteriorating because of NPS pollution. A restoration program was established in 1983 with the target of reducing nutrient loads into the bay (Novotny and Olem 1994). A recent report (USEPA 1997a) indicates that there has been a significant reduction in nutrient loads, 15 percent in nitrogen and 19 percent in phosphorus. The goal is to reduce phosphorus loading by 40 percent by the year 2000.

Approaches to Reduce NPS Pollutants

Control of NPS pollution is generally associated with the installation of a set of "best management practices" (BMPs) in a watershed or other region. BMPs are defined as "methods or systems of methods for preventing or reducing NPS pollution to levels compatible with water quality goals" (Dillaha, 1996). NPS pollution control policy is designed to increase BMP use through voluntary and mandatory programs. Voluntary programs rely on education, technical assistance, and financial incentives to encourage adoption of BMPs. Financial incentives include cost-sharing programs like the environmental quality incentives program (EQIP) and the rural clean water program (RCWP), as well as tax breaks and direct payments. Mandatory programs are generally based on either practice standards or performance standards. The former stipulate what BMP must be used for certain situations. The latter specify some tolerance for pollution, such as an erosion rate (Feather and Cooper, 1995).

Best management practices can be structural or management-oriented (Wolfe et al. 1995). Examples of structural practices are terracing, streambank stabilization, and animal waste systems. Examples of management-oriented practices are crop rotation, tillage practices, filter strips, riparian buffer zones, and use of green manure crops (Dillaha, 1996). Best management practices are site-specific, a BMP may be appropriate at one site but completely ineffective in another (Dillaha, 1990; Wolfe et al. 1995). Factors that affect the effectiveness of BMPs are soil characteristics, topography, geology, climate, and land use. Estimating the effectiveness of a BMP is therefore imperative to improve the likelihood of attaining water quality goals.

Methods to estimate the performance of BMPs are monitoring and modeling. Monitoring implies sampling and analysis to estimate pollutant change over time, but it is usually expensive. Modeling is often a preferred alternative because it is more cost-effective but it has some limitations like data availability and model accuracy (Wolfe et al. 1995).

NPS Models

NPS model can be classified into different categories according to different criteria. Models can be *event-based* or *continuous*, depending on whether they simulate storm events or use a continuous time scale. Models can also be classified into *lumped* or *distributed-parameter* models. Lumped-parameter models are those that synthesize watershed processes into a single unit. Distributed-parameter models include spatial variability by specifically considering a watershed as an entity formed by several subunits, each with its own parameters (Novotny and Olem, 1994). There is a large literature concerning NPS models and NPS modeling. For example, the USEPA (1997b) gives an account of different watershed-level management tools, including NPS models. Here I restrict the discussion to some agricultural NPS models.

USLE

The Universal Soil Loss Equation (USLE) (Weischmeier and Smith, 1978) is a model that predicts soil erosion for a given cropping system, management practice, soil type, rainfall pattern, and topography. The model was empirically derived from experimental observations of approximately 10,000 plot-years of basic runoff and soil loss data. The USLE is not an NPS model but an erosion model that computes long term soil losses from sheet and rill erosion under specific conditions, but it does not predict soil deposition or sediment yields from gullies or streambank erosion. The model was intended to be used for conservation planning of farm fields or construction sites.

GLEAMS

The Groundwater Loading Effects of Agricultural Management Systems (GLEAMS) model (Leonard et al. 1987) was developed by the USDA-ARS to add a vertical flux pesticide component to the CREAMS model (Knisel, 1980). GLEAMS evaluates the effect of agricultural management systems on the dynamic flow of pesticides and soil erosion. The model is designed to be applied to field-size, relatively homogeneous areas. GLEAMS is a continuous, lumped-parameter model. The model includes hydrology, erosion, nutrient, and pesticide dynamics components (Leonard et al. 1987).

EPIC

The Erosion Productivity Impact Calculator (Williams et al. 1983) is a lumped-parameter, continuous model. EPIC computes crop growth, soil erosion, and nutrient and pesticide dynamics as affected by crop management practices. It also estimates soil productivity changes due to soil erosion. EPIC was used to prepare a study on the effect of soil erosion on crop productivity on a nationwide scale (Williams et al. 1983; Williams et al. 1990).

AGNPS

The Agricultural NonPoint Source (AGNPS) model (Young et al. 1989) is a distributed-parameter, event-based model. It was developed by the USDA-ARS to estimate the effect of point and nonpoint source pollutants on water quality. AGNPS simulates runoff, sediment, and nutrient transport over landscapes of almost any size, from just a few hectares in size up to 20,000 ha (Young et al. 1989). The event-based version has been discontinued, but a continuous version is in development (Voorhees, W.B. 1996. Personal communication. USDA-ARS North Central Conservation Research Laboratory. Morris, Minnesota).

EUROSEM

The EUROpean Soil Erosion Model (Smith et al. 1995) was developed by the European Union to address the particular conditions of European soils. EUROSEM is an event-based, distributed-parameter model that simulates soil erosion and produces hydrographs and sediment graphs on an event basis (Morgan et al. 1994; Morgan et al. 1998). EUROSEM's hydrology component is based on the KINEROS model (Smith et al. 1995).

ANSWERS

The Areal Nonpoint Source Watershed Environmental Response Simulation (ANSWERS) is a distributed-parameter, continuous simulation model developed to estimate the impact of management practices on soil erosion and runoff (Beasley et al. 1980). Bouraoui (1994) updated the model to make it continuous. The model was also expanded and improved by including a nutrient component and changing some hydrological components like infiltration.

WEPP

The USDA-ARS National Soil Erosion Research Laboratory developed the Water Erosion Prediction Project. WEPP is a continuous, distributed-parameter model. It is used to simulate the effect of crop management practices on soil erosion. The model estimates soil losses for "field-sized areas" or conservation treatment units, but it also works for watershed-level areas (Flanagan et al. 1995). WEPP was intended to replace the USLE model for conservation planning (Laflen, 1997).

CLIGEN

The CLImate GENerator is a weather generator, not an NPS model. However, it is discussed here because most of the NPS distributed-parameter models require daily meteorological data, like temperature and precipitation, to make long-term simulations. Weather data for such runs are often obtained by simulation (Richardson and Nicks, 1990). CLIGEN is based on the weather generator used by the EPIC model (Nicks et al. 1995). The model has been well tested in several locations of the U.S.

The required input exists for nearly 200 stations, and parameter estimation software and techniques are available (Richardson and Wright, 1984; Richardson and Nicks, 1990; Nicks et al. 1995). CLIGEN was chosen for this research because it provides the data required by ANSWERS. It is also the weather model used by the WEPP and FARMSCALE systems (NSERL, 1994; Nicks et al. 1995; Wolfe et al. 1995).

Chapter 3 Design and Implementation of a Whole-Farm Planning Decision Support System

Introduction

There was nothing in theory that prevented the creation of an integrated decision support system that could use data stored in a GIS and in other formats to initialize and run first a whole-farm planner, and then an NPS simulation model. In practice, however, the construction of such an integrated system presented many obstacles, from converting data formats to transferring data between computer platforms, to data visualization. As part of this research project, therefore, the development of a problem solving environment (Houstis et al. 1997) with which to set up experiments and evaluate results became of paramount importance.

System Components

The system is comprised of five components (Figure 3-1): a database module (databases and the Borland database engine, BDE), a planning module (CROPS-LT), an NPS module (CLIGEN and ANSWERS), an utility programs module (FTP and TELNET), and a farm-level management module (FLAME). The database module was used to store and retrieve data used by and shared between modules. The planning module was composed of a modified version of the CROPS planning engine with some new capabilities and constraint types. The NPS module was composed of CLIGEN and ANSWERS. CLIGEN is a weather generator that generates weather data to ANSWERS. ANSWERS is a distributed-parameter model selected to evaluate plans. Modifications to these programs were done to suit the system. Programs in the utility module served to transfer files and establish remote connections to machines. The farm-level management module had the task of data definition (create tables or indices, for example), data manipulation (edit tables or launch queries, for example), visual display of information, and economic analyses. Databases, database management components, and simulation models are common to many decision support systems (Turban, 1995).

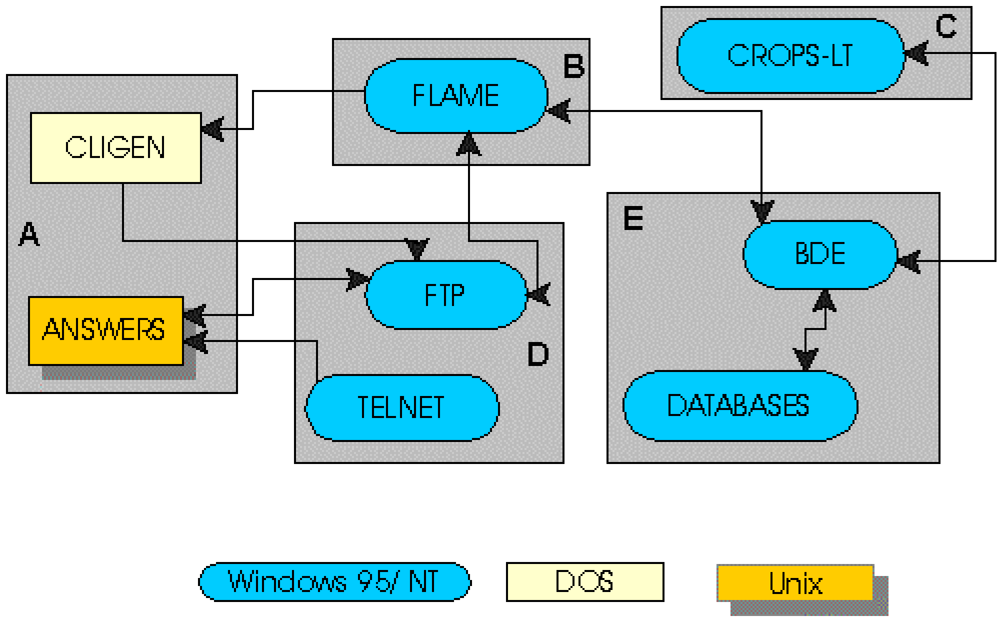

Figure 3-1 Main components of the whole-farm planning decision support system. Related components are grouped in modules (as gray boxes). (A) NPS module. (B) Farm-level management module (FLAME). (C) Planning module (CROPS-LT). (D) Utility programs module. (E) Database module. Blue components run under Windows 95/ NT, yellow components run in DOS, and orange components run in Unix. Arrows represent data flow.

The NPS programs were taken from the FARMSCALE system (Wolfe et al. 1995). Commercial FTP and TELNET programs were used for the utility programs module. The database module, FLAME, and CROPS-LT were created specifically for this project. Delphi version 2 was chosen as the programming environment to develop the database module, FLAME and CROPS-LT. Delphi is a rapid application development (RAD) tool (Borland International, 1995). I chose Delphi because the underlying programming language, Object Pascal, has object-oriented features, is faster to compile and to execute than interpreted languages, and provides many database management facilities and visual components to make Windows programming relatively easy. A detailed description of each module follows.

Database Module

The database module was composed of databases and the Borland Database Engine (BDE). The database design for the system is represented as an entity-relationship model diagram (ERD)

(Figures 3-2 to 3-5). The data are classified into *crop* and *farm* data. Common relations between entities (or tables) are one-to-one and one-to-many (see legend, Figure 3-2) (Mittra, 1991). Primary keys (PK) are elements used to uniquely identify records within a table. A complete description of each entity is presented in a data dictionary available on the Internet at: http://www.isis.vt.edu/~jlopezco/research/dissertation/data_dictionary.htm.

Data related to agronomic crops are shown in Figs. 3-2 and 3-3. Entities represent information about crop management and crop budgets. The ERD was modified from the one used by CROPS. Rotations, soil properties (*manning* and *surface_storage* entities), and soil erosion properties (*SLR_USLE*) were included as new entities (see data dictionary). Data stored in the tables *crop_operations*, *op_inputs*, and *inputs* were used to build the crop budgets (Figure 3-2). The tables *rotations* and *year_rotations* provided the data that the planning system needed to know about each crop rotation, for example, to calculate its environmental risk. Tables *crop_gleams*, *SLR_USLE*, *manning*, and *surface_storage* provided information needed to run ANSWERS (Figure 3-3A). The output generated by ANSWERS was stored in the entities *answers_output* and *answers_channel* (Figure 3-3B). The *machinery_input* table stored data about the cost and labor requirements incurred through the use of different types of machinery, including cost of fuel, oil and lubricant, repairs, depreciation, and labor. This entity is related to the *inputs* entity by the *inp_id* item (Figure 3-2).

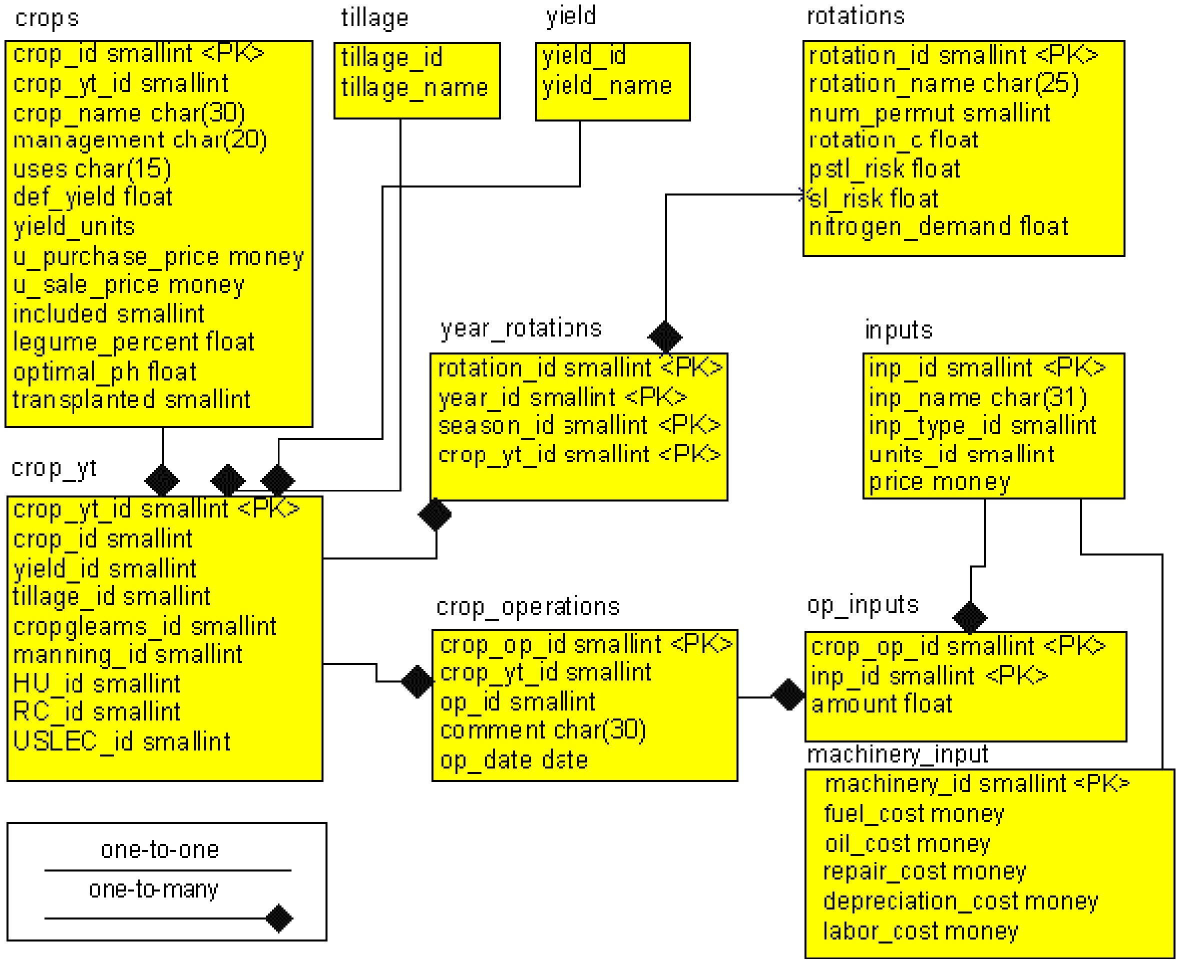

Figure 3-2 Entity-relationship model for crop information. Data are used for crop budgeting and planning. Primary keys (PK) are elements used to identify entities and for fast retrieval. Legend shows the correspondence between one entity to another.

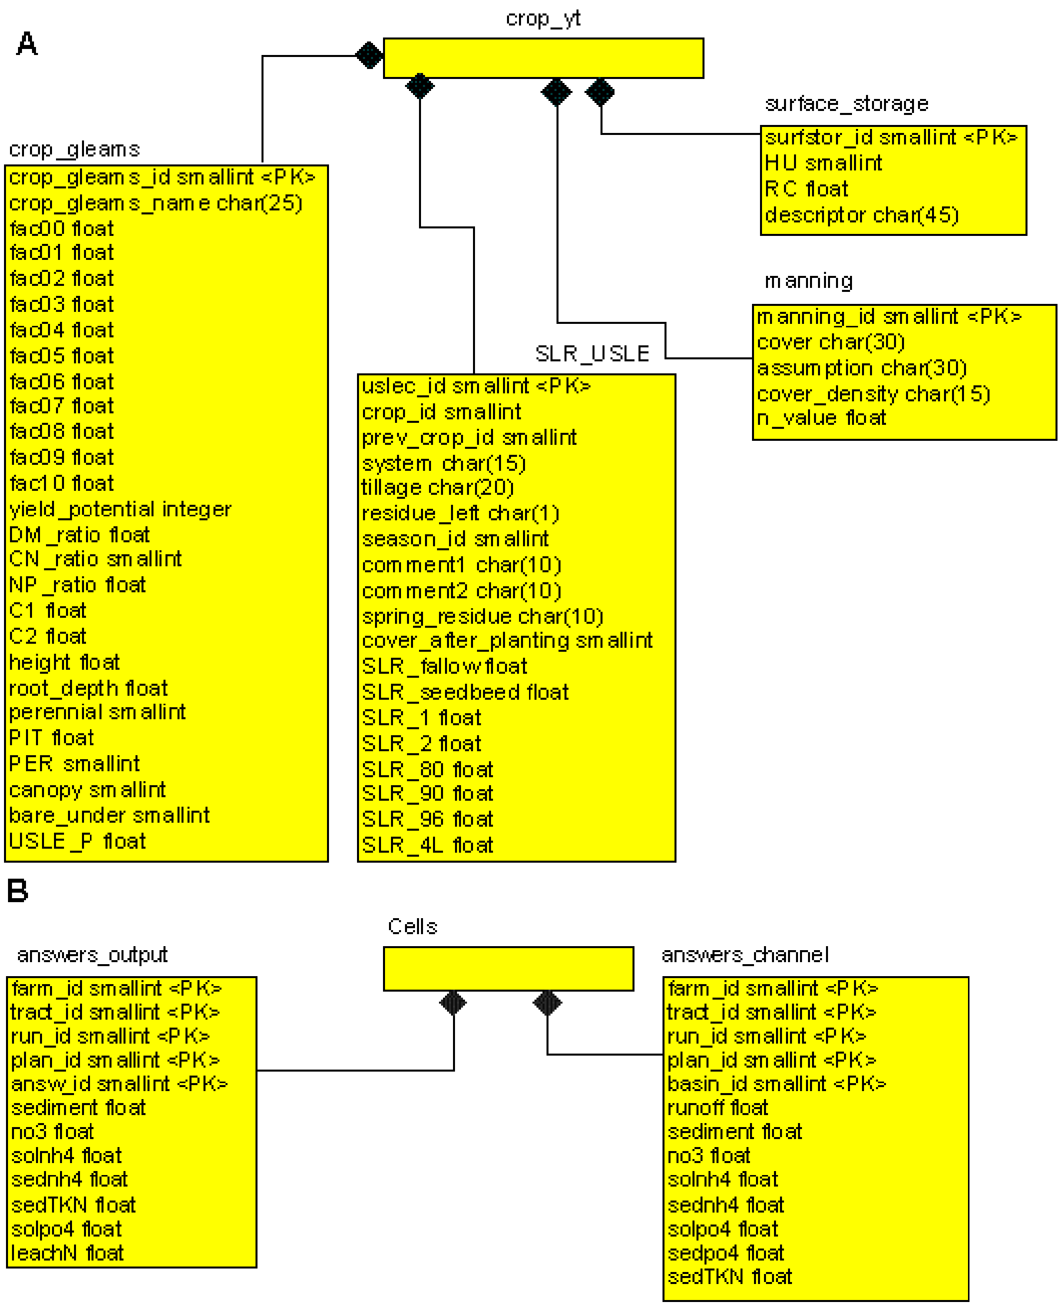

Figure 3-3 Entity-relation model for data related to ANSWERS. (A) Crop parameter data required by ANSWERS. (B) Output results from ANSWERS. The *cells* entity is described in Figure 3-4.

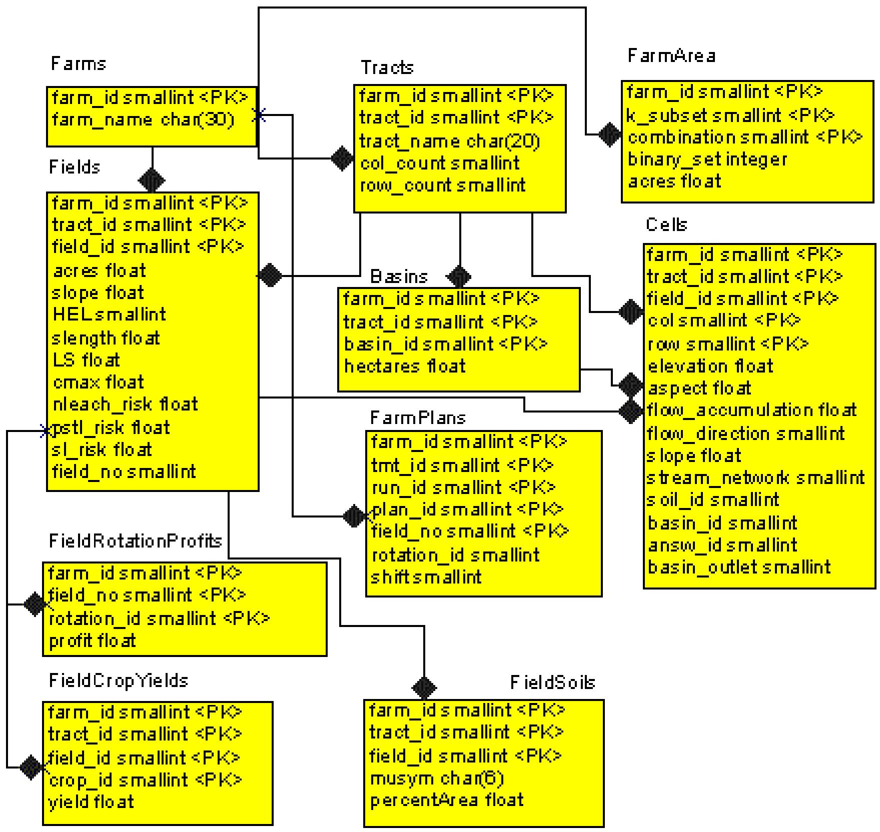

Figure 3-4 Entity-relationship model used to represent information for farms.

Farm entities are displayed in Figure 3-4. Data in this category describe farms, tracts, fields, basins, cells, and soil types per field, crop yield per field, crop rotation profit per field, field-combination areas, and farm plans. The finest level of resolution is the cell level. Cells are homogeneous squares arranged in a grid model of a landscape; they store topographic and soil information (*cells* entity). A grid structure is necessary because ANSWERS represents soil types, topography, and crop coverage at a cell level.

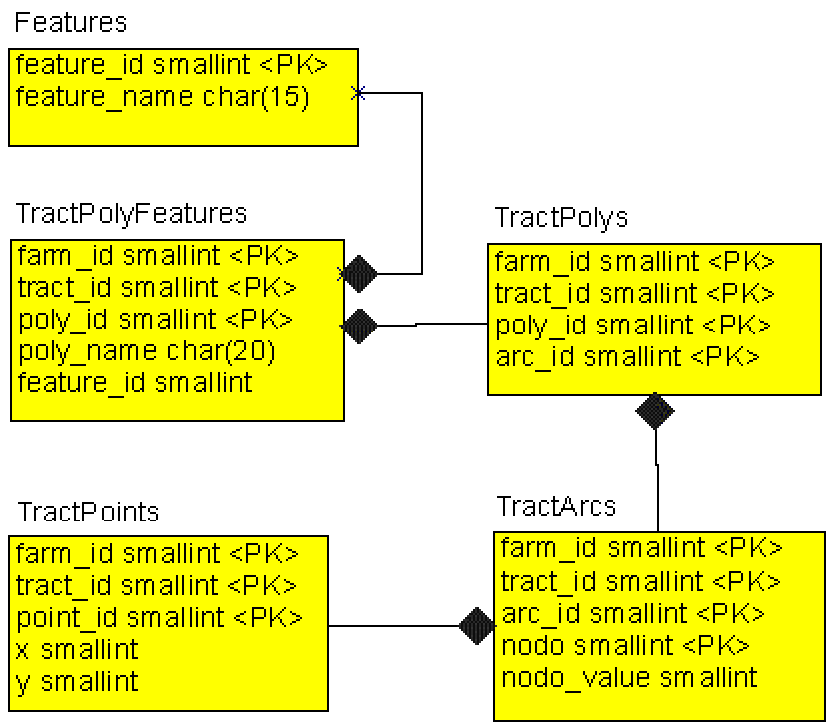

Figure 3-5 Entity-relationship model to represent farm polygons.

Another set of farm entities represents land boundaries. Land-unit boundaries (tract, fields, soils, and basins) are represented as polygons (Figure 3-5). A polygon is defined as a set of arcs. An arc is a directed line between two points. This is a common representation used by GISs and was used here to display such features (ESRI, 1997).

Table 3-1 Physical databases for the Whole-Farm Planning Decision Support System.

Database	*Purpose*
DBANWSCROPS	Provides data for the ANSWERS simulation model (GLEAMS, Manning, surface storage, and soil erosion parameters)
DBMOVE	Storage and transfer of text files to tables
DBCLIGEN	Provides data for the CLIGEN model
DBCROPS	Provides data relative to crops, crop budgets, and crop rotations
DBFARMS	Stores and manages data related to the farms: topography, soil composition, crop yields per field, rotation profit per field, and plans
DBPOLYGONS	Provides data to represent farm polygons: tracts, fields, basins, and soil types
DBVALUES	Stores the VALUES database (Simpson et al. 1993)
DBSOILS5	Stores the SOILS5 database (Goran, 1983)

All the database entities were implemented as tables in Paradox (Borland International, 1995) and stored as several physical databases for easy maintenance (Table 3-1). The DBSOILS5 and DBVALUES databases are not directly used by the planning system or simulation model but provide useful information related to soils, for example, they include crop yield data by soil type and other information required by ANSWERS.

Entity integrity refers to the representation of unique entities and was implemented by using primary keys (PK) in all the tables. Domain integrity refers to the condition that an item (column) in a table has to have a certain value. To enforce domain integrity, minimum, maximum, and default values were defined for items when appropriate. Referential integrity means that values in one or more columns (known as foreign keys) in a "child" table are obtained from values of columns (primary keys) in a “parent” table. In other words, the parent table provides valid values to the child table. If a record in the parent table is deleted and its value is used in the child table, to maintain referential integrity, all the associated records in the child table would also be deleted. Referential integrity was implemented only at the data entry level. Values for child tables were obtained from the parent tables by connecting both tables through appropriate windows within FLAME.

The system uses the Borland Database Engine (BDE) as the database management engine. The BDE is a set of dynamic link libraries that enable access to local and remote databases. Most data manipulations are handled by launching dynamic SQL statements embedded in Pascal code. SQL is the standard query language for database management (Gruber, 1993). In addition, a visual interface component in FLAME, the SQL manager window (Figure 3-24) allows the user to create a query using SQL statements. Visual editors allow a user to edit all the tables. These editors were developed as part of the FLAME component.

CROPS-LT: The Planning Module

The planning module consists of three elements: a planning engine or planner, an economic model, and a visual interface. The planning engine searches for whole-farm plans while an economic crop budget manager provides economic information to the planning engine. The visual interface allows the user to run the planner.

Planning Engine

The purpose of the planner is to search for and select plans at the farm level. A plan consists of crop rotations with associated management practices assigned to fields subject to different constraints and goals. The planner’s design was based on the CROPS planning engine (Buick et al. 1992; Stone, 1995). The engine uses a constraint satisfaction problem (CSP) approach to represent and solve the farm planning problem. I chose CROPS as the model to build the planner because its design is available and because of its capabilities to deal with qualitative variables common to the problem of farm planning (Stone, 1995). CROPS itself was not used because some changes in the algorithms and new types of constraints were needed. This section introduces the general CSP, the CSP terminology and the basic algorithms to solve it. Next, a description of how the planner solves this problem is outlined, as described by Buick et al. (1992) and Stone (1995). Finally, the farm planning problem is formulated as a CSP, and the implementation is described using an object-oriented approach.

A constraint satisfaction problem can be represented as a network graph (Figure 3-6) where the vertices or nodes represent a set of variables, **V**, each associated with its corresponding domain, **D**, of possible values (Mackworth, 1987; Weiss, 1994). Constraints are logical conditions used to determine the validity of a particular domain item, d, for a given node, V. Unary constraints apply to a single node. They test whether that node may be assigned a given value from its domain. The problem is solved by assigning a value to each node from its domain, such that all constraints are satisfied. Node consistency (NC) is the process of testing values of the domain against all unary constraints defined for a node. For example, consider a node *color_of_shoes*, with domain, {red, green, blue, yellow}, and a unary constraint that restricts the values to red, green or blue. When NC is applied to *color_of_shoes*, the value yellow would be eliminated.

A binary constraint defines a valid relationship between the values of two nodes (Figure 3-6). It determines whether a value from one node is valid given that the second node has a particular value, already assigned. Two nodes, A and B, with a binary constraint from A to B, are said to be arc-consistent when for each item remaining in the domain of B, there is at least one value in the domain of A that if assigned to A would satisfy the constraint. Removing values from domains that can find no supporting values in domains of nodes they depend on is called establishing or applying arc-consistency (AC). Suppose for example that a second node, *color_of_dress*, has values {red, green, purple} and that a binary constraint from *color_of_dress* to *color_of_shoes* exists requiring that shoe color match dress color. To establish AC, *color_of_shoe* will eliminate blue and yellow. Full arc consistency is obtained when all the arcs in a network graph are made consistent (Mackworth, 1987).

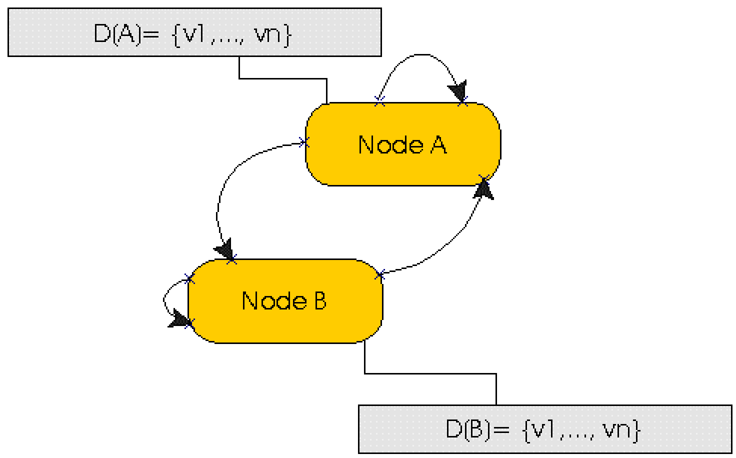

Figure 3-6 Graphical representation of a constraint satisfaction problem (CSP). Rounded rectangles are nodes that correspond to variables and square rectangles represent a set of values or domains. Directed arcs (arrows) represent constraint relations. Unary constraints point to a single node. Binary constraints define a dependence of one node on another node.

Different methods exist to solve CSPs. The general algorithm, called depth-first or tree-search (Nadel, 1989; Russell and Norvig, 1995) is to start with a node, assign one of its domain items to it, check constraints, and, if no constraint has been violated, proceed to the next node. If a constraint is violated, the next domain item is tried. In practice, applying node-, arc-, or other consistency checks before or during the assignment can greatly improve performance (Nadel, 1989). For example, AC before tree-search eliminates domain values from each node that are not part of any solution. After applying AC, if any node has an empty domain, no solution exists. Otherwise, a tree-search is used to find solutions from the remaining values (Mackworth, 1977; Kumar, 1992). Also, the order in which nodes and their domain items are selected can have dramatic effects on search efficiency (Sadeh and Fox, 1995). In all cases, it is assumed that node consistency is executed first to remove unsatisfactory values.

The farm planning problem was constructed as a CSP following Buick et al. (1992). The goal was to assign a six-year crop rotation to each field of a farm in a way that satisfies the farmer's goals, represented as constraints. As in the CROPS system, the planning engine uses a three-level hierarchical approach for planning (Figure 3-7). Before starting the search, target crops have to be selected and appropriate yield or acreage targets established. The first level attempts to satisfy annual requirements. The planner finds assignments of target crops to fields within a single year (called *field combinations*) so that annual target acreages or yields are met (See Figure 3-7A for an example). The second level addresses field-level requirements. From a master list of crop rotations and associated management practices, the planner assigns rotations to each field and then performs NC using constraints that limit soil erosion, and nitrate and pesticide leaching and runoff risks to acceptable values (Figure 3-7B). At this level, the CSP nodes are the fields, and rotations comprise their domains. In level three, the planner attempts to assign rotations to fields and combinations to years to complete a plan. Binary constraints ensure that crops from the field combinations and the rotations match appropriately (Figure 3-7C). Two constraints accomplish this[1]: a field-to-crop constraint establishes that the crop in the y-year position in a rotation assigned to a field must match the target crop in the same field in the combination for that year (Figure 3-7C). A crop-to-crop constraint states that two crops cannot be planted at the same time in the same field unless they are double cropped. Selection of compatible rotations and combinations is performed using a tree search combined with forward checking and partial look future algorithms (Nadel, 1989). Each level includes checks to detect failures in the search.

[1] In Stone (1995), the field-to-crop constraint is named *Crops Must Match* and the crop-to-crop constraint is called *Combo-to-Combo*.

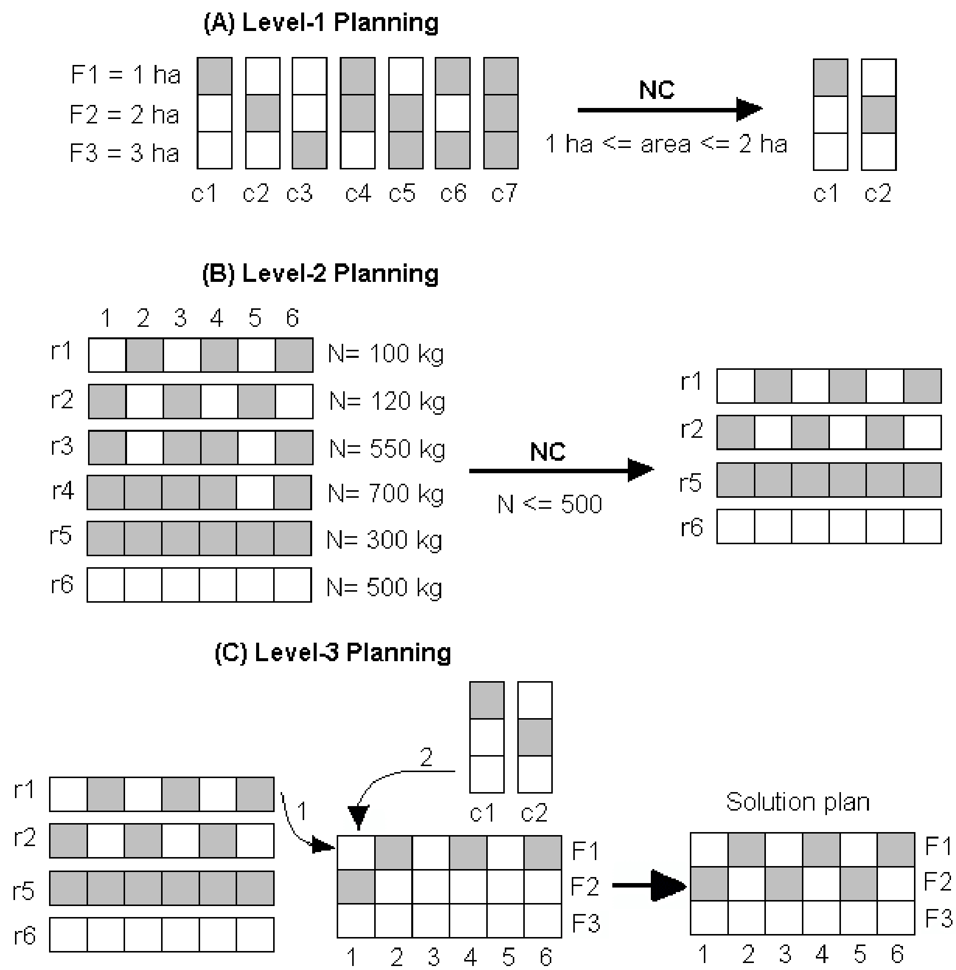

Figure 3-7 Three-level hierarchical search to solve the farm planning problem for a single target crop, a 3-field farm and a 6-year period. (A) In level-1 the planner creates and selects fields combinations to satisfy annual area and yield requirements for the target crop. The example shows combinations that include the target crop in the gray field boxes. Field combinations are selected to satisfy the annual target acreage of 1 to 2 ha. Node consistency (NC) is applied to make the selection. (B) In level-2 the planner addresses the selection of rotations by each field to meet preferences on soil erosion, nitrate leaching, and pesticide leaching and runoff. The example shows rotations that are selected for nitrogen (N) to meet the condition that nitrogen consumption be less than or equal to 500 kg. (C) In level-3, the planner performs a tree-search (TS) to assign rotations to fields and combinations to years to fill a field-year matrix. Rotation r1 is assigned to field 1 (arrow 1), and combinations are tested to match the crop in year 1. Combination c2 matches and is assigned to year 1 (arrow 2). The search continues until all the assignments are done. A solution is presented on the right.

The planner was written in object-oriented form; its class hierarchy is presented in Figure 3-8. The hierarchy structure uses a single inheritance approach; thus, all the classes are derived from the built-in ancestor **TObject**. The engine is an instance of the **TCSPSolver**, a descendent of **TBinaryConstraintManager**. A **TBinaryConstraintManager** contains lists of arcs, binary constraints, and nodes. Arcs are used to establish constraints among nodes.

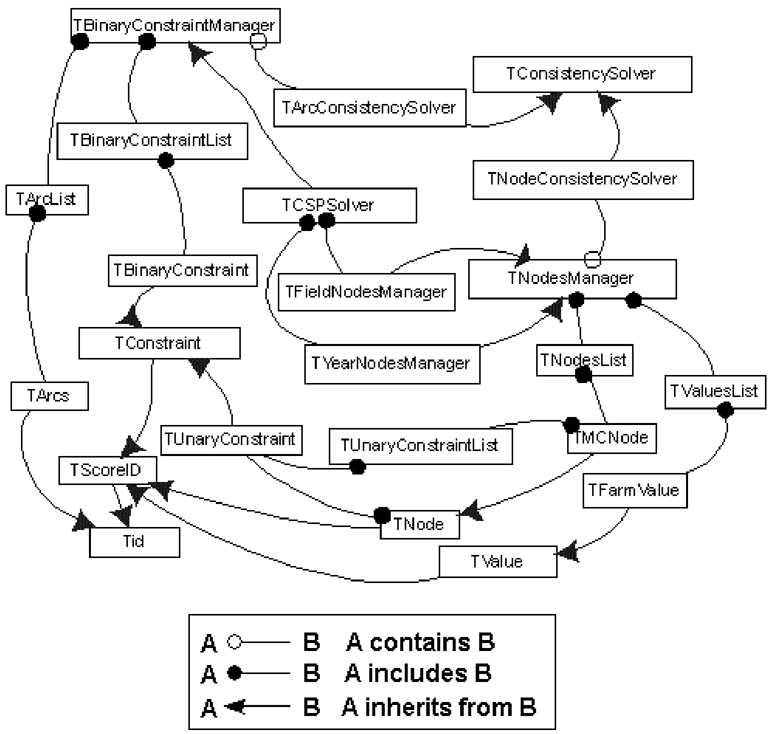

Figure 3-8 CROPS-LT planning engine class hierarchy. Legend shows the relations between classes based on Booch (1994). ***Contains*** **entails a short-term use of a server class by a client class. The** ***includes*** **relation means the inclusion of a class within another class. The** ***inherits*** **relation establishes a child-parent link, classes are derived from a parent and inherit its properties and behavior.**

The **TBinaryConstraintManager** uses an arc consistency solver to perform arc consistency. The **TCSPSolver** contains a **TFieldNodesManager** and a **TYearNodesManager**, both subclasses of the **TNodesManager**. The **TNodesManager** is a class that manages generic nodes and values. A single copy of the values is kept in a list by the **TNodesManager** and nodes contain pointers to this list to avoid duplication. The **TFieldNodesManager** manages field nodes and their rotations, while the **TYearNodesManager** manages target crops and combinations. Field and crop nodes are represented by the **TMCNode** class, which records assignment of unary constraints to nodes.

The **TConstraint** class encapsulates the behavior of a generic constraint, that is, it encapsulates a logical check method, and properties to contain **TValues** from the domain of the nodes. The **TUnaryConstraint** is a specialized subclass to handle unary constraints like soil erosion limits or annual target acreage ranges. Each **TMCNode** includes a list of **TUnaryConstraint** objects. Field constraints restrict nitrate leaching, soil erosion, and pesticide leaching and runoff. Buick et al. (1992) describe the computations of these variables. Crop constraints are defined by the planner to select annual target acreages and yield targets specified by the user. In addition to the constraints in CROPS (Stone, 1995), some new constraints were added (Table 3-2). All the new constraints restrict some numeric value to be less than or equal to a threshold. The corn rootworm (CRW) constraints are used to reduce the risk of CRW damage and are described in Chapter 5. The HEL percent constraint requires field combinations to have at least certain percentage of highly erodible land (HEL). The HEL percent constraint is enabled by the user to low-erosion crops, such as alfalfa or pasture. The field combination computes the percentage of HEL fields that exist in the given set of fields. For example if a field combination contains three fields with acreages: field-1 = 20 ha, field-4 = 20 ha, and field-7= 30 ha, and fields 1 and 7 are HEL, the percent of HEL fields in this combination is (50/70)100= 71.4 percent. If the planner applies this constraint, it will eliminate field combinations with percentages lower than 90 percent and this field combination would be rejected. The patch ratio constraint is used to limit the number of adjacent fields that contain a target crop.

Table 3-2 Unary constraints added to those defined by Stone (1995) for the farm planning problem. Level 1 constraints operate at the field level, level 2 constraints operate at the annual or farm-level.

Name	*Level applied*	*Description*
CRW-Field risk	2	Restricts rotations to those with a CRW risk below a threshold
HEL Percent	1	Restricts field combinations to have high HEL percent acreage
CRW-Combo Risk	1	Eliminates field combinations with a CRW risk above a threshold
Patch ratio	1	Eliminates field combinations with either too many or too few adjacent fields planted with the target crop

The **TNodeConsistencySolver** performs node consistency. It receives a **TMCNode** from the **TNodesManager** and applies NC to its values using the constraints owned by the node. The normal NC procedure removes values that do not satisfy the unary constraint. However, if after NC the domain is empty, no solution exists and the planner fails. The farm-planning problem is complex because it entails the satisfaction of multiple constraints; with each new constraint, the possibility of failure increases. The CROPS system included an automatic constraint relaxation scheme to handle this problem (Stone, 1995). However, in this project it was necessary to manipulate specific constraints to address specific problems. Consequently, two procedures were implemented in the **TNodeConsistencySolver** to allow the researcher to change the thresholds in and behavior of unary constraints. The *autorelax* procedure allows the planner to relax the constraint if the domain size of the node falls below a minimum count. The algorithm relaxes the constraint until the selected values are above a minimum acceptable number of values (*m* in Figure 3-9). The second procedure, *tight*, is the inverse of the *autorelax* algorithm. It tightens the constraint to eliminate more values, down to a minimum threshold (Figure 3-10). Each **TUnaryConstraint** can be set to use either procedure.

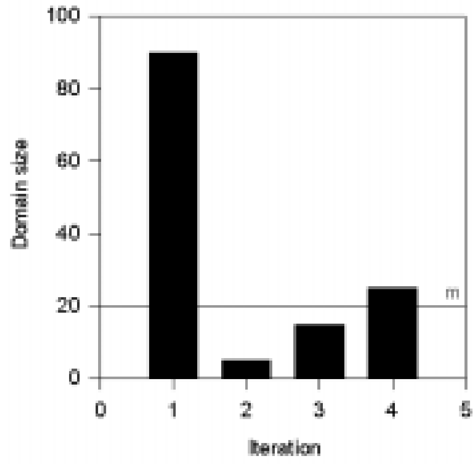

Figure 3-9 Domain size after each iteration of the *autorelax* algorithm operating on numerical constraints. The procedure works in a loop. If the domain falls below a minimum threshold (*m*) as in iteration 2, the procedure relaxes the constraint and iterates until the selected values are above *m*.

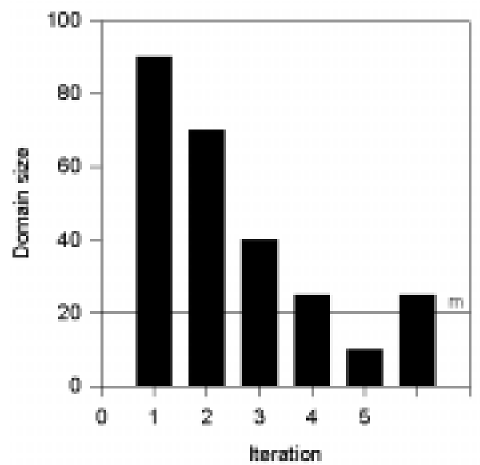

Figure 3-10 Domain size after each iteration of the tightening algorithm operating on numerical constraints. The constraint is iteratively tightened until the domain size falls below a minimum threshold (*m*) at which point (iteration 5) the previous domain is restored.

Constraints, nodes, and values (domain items) all inherit behavior from the **TScoreID** class. This class implements a numerical score property and a quick-sort algorithm (Weiss, 1994) sorts groups of these objects.

At run-time, the **TCSPSolver** creates the **TNodesManager** objects. When created, each object loads its required data from the databases. The **TYearCropManager** manages the first level search. It creates the crop nodes when target crops are selected. **TArc** objects for the crop-to-crop constraint are created by the **TCSPSolver** at this stage. The **TYearCropManager** loads the field combinations with annual target acreages between the absolute minimum and maximum acreage across all the target crops. Once the combinations are loaded, constraints are sorted by the nodes and node consistency (NC) is applied to select only those values that satisfy the unary constraints of each target crop. Once the constraints are applied, the **TCSPSolver** attempts to make all crops for a single year fully consistent using Mackworth's (1977) AC3 algorithm. Next, the **TYearCropManager** sorts the crop nodes by domain size and the field combinations by profit or yield if required. Once the sorting is done, the crop nodes are copied to nodes representing all the planning years.

Although it has been reported that full arc consistency is not efficient (Nadel, 1988, 1989), I implemented full AC because it has the advantage of quickly determining when no solutions exist. If there is no solution, the user can change the settings of the planner.

The second level search is handled by the **TFieldNodesManager** and the **TCSPSolver**. Each field node sorts its constraints and applies NC to rotations through a **TNodeConsistencySolver**. Field constraints are applied sequentially within each field. Relaxation or tightening is executed locally within each node. The next step is to make the fields and crops arc-consistent. The **TCSPSolver** creates **TArc** objects to represent binary constraints between field and crops, and a **TArcConsistencySolver** then applies full AC to the field and crop nodes. Once crops and fields have been made consistent, the **TFieldNodesManager** sorts fields by domain size and rotations can be sorted by profit, USLE C-value, pesticide leaching risk, or nitrate leaching risk. Sorting heuristics have been proposed to speed up the search procedure (Kumar, 1992; Sadeh and Fox, 1995). However, here they were also used to improve the quality of the plans. For example, sorting rotations by profit can make it more likely for the planner to find high profit plans quickly. If any fields are not included in valid field combinations after AC is established, the **TFieldNodesManager** excludes that field from level three planning. For example, in Figure 3-7, field three has an area of three ha and cannot participate in any solution because is not within the target range (1 to 2 ha). The rotation with the highest profit but including no target crop is assigned to this field.

Level three planning is executed by the **TCSPSolver** using the tree-search and partial look-future algorithms from Nadel (1989). For each plan found, the **TCSPSolver** restores any excluded field with its best rotation. The plan net profit is computed and compared against a benchmark plan profit. Only plans with a profit equal to or higher than a certain percent of the benchmark are accepted. If the plan is accepted, it is included in a list of valid plans and sorted by profit. The planner stops when a user-defined number of plans have been found or the allotted planning time has expired. Each plan is represented by a **TFarmPlan** object that contains id's for fields and rotations. An additional id is also required for rotations to indicate the year it starts to cycle. Plans are stored in the *FarmPlans* database table with additional indices to identify which planning run generated the plan (Figure 3-4).

Economic Model

The economic model is represented by the **TCropBudgetManager** and associated classes (Figure 3-11) and is responsible for most of the economic computations required by the system. The analyses are computed using a crop enterprise model (Boehjle and Eidman, 1984).

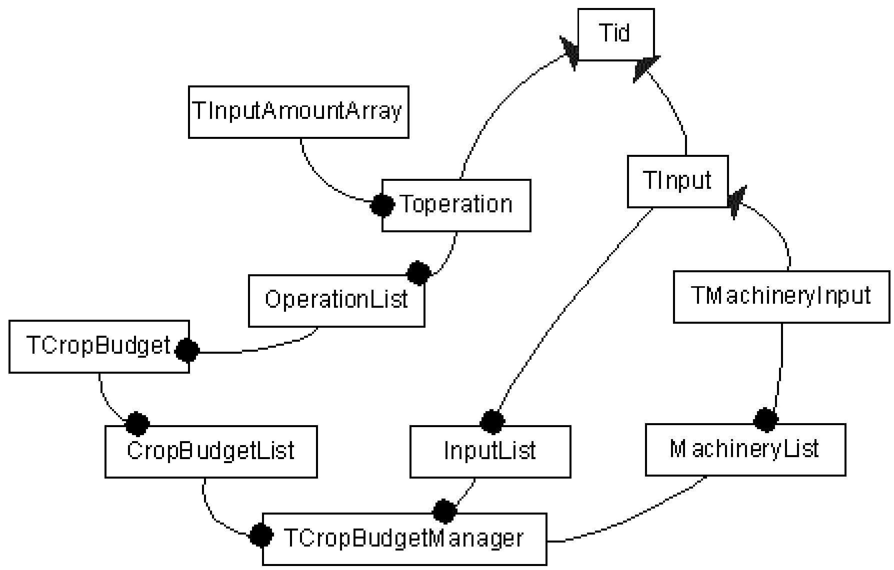

Figure 3-11 Class hierarchy for the crop budget manager. Boxes represent classes. Arrows establish child-parent relationships (*inheritance* relationship). Lines with black circles represent the inclusion of a class within another class (*has* relationship).

The budget manager has a list of **TCropBudget** objects and maintains unique copies of **TInput** objects divided into generic inputs and machinery inputs. Machinery is a special input because it has several costs associated with it: labor, fuel and oil, repair, and fixed costs. A crop budget is represented as a list of operations by the **TCropBudget** class. Each operation has an operation date and an array of input-amount pairs.

The **TCropBudgetManager** calculates profit at the field level and for a given rotation. The crops that compose a given rotation are selected from the databases; then, crop budgets for each crop are computed using yield and input requirements specific to the field. Profit is computed for each crop as return above variable plus fixed costs (see example in Table A-1, Appendix A). The profit for the rotation is the sum of profit of each of the crops that compose the rotation. Because a rotation is a sequence of several (possibly distinct) crops, the crop manager has to have access to data for all the crops in the system.

A **TCropBudgetManager** is used by the **TYearNodesManager** to compute the profit of each of the field combinations so it can sort field combinations by profit. Likewise, the

TFieldNodesManager uses a **TCropBudgetManager** to compute the profit across all the rotations when sorting domains. The **TCSPSolver** itself uses a **TCropBudgetManager** to determine the overall profit of each potential plan.

Visual Interface

The CROPS-LT program is the visual interface for the planning module. The main window has menu, button, and status bars. It also has a two-page component and a memo control in the bottom part (Figure 3-12). This program was written separately from the FLAME module to run independently.

The menu and button bars are used to create, reset, and destroy the planner. A planner can be instantiated for a given farm selected from a list box. Once the planner is created, the system displays the total area for the selected farm. The user can choose to view information for fields, target crops, rotation, and field combinations using the menu bar. The *Identifiers* page has controls allowing the user to set the file name for storing the results and to store the results to a database if desired. The interface also allows the user to set the solutions (plans) that the planning engine should find before stopping as well as an overall time limit for the planner.

The target crop input page allows the user to specify acceptable crop acreage range as annual targets for the planner (Figure 3-12C, H). The window maintains a list of all selected target crops and their acreage (Figure 3-12D). It also shows the total crop acreage assigned to each of the growing seasons based on the user’s targets. If the user exceeds the total crop acreage, a warning message is displayed.

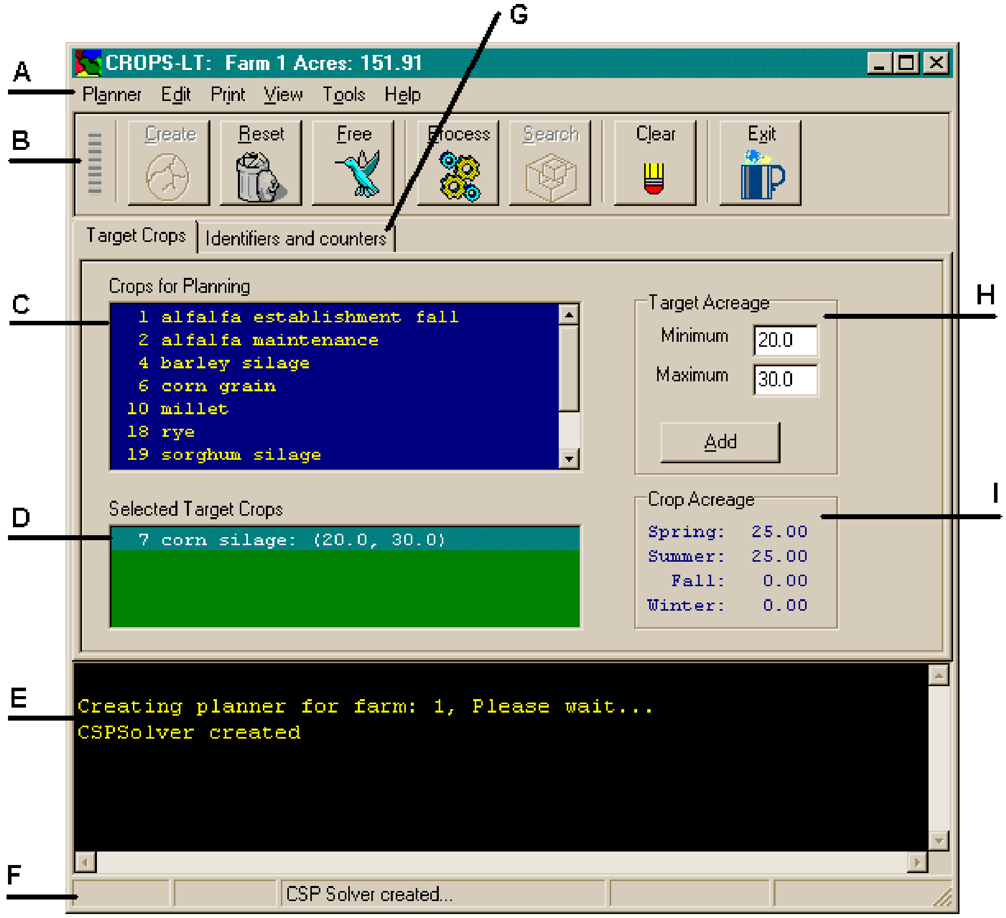

Figure 3-12 CROPS-LT Main window. (A) Menu bar. (B) Button bar. (C) Crops for planning. (D) Target crops with requested acreage. (E) Memo control. (F) Status bar. (G) Identifiers page. (H) Acreage selector. (I) Total acreage used by the target crops.

CROPS-LT contains three editors to edit fields and target crops: Active Fields Editor, Field Editor, and Target Crop Editor. The Active Fields Editor (Figure 3-13) is used to remove fields from planning if desired. This editor can be used, for example, to remove fields that are in permanent pasture to prevent the planner from putting them into crop production.

Figure 3-13 Active Fields Editor window. Fields from the left list can be removed from planning and put into the excluded list on the right.

The Field Editor (Figure 3-14) is used to change different field properties related to planning, for example, the minimum number of values to remain after performing node consistency (m in Figures 3-9 and 3-10). Some constraint properties, like consistency mode and rank, can be changed as well (Figure 3-14E). The users may choose to sort the rotations (domains) by different criteria: profit, C-factor from USLE, pesticide leaching risk and nitrate leaching risk or leave domains unsorted (Figure 3-14B). The user may also choose to sort the fields themselves by either domain size (default) or a user-defined field ranking (Figure 3-14C). The user can also activate or inactivate the field constraints, change the sorting order or change the node consistency mode if desired (Figure 3-14E).

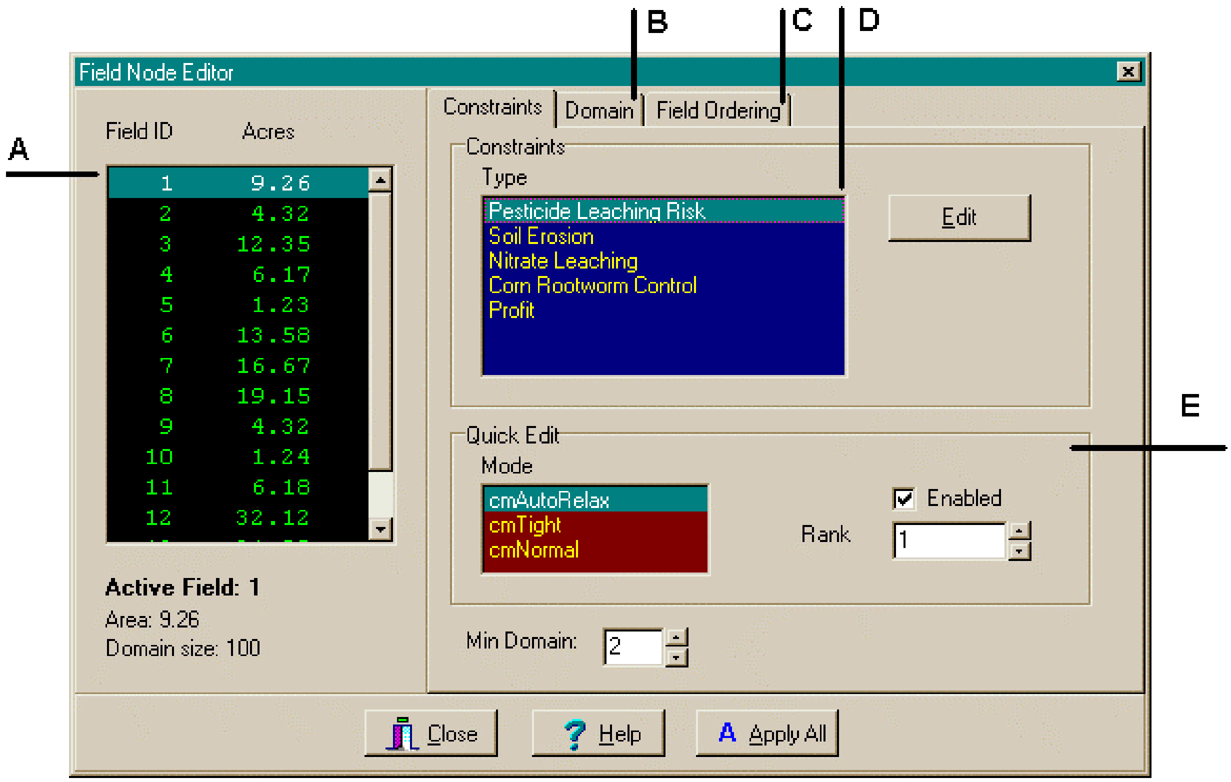

Figure 3-14 Field Editor window. (A) List of available fields for the current farm. (B) Domain sorting preferences page. (C) Field sorting preferences page. (D) Constraint list. (E) Quick constraint editor.

The Crop Editor (Figure 3-15) is used to edit the target crops. As for fields, the user can set a minimum acceptable number of target combinations after node consistency. The user can also enable or disable specific constraints for each crop; for example, the yield constraint can be enabled to restrict values within a range. The window displays the annual average yield and area for the selected target crop in the benchmark plan to serve as a guide. The user can specify sorting the combinations by profit, yield or at random.

At any time except during the third level of search, the user can reset the planner by pushing the *Reset* button (Figure 3-12). The reset action removes the target crop specifications, restores the default values for each field and replaces any removed field. After the target crops have been selected and reviewed, planning is initiated by clicking the *Process* button to apply the first and second planning steps.

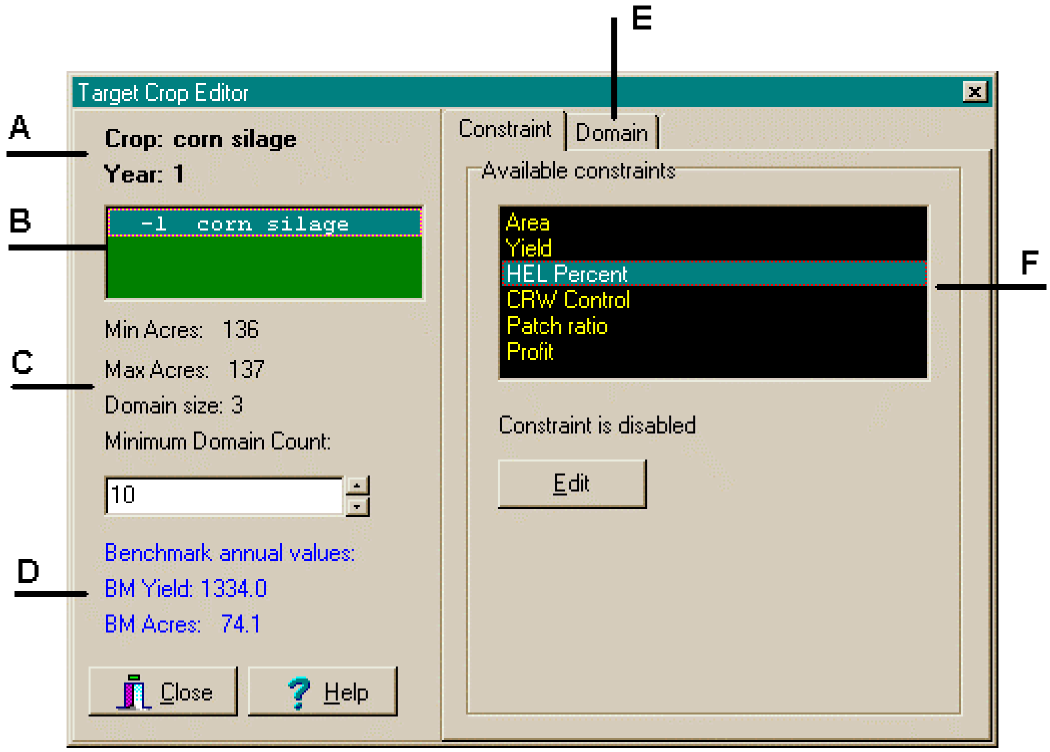

Figure 3-15 Target Crop Editor window. (A) Selected target crop. (B) Available target crops. (C) Target crop properties. (D) Yield and acreage for the selected target crop in the benchmark plan. (E) Domain sorting page. (F) Constraints for the selected target crop.

After planning steps 1 and 2 are completed successfully, the user can click on the *Search* button to launch the final search for whole-plans. During the search, information about the planning process is displayed in the status bar at the bottom (Figure 3-12F). The status bar displays the remaining planning time and the number of acceptable plans found. Planning finishes when either the planning time expires or the planner has found the number of plans requested. The system displays messages at the window's footer to let the user know about the status of the system (Figure 3-12F).

NPS Module

The components of the NPS module are ANSWERS and CLIGEN. ANSWERS was selected as the NPS model because it has several advantages over other models. ANSWERS is a continuous, distributed-parameter model, and can simulate a single tract during long times (years). A tract is defined as a set of contiguous fields. Field-level models like GLEAMS and EPIC were discarded because they cannot simulate variations among fields or even within a single field. Other distributed models, such as AGNPS and EUROSEM are event-based and therefore not appropriate for long-term simulation. WEPP is a distributed model like ANSWERS but was not yet available in an appropriate form when this research began. Finally, ANSWERS does not require calibration of parameters against observed data (Bouraoui, 1994).

ANSWERS represents the simulated landscape as a rectangular grid of homogeneous square cells. The cell size can be set to achieve the desired resolution. To simulate a farm of 100 ha, for example, a cell length of 0.5 ha would be appropriate (Bouraoui, 1994). However, the current version allows the management of 2000 cells only. Another restrictions are the high data requirements and the running time, which increases as the cell number increases (Bouraoui, 1994). Detachment, transport, and deposition of soil particles are simulated at the cell level. ANSWERS conveys sediment and nutrient losses to a channel network that discharges at discrete outlets. Lumped models like EPIC and GLEAMS provide only detachment of soil particles. Bouraoui (1994) describes ANSWERS in detail. Here, only the data requirements and data flow are described.

The ANSWERS data flow for a single run is illustrated in Figure 3-16. The model uses three input files for a single tract: *weather.inp*, *answers.inp*, and a *fertilizer.inp* (optional). Because ANSWERS simulates a single tract per run, if a farm is composed of several tracts, ANSWERS will need as many sets of input files as tracts are in a farm.

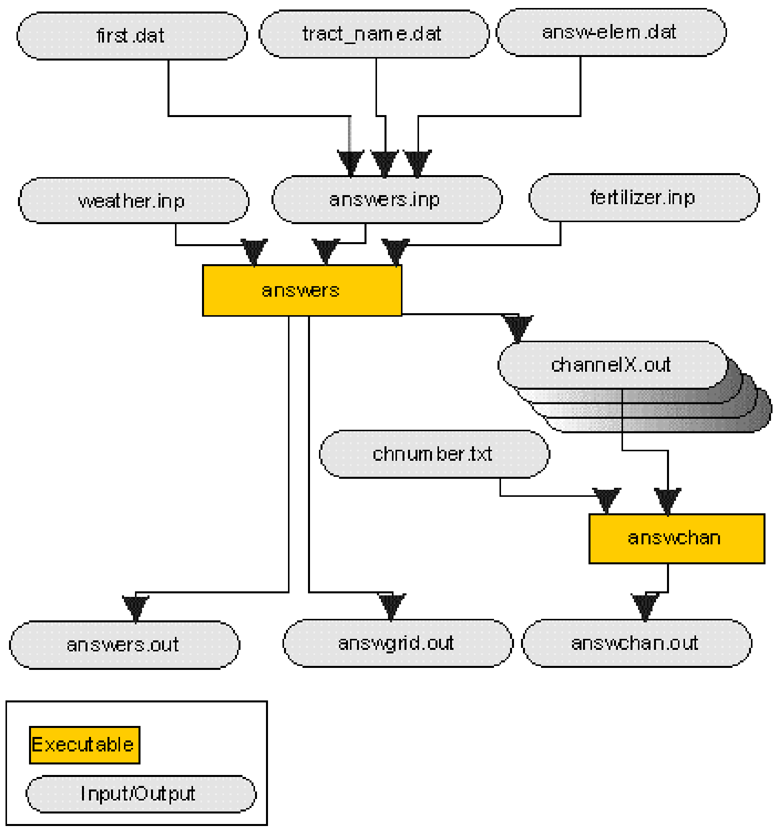

Figure 3-16 ANSWERS components and data flow. Input files are depicted as rounded rectangles and executable files as square rectangles. Input and output files are required for a single run and a single tract.

The *weather.inp* file contains temperature, solar radiation, and rainfall data. Rainfall is organized in breakpoints (Bouraoui, 1994). Breakpoint data is represented by cumulative time from the beginning of a storm and associated mean rainfall intensity between successive time intervals (Nicks et al. 1995). ANSWERS allows up to four types of storm patterns for a single cell but one is generally used (Bouraoui, 1994).

The *fertilizer.inp* file contains information about fertilizer applications: application date, crop receiving the fertilizer, type of fertilizer (nitrogen, phosphorus, and potassium), and amount applied. ANSWERS simulates the application of fertilizer at each date for the receptor crop in the cells it is "planted". Fertilizers enter the nutrient transformation and transport models (Bouraoui, 1994).

The *answers.inp* is a composite of three files generated separately: A header file (*first.dat*), crop and management practices data (*tract_name.dat*), and cell data (*ans-elem.dat*). The header file contains the name of the tract, simulation period, and soil types and associated properties. Soil properties include, among others, particle size, porosity, field capacity, organic matter content, and the K value from the USLE model (Bouraoui, 1994). Soil data can be extracted from the *DBSOILS5* database. The *tract_name.dat* contains information related to crop parameters required by the model (Figure 3-3A), number of fields and the crop rotation sequence for each field in the tract. Crop parameters include planting and harvest dates, leaf area index, yield, and canopy and bare area cover. This file and the simulation period in the header file vary according to the rotations in a farm plan. The *answ-elem.dat* file contains information for each of the cells: slope, flow direction, soil type, and basin and field identifiers.

The output files from ANSWERS are *answers.out*, *answgrid.out*, and one or more *channelX.out* files. The *answers.out* file has the same information as the *answgrid.out* file, that is, soil detachment and nutrient losses at the cell level, but the *answgrid.out* file has a format suitable to be directly loaded to a table. The *channelX.out* files contain soil deposition and nutrient losses at each of the basin outlets. In WFP, the *channelX.out* files are grouped into a single file (*anschan.out*) to simplify operations. The *chnumber.txt* file contains the number of channels (basins) for the tract. The results from each simulation are stored in two tables, one for cell-level results and the other for basin-level results (Figure 3-3B). ANSWERS is called manually via a TELNET connection to its host machine (Figure 3-1).

Some modifications were made to ANSWERS for this project. The original *ANSWERS.OUT* file was renamed *answgrid.out* because it had the same name as the *answers.out* but in upper case. This would have created conflicts in the Windows environment. Another modification was to increase the size of the arrays from 20 to 99 to allow for more fertilizer applications.

The weather data are obtained from the CLIGEN module. Components and data flow for CLIGEN are displayed in Figure 3-17, showing files required as input or generated as output for each tract simulated. The *stations* file is a database containing a list of weather stations in the U.S. The *state.lst* file contains the parameters required by CLIGEN to run, listed for each of the stations in the *stations* file. The *cligen.dat* file contains the user-defined requirements: starting date, simulation period, station id, and a seed number to start the generation of random numbers. CLIGEN uses the station id in the *cligen.dat* file to read the parameters from the *state.lst* file.

The output file from CLIGEN is not directly used by ANSWERS because it gives values on a daily basis and ANSWERS requires rainfall as breakpoints. The executable programs (bp_*) that process CLIGEN's output generate the final *weather.inp*.

The CLIGEN program was modified to include a new random number generator. The previous random number generator was changed because it did not allow setting an initial (seed) value. Initial values are needed to repeat the same stream of random numbers when statistical analysis is required (Law and Kelton, 1991). CLIGEN's random number generator was replaced with UNIRAN, which allows the control of stream numbers and has been thoroughly tested (Marse and Roberts, 1983). The CLIGEN module programs are run from FLAME by calling Windows application programming interface (API) functions.

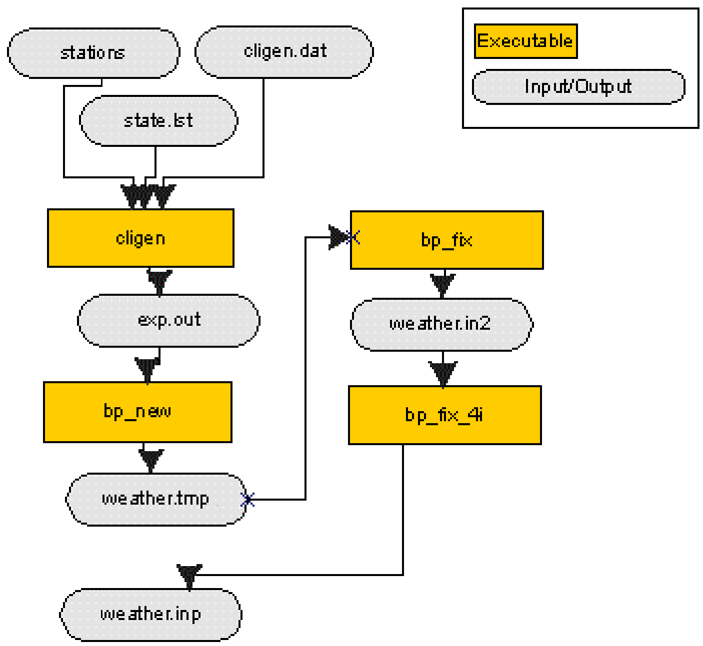

Figure 3-17 CLIGEN components and data flow. Input files are depicted as rounded rectangles and executable files as square rectangles. Input and output files are required for a single run and a single tract.

Utility Programs Module

This module is composed of two programs: FTP and TELNET. Commercial programs were used for this project. An FTP program is used to send and retrieve files between computers with different operating systems (Figure 3-1). A TELNET connection is normally required to run ANSWERS from its host machine.

FLAME: The Farm-Level Management Module

FLAME was developed as part of this research and is composed of an executable file and a dynamic link library (flameDlg.dll). FLAME is used to manage data, generate input files and process output files, display graphics, and call the CLIGEN programs (Figure 3-17). In particular the tasks that FLAME executes are the following:

- Management of databases required by the whole system
- Display of topographic and physical properties at cell, field, soil, basin, and tract levels

- Display of rotation plans
- Display of crop sequence at farm level
- Display of sediment and nutrient loses estimated by the simulation model
- Display of general information in XY coordinates
- Creation of the *tract-name.dat* and *fertilizer.inp* input file for ANSWERS
- Creation of the *cligen.dat* input file for CLIGEN
- Running the CLIGEN module programs.

FLAME uses the parent-child model proposed by Microsoft to build programs for the Windows operating system, following guidelines published by Pacheco and Teixeira (1996) and Ambler (1998). The child windows are shown in Figure 3-18 enclosed within the main window. In addition, floating and dialog windows can be displayed from the menu or button bars. A description of these windows follows.

The main window (Figure 3-18) contains the menu, button, status bar, and child windows. The menu bar provides services to create child windows, copy graphics to the clipboard, call database editors, or call the crop budget explorer to perform economic analysis. Input files for CLIGEN and ANSWERS are created from the menu bar through dialog windows. The button bar has buttons to create new child windows, save graphics to clipboard, and print graphics. The status bar displays information about the state of the system.

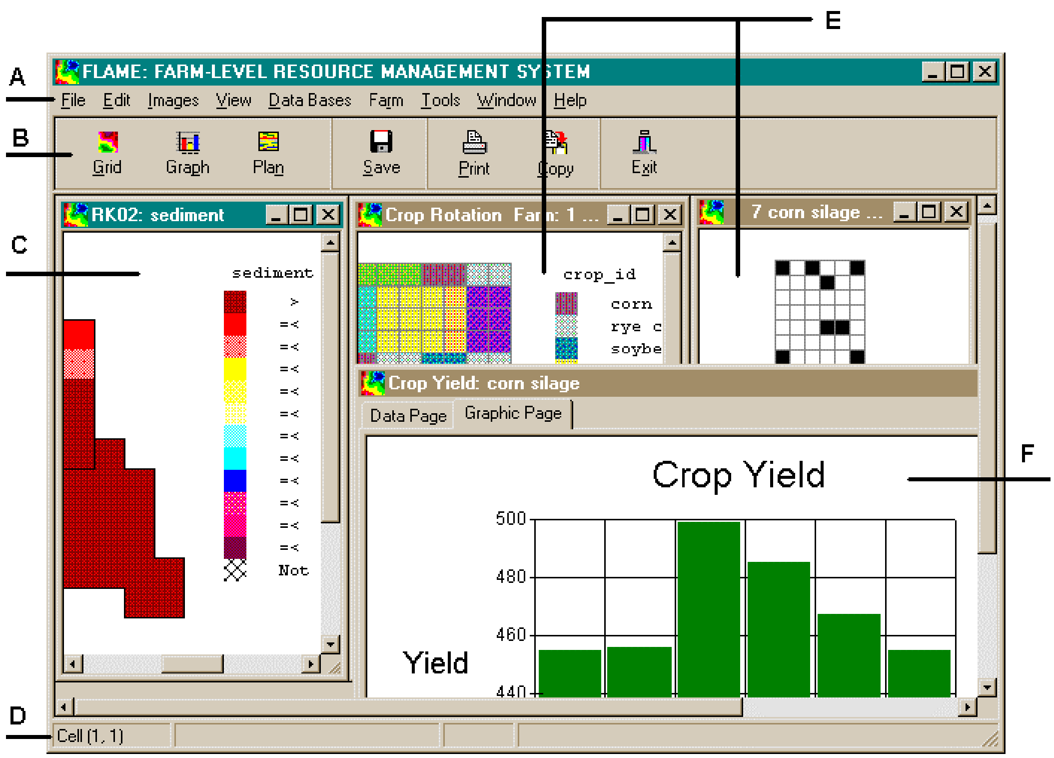

Figure 3-18 FLAME main window. User bars: Menu bar (A), Button bar (B), and Status bar (D). Child Windows: Grid window (C), Plan windows (E), and Graph window (F).

There are three types of child windows: grid, graph, and plan (Figure 3-18). The grid window displays information using a tract as the display unit, for example, biophysical properties like slope, or simulation results, such as sediment loss (Figure 3-19). The graph window displays generic data in an X, Y format (Figure 3-20). The plan window shows the distribution of crops of a selected plan across the planning years and fields for a single tract (Figure 3-21). The windows can have multiple instances and are created at run time to optimize memory.

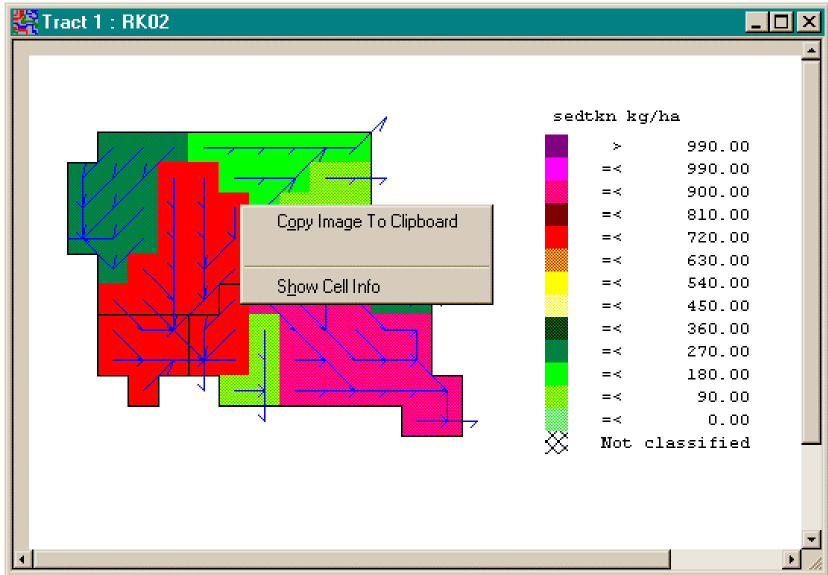

Figure 3-19 Grid window showing total Kjeldahl sediment-adsorbed (TKN) nitrogen values per basin (colored cells) and flow direction (blue arrows). The floating menu in the center is used to copy the graphic to the clipboard or to show cell-level information.

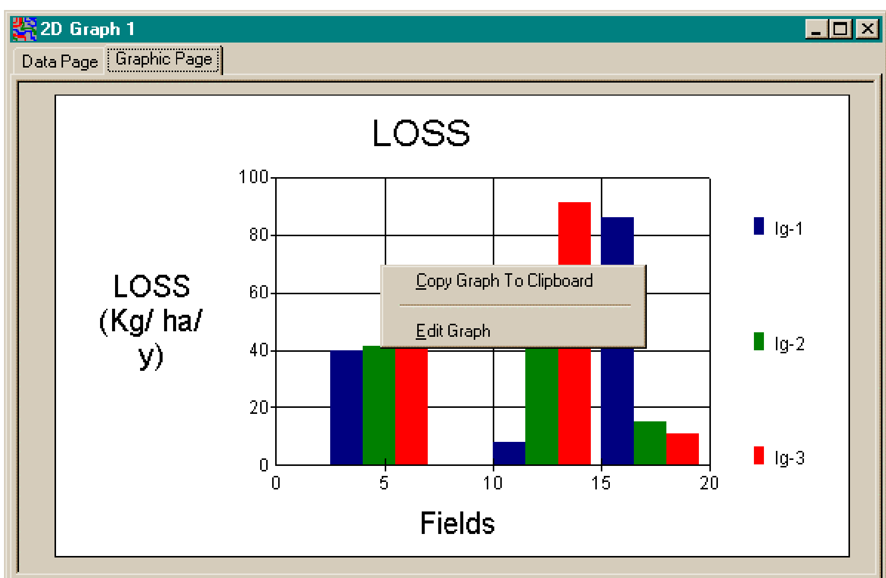

Figure 3-20 Graph window. This window includes a (hidden) spreadsheet-like data page and the graphic page. The data page is used to input numeric data. The graphic page displays the data using predefined types (column, line, bar, pie, and scatter). A floating menu allows to copy or to change the type of the chart.

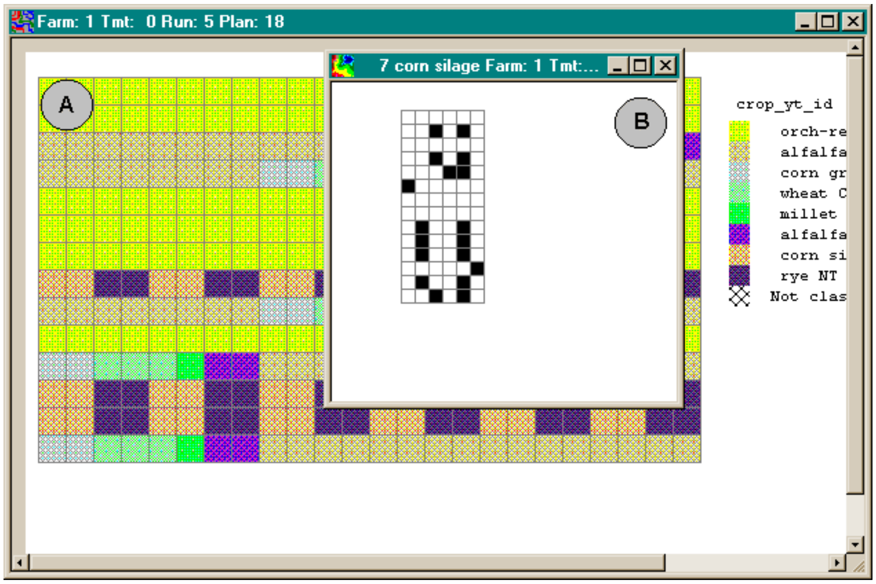

Figure 3-21 Plan windows represent the sequence of crops of a plan for a single tract. Rows in the grids correspond to fields ordered from top to bottom. (A) Sequence of crops per field. Columns correspond to seasons (6 years × 4 seasons = 24 columns). (B) Single crop sequence per field. Columns correspond to years only.

Floating windows can "float" on top of other windows on the screen. The floating windows are used to edit tables: crops, crop budgets, and farm information. Figure 3-22 shows a floating window to edit farm information. Floating windows also include the Crop Budget Explorer and the SQL manager.

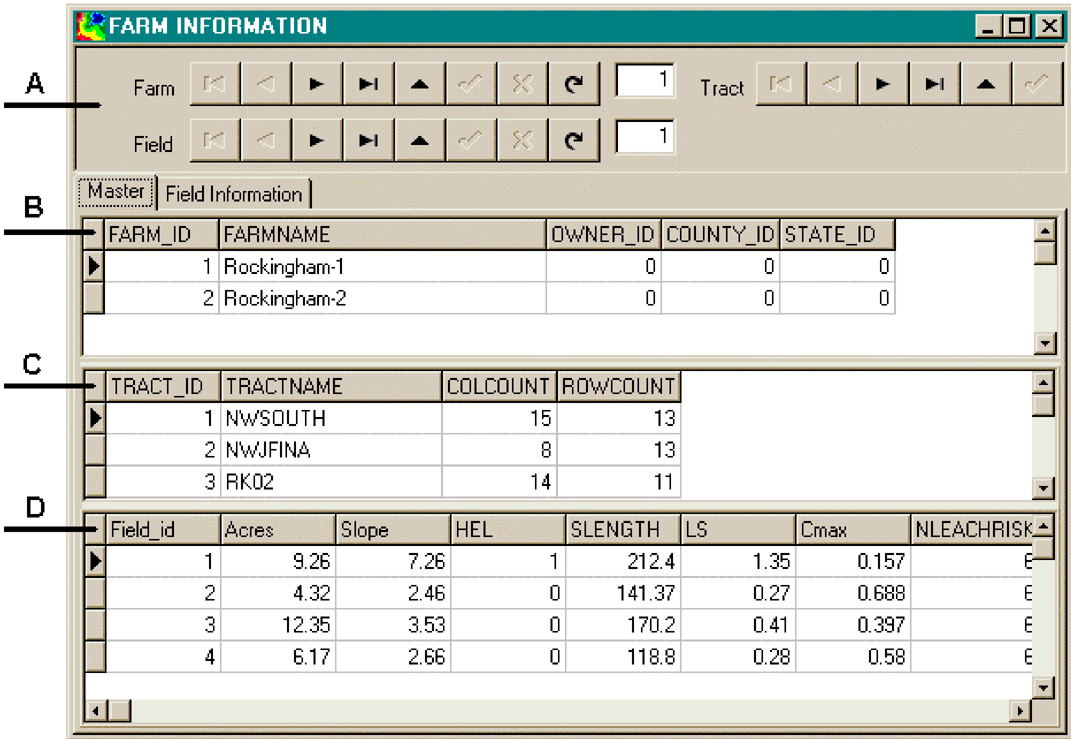

Figure 3-22 Floating window showing farm information. (A) navigator buttons. (B) Farm-level information. (C) Tract-level information. (D) Field-level information. By clicking on the navigator buttons, the user can select different farms, tracts or fields. The second page (hidden) presents more information at field-level. The tables are arranged in a master-detail way. Selecting a farm automatically selects its tracts and fields.

The Crop Budget Explorer (CBE) (Figure 3-23) is another floating window. It allows the user to perform economic analysis for specific fields. This window encapsulates the behavior of the **TCropBudgetManager**. Data on farms stored in the *fields* table of the *DBFARMS* database (Figure 3-4) are displayed in a selection box, subdivided into tracts and fields. A list of available crop budgets is presented below (Figure 3-23C). The user can match a crop with a field to compute a crop budget. The result is shown in the second page of the window (Figure 3-23D). The user can also compute crop budgets for an arbitrary field size, crop yield, and annual percentage rate (APR). The APR is needed to compute the production interest cost, which is an opportunity cost on capital invested in the variable costs (Kay, 1986).

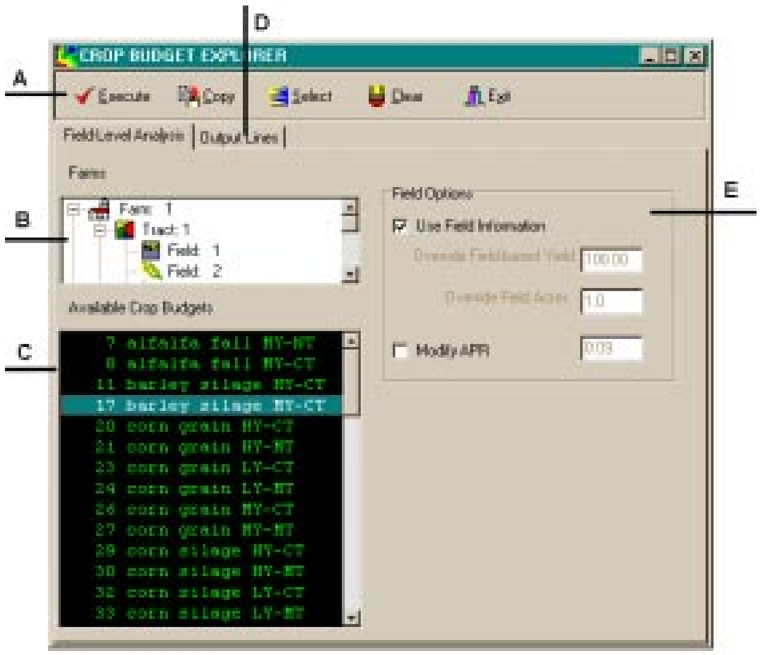

Figure 3-23 Crop Budget Explorer window. (A) Button bar. (B) Field selector. (C) Crop budget selector. (D) Output page. (E) Crop budget options.

The SQL Manager is another floating window as well (Figure 3-24). It is the central component for database management. Two multiple pages controls exist inside the form. On the left side is the database management selector (Figure 3-24D). It contains four pages. The first displays the available databases (Figure 3-24B) and their corresponding tables (Figure 3-24C). The second page is a system directory. The third page allows the user to change the settings for SQL statements (browse mode, edit mode, and EXE mode). The browse mode is used to retrieve read-only records and the edit mode allows the user to edit the retrieved records. The EXE mode is used to launch SQL statements to create or modify tables, or to create or drop indexes. The fourth page allows the user to empty the content of a table. The second page control is on the right and contains four pages as well (Figure 3-24E, F, G, and H). The first page is used to write SQL statements. The second page displays the results of the SQL query. The third page allows the user to write and save notes. The fourth page displays a table if it is selected from the table list on the left side. The user can launch commands from the button bar or a floating menu (Figure 3-24A, I).

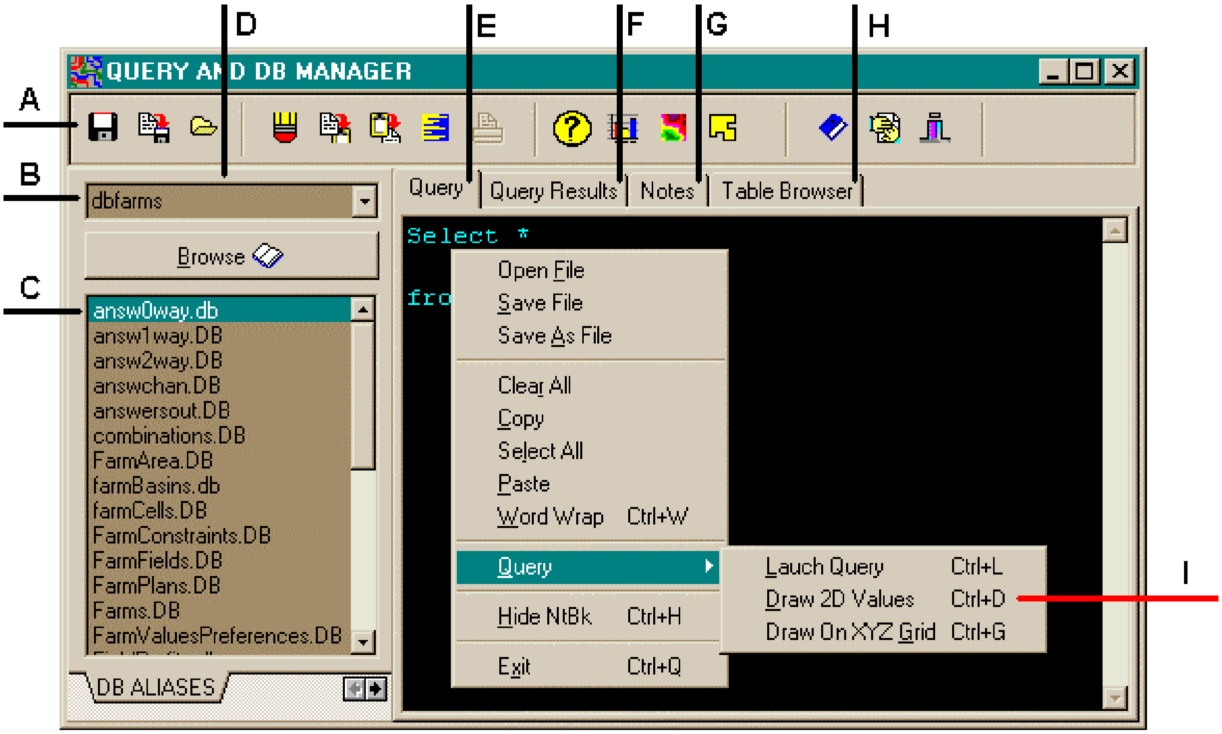

Figure 3-24 The SQL Manager window. (A) Button bar. (B) Database selector. (C) Table selector. (D) Database management selector. (E) Query editor. (F) Query result page. (G) File editor. (H) Table browser. (I) Floating menu.

Dialog windows are used to obtain input from the user to perform specific tasks. Dialog windows are described in Table 3-3. Other dialog windows not described but included in the system are standard window boxes to print, select system colors, and open/save files. The user must close a dialog window before he/ she can work with other parts of the program.

Table 3-3 FLAME dialog window description.

Window	*Use*
About	Displays version, copyright, and memory.
ANSWERS input editor	Allows users to select the tract to create the *fertilizer.inp* and *the tract-name.dat* files for a given plan.
CLIGEN input editor	Creates the *cligen.dat* file.
Tract selector	Selects the active tract for the grid window.
Graphics editor	Lets users edit graph properties (graph type, active series, legends) for the graph window.
Cell information	Displays topographic information for the current cell when selected from the grid window.
Append ANSWERS output file	Appends the output files from ANSWERS (*answgrid.out* and *answchan.out*) to tables in the database.
ASCII to table	Transfers data in ASCII format to any table following a predetermined schema (Pacheco and Teixeira, 1996). The schema lists the name of items in the text file, the type of data (integer or float for example), the size of the string data (10 characters for example) and the format for numeric data.
Display ANSWERS results	Lets user select resolution level (cell, field, soil, basin, tract) and result type (soil loss or nutrient loss) to be displayed in the grid window for the active tract.
ANSWERS summaries	After storing the output ANSWERS files in tables, processes the cell-level results to obtain averages for field, basin, and tract levels.
Database editor	Allows user to edit tables across all the registered databases.

The *cligen.dat* file is created from a dialog window. The user has to specify the starting date, simulation period, station id, and a seed number to start the generation of random numbers. The *tract_name.dat* and *fertilizer.inp* input files for ANSWERS are created from a dialog window as well. The *tract_name.dat* file contains a list of crop parameters and the crop sequence per field. The *fertilizer.inp* file contains a list of fertilizer applications, as mentioned in the NPS module section. To create these files, the user selects a plan from a list of available plans. Farm plans are stored and extracted from the *FarmPlans* entity (Figure 3-4). After a plan is selected, the crops that compose the rotations for the plan are selected. Next, the crop sequence for all the rotations is extracted. Crops and crop sequence are extracted from the *year_rotations* table (Figure 3-2). For each crop, planting and harvest dates, and fertilizer requirements are extracted from the crop information database (Figure 3-2). The crop parameters are extracted from the crop information database as well (Figure 3-3A). Finally, the files are created in the appropriate format for ANSWERS. FLAME creates input files per tract, as required by ANSWERS; therefore, to obtain a full evaluation of a plan the user has to create as many input file sets as tracts are in a farm.

Assumptions of the System

The system was written with the following assumptions and limitations:

- The system provides crop production plans to support dairy farm enterprises. The reason is that the crop network, from which the rotations were created, is composed mostly of crops that provide silage or hay to support dairy operations (Buick et al. 1992; Stone et al. 1992).
- Acceptable plans proposed by the system are evaluated in terms of returns above variable plus fixed costs across all the crop production enterprises.
- Generated plans are evaluated in comparison to a benchmark (control) plan that must be entered for each farm.
- Evaluation of profit, environmental risks, and simulated pollution are valid only in relation to other plans and to the benchmark. They are not intended to be used as absolute estimates of profit, soil or nutrient losses.

Operation of the System

For completeness, this section describes the general operation of the system; the operation of particular components was described previously. A flowchart describing a typical session is presented in Figure 3-25. The first step is to populate the databases with data. The user can input the data by opening the database editor (Table 3-3). Next, the user can either look for plans by using the CROPS-LT component or compute crop budgets, edit tables, launch queries, etc. (right side box in Figure 3-25). Once the planner has generated at least one acceptable plan, the next step is to assess the plan using ANSWERS. FLAME provides the interface (dialog windows) to create the required input files for CLIGEN and ANSWERS, as described before. The results of the simulation can then be stored in databases, and the user can display the results on the grid, graph or plan windows.

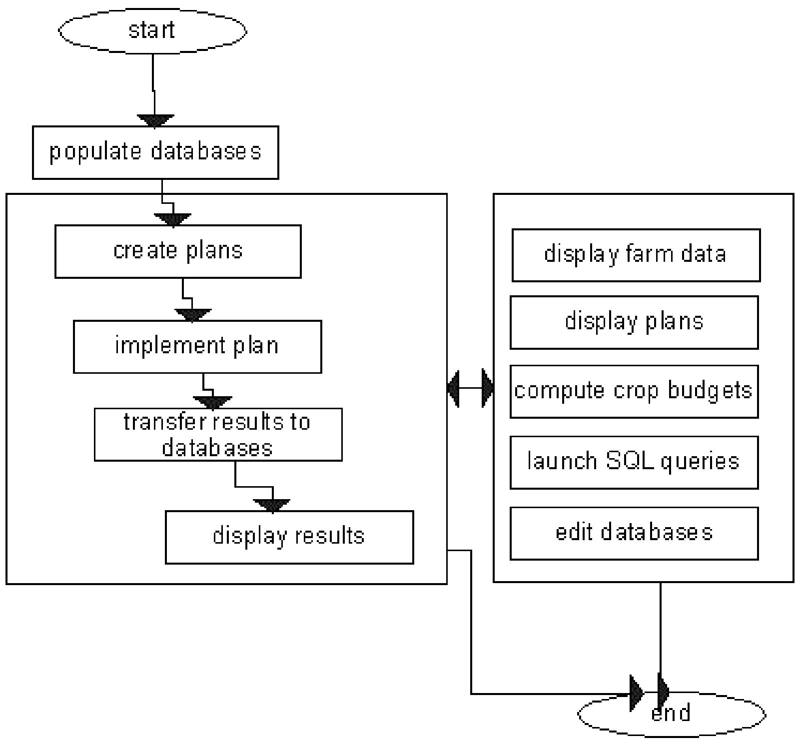

Figure 3-25 Operation sequence of the whole-farm planning decision support system. Once the databases are populated with data, the user can generate plans with planner and display results (left box) or display farm data, compute crop budgets or edit databases (right box).

Verification and Performance of the System

Some system components, such as CLIGEN and ANSWERS, have been verified and validated elsewhere (Nicks et al. 1995; Bouraoui, 1994). Each piece of the CROPS-LT planning engine was coded independently and tested separately. The arc consistency algorithm, AC3 (Mackworth, 1977) and tree search algorithms were tested with an example provided by Mackworth (1987). The unary constraints and consistency solvers were tested to verify the constraint modes and the removal of values. The tree search algorithm has been extensively verified by checking the generated plans with the corresponding annual target acreage across all the planning years. The *Crop Budget Explorer* and *TCropBudgetManager* were tested with a set of crop budget enterprises published by the Virginia Cooperative Extension (VCE, 1997).

Russell and Norvig (1995) mention four criteria to evaluate searching strategies, specifically, completeness, time complexity, space complexity and optimality. Completeness refers to the ability of the strategy to find a solution when it exists. Time complexity evaluates the time the algorithm takes to find a solution. Space complexity evaluates the memory requirements. Optimality refers to the ability of the strategy to find the best solution when there are many available. The general backtracking algorithm is considered complete for CSP problems so the first criterion, completeness, is satisfied (Tsang, 1993; Russell and Norvig, 1995). However, for practical purposes, the algorithm is not complete because planning is usually performed for a limited time, and the whole search space is not sampled. This is the main reason for implementing domain and node-ordering heuristics, to improve the chance of obtaining good plans. Specifically, domain ordering and node exclusion are included in the planning system to address the problem of optimality under a limited planning time constraint. These heuristics are evaluated in Chapter 4. The time complexity of node and arc consistency algorithms, including the ones used in the current planner, is known to be polynomial (Mackworth and Freuder, 1985).

The time complexity for backtracking is exponential, which means that the time to explore the whole search space increases in an exponential way as the number of variables (nodes) increases (Mackworth, 1977; Tsang, 1993). The space complexity of the arc consistency used here (AC3) and backtracking is linear so the memory requirements are modest (Tsang, 1993; Russell and Norvig, 1995). Finally, it is appropriate to mention that during the tree search, when a plan is found, it is stored in a list of accepted plans only if it is different from the previous ones. In addition, when a plan is added to the list of unique plans, it is inserted in a position relative to its profit; the plan at the tail of the list is thus removed. Therefore, if the planner examines all the search space, the selected plans are optimal in terms of profit and distinct from each other.

Chapter 4 System Configuration Analysis

Introduction

Constraint satisfaction techniques have been used to search large potential solution spaces. The techniques are divided into preprocessing methods and backtracking methods (Nadel, 1989). Preprocessing methods, such as arc consistency, search for and eliminate values in the domain of the variables that are not involved in any solution. That is, they try to reduce the size of the search space by eliminating values. Backtracking methods are used to avoid redundancy during a search through a tree-like search space. The methods are described in Chapter 3 in the section about the planning engine. A third technique is to combine both methods to improve the planning process. The efficiency of these algorithms has been well studied. For example, Tsang (1993) has compiled the efficiency of arc consistency algorithms. The analysis of efficiency may consider the worst case scenario as well as the average (Mackworth, 1977). For mixed backtracking algorithms, numerical analysis is required because analytical methods are not available (Haralick and Elliot, 1980; Nadel, 1989). For example, Nadel (1989) examined the efficiency of several hybrid algorithms to solve the n-queens problem, referring to placing n-queens on a chessboard of n×n spaces so that no queen attacks any other. The n–queens problem is a common problem in the artificial intelligence community (Tsang, 1993; Russell and Norvig, 1995). However, some authors have remarked that efficiency might be problem-dependent, and the n-queens problem may not provide enough insight into how appropriate a method is (Nadel, 1989; Tsang, 1993). In addition, domain-specific factors can greatly affect the performance of these algorithms. For example, domain-based heuristics can often speed up the searching process or improve the quality of the solutions found (Kumar, 1992; Tsang, 1993).

The current planning system includes some heuristic algorithms to improve the quality of the solutions. These heuristics are domain ordering, to increase the chance of quickly obtaining good plans, and field exclusion, used to increase the profit and reduce the search space. This chapter presents empirical numerical analyses of the heuristic ordering (sorting) algorithms included in the planning system. The effect of target and range acreage on the time the planner takes to find the first plan was studied as well. The annual target acreage is set by the user as a range. Therefore, it is important to know how target and range acreage affects the planning process. The planner's feasibility and flexibility to plan for different target crops were also subjectively evaluated.

Experimental Conditions

The analysis was applied to a farm, named farm-1, composed of 14 fields divided into 3 tracts. A tract is defined here as a set of contiguous fields. The farm is located in Rockingham County, Virginia, and contains about 62 ha of cropland with four fields classified as highly erodible land (HEL) and requiring soil erosion control practices. Data describing the farm's field properties were extracted using the FARMSCALE system (Wolfe et al. 1995) from digital SOILS5 and topographic data from Rockingham County, Virginia. (Goran, 1983). Computations of pesticide leaching risk, pesticide runoff risk, nitrate leaching risk, and Cmax were performed as described by Buick et al. (1992). Cmax is a field value such that a crop with an USLE-C less than Cmax would produce soil erosion in that field with a value less than T, which is the soil formation rate (Buick et al. 1992). Crop yield data were calculated based on the VALUES database (Simpson et al. 1993). The field properties of the farm are included in Table A-2. A further description of the farm is presented in Chapter 6. One hundred crop rotations were used for planning, a subset of those used by CROPS deemed appropriate for dairy farms in the Shenandoah Valley of Virginia. The rotations were generated using a crop rotation network based on the most common cropping practices in Virginia (Buick et al. 1992; Stone et al. 1992). The distribution of crops across the rotations is presented in Table A-3 (Appendix A). Crop budgets were prepared from information obtained from the Virginia Cooperative Extension (VCE, 1997). All runs of the planner were executed on a PC with a 233 MHz AMD central processor unit, 64Mb of RAM and a 3 Gb hard disk.

Efficiency of Domain Ordering

Ordering the possible values of a domain before the tree-search can have strong impacts on the efficiency of a CSP planner (Nadel, 1989; Stone, 1995). Here three heuristic ordering methods and a control were compared based on their effects on the speed of the planner and the profitability of plans found in a limited time. Three configurations and a control were tested with some selected crops across different annual target acreages. In the control (Ctrl), both field domains (rotations) and target crop domains (field combinations) were randomly ordered. The first treatment was to plan with the rotations sorted and the field combinations randomized (R). The second configuration was to plan with the field combinations ordered and the rotations randomized (C). The third configuration was to plan with both rotations and field combinations ordered (RC). For rotations, the field domains were sorted by profit, Cmax, pesticide leaching risk, and nitrate leaching risk. For field combinations, the domains were sorted based on profit. I chose profit as the sorting criterion because economics is an essential constraint for accepting plans once all other constraints have been satisfied. Profit for rotations and field combinations was computed as gross returns minus variable and fixed costs across all the crops in a single rotation or across all the target fields for a field combination. The quick-sort algorithm was used to sort the domains (Weiss, 1994). Two sets of experiments were conducted. The first one examined how quickly plans were found and how quickly the higher profit plans were discovered. The second set of experiments was conducted for a limited time. It had the objective of determining the best sorting strategy for different target crops. In both sets of experiments, plans were accepted regardless of their profitability, i.e., the overall economic constraint was disabled.

Planning to Obtain All Plans for Domain Ordering

To determine the efficiency of the search algorithms using the different domain-ordering strategies, the searches were run to completion, meaning that all possible plans were found. This required limiting the farm size to seven fields (from 14 total) so that planning time was reasonable (within a 24-hour limit). Preliminary runs were used to determine an appropriate number of fields and to choose limits for the domain sizes for field rotations and target crop field combinations. Corn silage was chosen as the target. Six annual target acreage ranges were tested for plans (Table 4-1) each with a target range of 2.02 ha. The target range is the interval the target acreage can vary, for example, if the target acreage for crop A is 10 ha and the target range is 2 ha, then crop A must be planted between 9 and 11 ha to comply with the target requirements.

Table 4-1 Annual target acreage as percent of total farm area, total number of plans and mean planning time (min) for different sorting modes. Target crop is corn silage. Numbers in boldface represent the shortest mean planning time for each target acreage. Sorting modes: (Ctrl) control, (R) field domain, (C) target crop domain, (RC) both domains.

Target acreage (Percent of total)	*Total plans*	*Mean planning time (min)*			
		Ctrl	*R*	*C*	*RC*
Very low (43.1)	6476	**14.6**	17.9	15.7	21.8
Medium low (50.9)	5430	38.2	20.8	30.2	**18.5**
Medium (59.1)	8877	**24.9**	38.0	43.2	30.7
High-medium (66.8)	8480	16.1	**8.5**	22.3	8.6
Medium-high (74.7)	26781	111.8	112.7	**62.6**	147.6
Very high (82.4)	7236	26.1	27.1	**21.2**	46.2

The data were analyzed in the following ways. Histograms of number of plans found versus planning time were used to examine the efficiency of search. To determine the effect of sorting, the mean planning time was computed across all the configurations for plans with profit higher than or equal to 95 percent of the maximum profit. Maximum profit is the profit of the plan with the highest profit across all runs for a given target acreage. The 95 percent value was chosen in consideration that users would be more interested in getting plans in the high tail of the distribution. The distribution of profit was studied by fitting the Normal and Weibull distributions to the data. Parameters were estimated by maximum likelihood for both models (Law and Kelton, 1991). A χ^2 goodness of fit test was performed on the grouped data (Steel and Torrie, 1980).

Planning time

The total number of plans found under the different target acreage and sorting modes is given in table 4-1. The number of plans ranged from 5430 (50.9 target percent) to 26781 (74.7 target percent). For a single crop, all this plans satisfy the environmental constraints. The mean planning time ranged from 8.5 minutes (66.8 target percent and R mode) to 147.6 minutes (74.7

target percent and RC mode in Table 4-1). The results indicate that, in general, sorting the domains reduces the time spent in search, as compared to the control. That is, in four cases, the mean planning time was smaller for the sorting methods than the control and only in two cases (43.1 target percent and 59.1 target percent) the control had a smaller average planning time, probably because high-profit rotations or field combinations were not satisfactory in the last level of search and failed to produce plans.

Plans were found by the planner at relatively constant rates with some exceptional periods during which plans were discovered at a higher rate. In all cases, there were no significant gaps in the discovery of plans across the whole planning time (Figure 4-1). That is, plans were found by the planner during all the planning time.

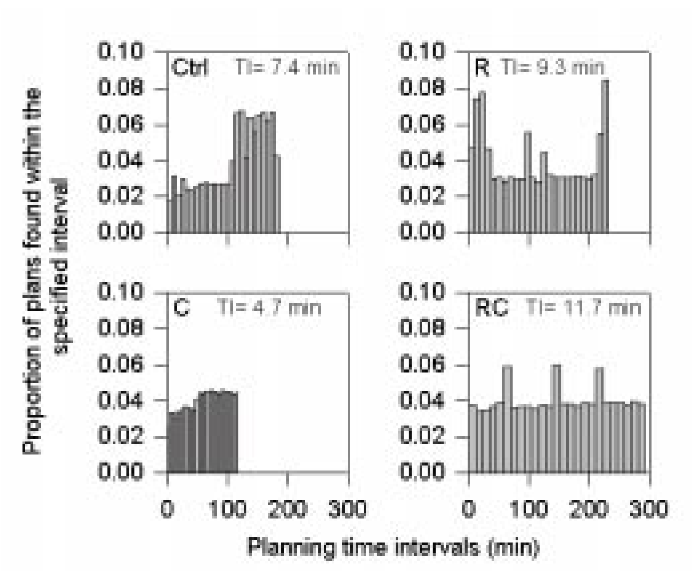

Figure 4-1 Distribution of planning times for different domain ordering modes. (Ctrl) Control. (R) Field domain. (C) Target crop domain. (RC) Both domains. TI= time interval. Annual target acreage: 74.7 percent.

When the planning time to obtain high profit plans was examined, except for the target of 59.1 percent, all the sorting modes yielded planning times lower than the control (Table 4-2). In general, sorting the domains reduced the planning time to obtain high profit plans, but no sorting mode was best over all the target acreage ranges. The results suggest that target-crop domain ordering works the best (best times for three of six target acreages); however, these results are for one farm, without all fields included.

Table 4-2 Annual target acreage as percent of farm area and mean planning time (min) for plans with a profit higher than or equal to 95 percent of maximum profit. Numbers in boldface represent the shortest mean planning time for each target acreage. Sorting modes: (Ctrl) control, (R) field domain, (C) target crop domain, (RC) both domains.

Target acreage (Percent of total)	*Ctrl*	*R*	*C*	*RC*
Very low (43.1)	7.2	4.7	2.5	**1.1**
Medium low (50.9)	1.6	1.5	**0.1**	0.4
Medium (59.1)	**6.2**	10.8	12.8	9.6
High-medium (66.8)	8.6	**4.1**	13.1	4.9
Medium-high (74.7)	100.7	67.7	**58.9**	99.7
Very high (82.4)	19.2	18.3	**1.2**	5.6

To further examine the search efficiency of these sorting strategies, it is useful to see when the best plans were found. Figure 4-2 shows the planning times for the best (highest profit) 100 plans for the target of 74.7 percent. This number of plans corresponds to plans with a profit higher than or equal to 93 percent of the maximum profit. In all the configurations, the plans were found in clusters. For this target acreage range, the best strategy was to sort the target-crop domains; the best plans were found with a mean time of 59.3 minutes while the second best strategy was to sort the field domains (67.5 minutes). Sorting both domains actually increased the planning time (101.7 minutes) suggesting some sort of interaction. The control gave a mean planning time of 99.7 minutes. The effect of sorting was to reduce the time between the discovery of clusters. Plans were also found in clusters in the other target acreages (not shown to save space). This behavior was also noted and exploited in the original CROPS planner (Stone, N. 1998. Personal communication. Virginia Polytechnic Institute and State University. Blacksburg, Virginia). When CROPS finds an acceptable plan, it backtracks to explore new clusters. This behavior is expected to allow the planner to speed up the search and maximize the chance of obtaining plans with a high profit. Whether it reaches the plan with the highest profit within each cluster is not known. These results are important because if, in general, high profit plans were discovered in clusters, then it would indicate that a single cluster contains plans well within 90 percent of maximum profit.

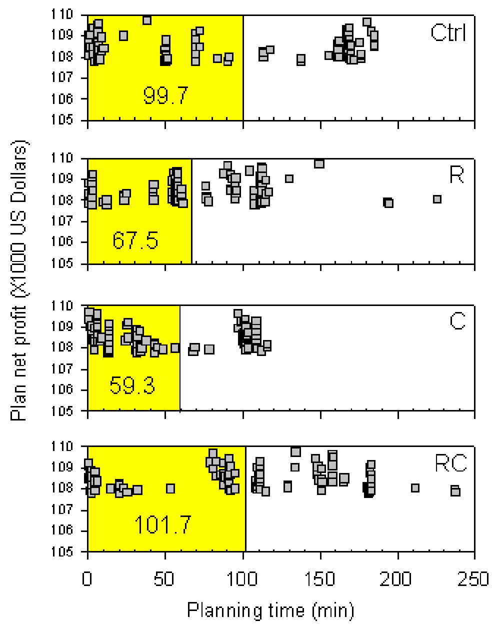

Figure 4-2 Planning time for the best (highest profit) 100 plans and different domain sorting modes. (Ctrl) Control. (R) Field domain. (C) Target crop domain. (RC) Both domains. Numbers within the colored region indicate mean planning time. Annual target acreage is 74.7 percent.

Ultimately, the goal of the planner is to find near-optimal plans in a reasonable amount of time. As shown in Figure 4-3, planning time required increases with the profitability of plans and the number of plans to be found. The best planning strategy was to sort the crop domains (Table 4-2), and in all cases, sorting domains improved the planner's ability to find near-optimal plans quickly.

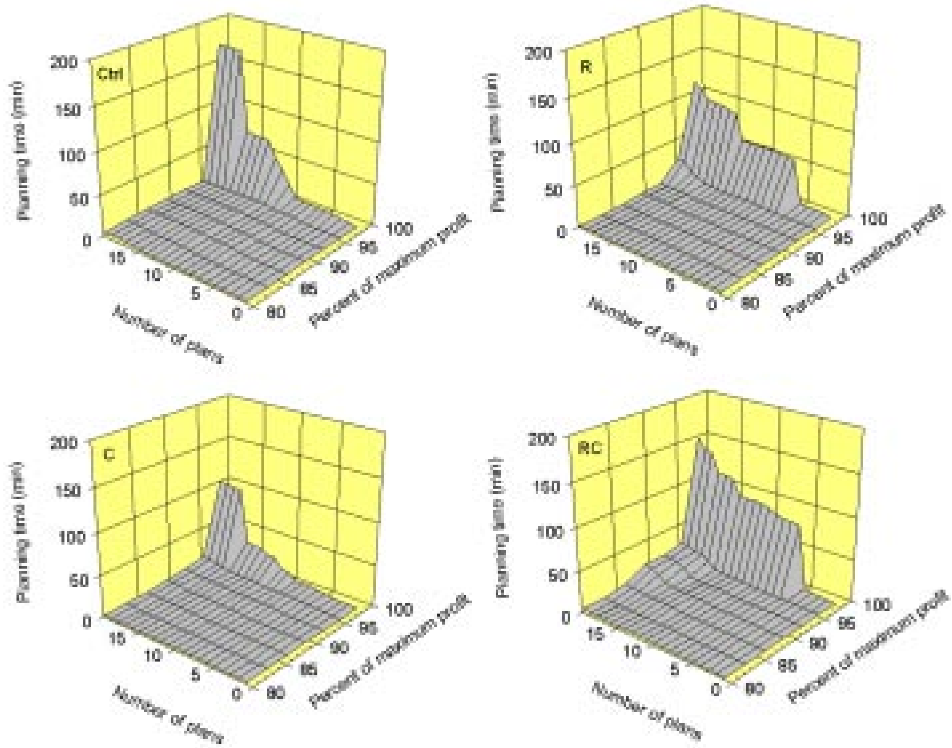

Figure 4-3 Relationships between planning time, number of plans, and percent of maximum profit for different sorting modes. (Ctrl) Control. (R) Field domain. (C) Target crop domain. (RC) Both domains. Annual target acreage is 74.7 percent of farmland.

Distribution of Profit

For a given target acreage, the distribution of plan profit for all plans found was bell-shaped but did not fit either the Normal or Weibull distribution (Tables 4-3, 4-4). Figure 4-4 shows the relative frequency distribution scaled by the width of the intervals (Law and Kelton, 1991) to adjust for the discrete nature of the data. Though neither the Normal nor the Weibull models fit the data statistically, it is apparent that profit does distribute around a mean value. In addition, values around the mean are more likely to occur than extreme values either higher or lower than the mean. The actual distribution of profit shows that high profit plans occur at low probabilities. A good planner should be able to find these plans in a reasonable time. From the previous section, it has been shown that profit distributes across the planning time depending on the sorting criterion.

Table 4-3 Model parameters for the distribution of plan net profit across different annual target acreages. Target crop is corn silage. Parameters estimated by maximum likelihood.

Target acreage (percent)	*Normal*		*Weibull*		*Number of plans (n)*
	Mean	*Standard error*	*Scale*	*Shape*	
43.1	60936.0	63.62	63315.67	9.7146	6476
50.9	66817.0	41.73	68132.48	23.3286	5430
59.1	80011.2	42.71	81969.82	17.6351	8877
66.8	83177.7	44.77	85263.01	16.3965	8480
74.7	93387.31	37.91	96291.31	16.4220	26781
82.4	100077.6	25.84	101168.24	36.1986	7236

Table 4-4 Chi square goodness of fit test for the distribution of plan net profit. *P*-values show nonsignificant fit of models. Target crop is corn silage.

Target acreage (%)	*Normal*			*Weibull*		
	χ^2	*P-value*	*Df*	χ^2	*P-value*	*Df*
43.1	1709.9	0.00000	15	2980.5	0.00000	13
50.9	4627.5	0.00000	15	2320.8	0.00000	16
59.1	864.1	0.00000	15	2398.4	0.00000	15
66.8	4084.1	0.00000	15	7170.7	0.00000	15
74.7	2133.2	0.00000	22	3675.8	0.00000	22
82.4	1534.0	0.00000	15	3186.8	0.00000	16

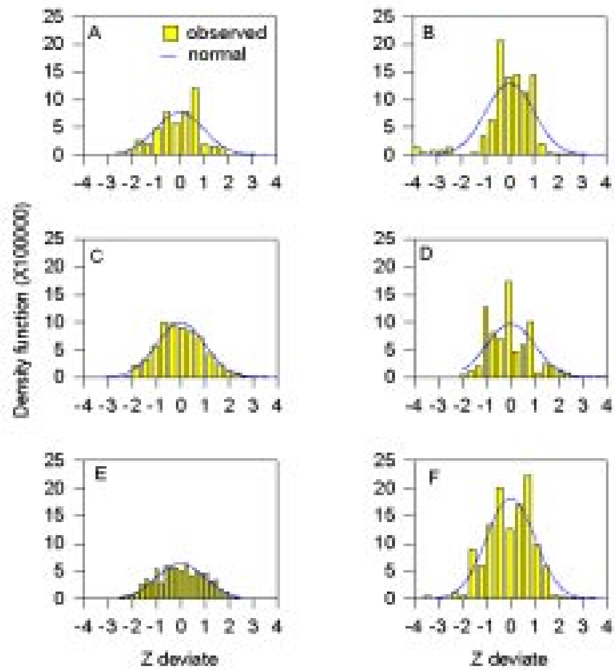

Figure 4-4 Plan net profit distribution for corn silage. Percent of farm acreage as target: (A) 43.1 percent. (B) 50.9 percent. (C) 59.1 percent. (D) 66.8 percent. (E) 74.7 percent. (F) 82.4 percent.

<u>Time-Limited Planning</u>

The objective of this section was to determine the best sorting strategy to find high profit plans for a limited planning time. Planning with limited time may be the more likely scenario for a real world use of a planning system. Alfalfa, corn grain, orchard grass-red clover, soybeans, and corn silage were used as target crops. Annual target acreages were 3.6 and 49.7 percent of the total farm area; the range acreage was 4.05 ha. Based on the results of the previous study, a time limit of 10 minutes was used. Domain sorting modes were the same as in the previous section. The plan net profit was calculated as the mean of the 20 best plans found. All fields for farm-1 were used in these experiments; therefore, the total area was 61.5 ha.

The planner did not find plans for soybeans (49.7 target percent) with the original planning time so the planning time was increased to 20 minutes. In general, sorting the domains improved the profit (Figure 4-5). The best sorting strategies were to sort both domains (six cases of 10) and to sort the field domains (four out of 10). Sorting the domains of target crops had an intermediate benefit and, in five cases, had a negative effect; that is, profit was lower than the control. These results differ somewhat from earlier results looking at whole planning time. In those runs, crop domain sorting worked better. However, those experiments were conducted using only corn silage as a target crop. In addition, the previous experiments showed that plans were found at different average times depending on the percent of maximum profit selected. The important similarity is that both sets of experiments suggest that domain sorting effectively increases the likelihood of obtaining plans with higher profit sooner.

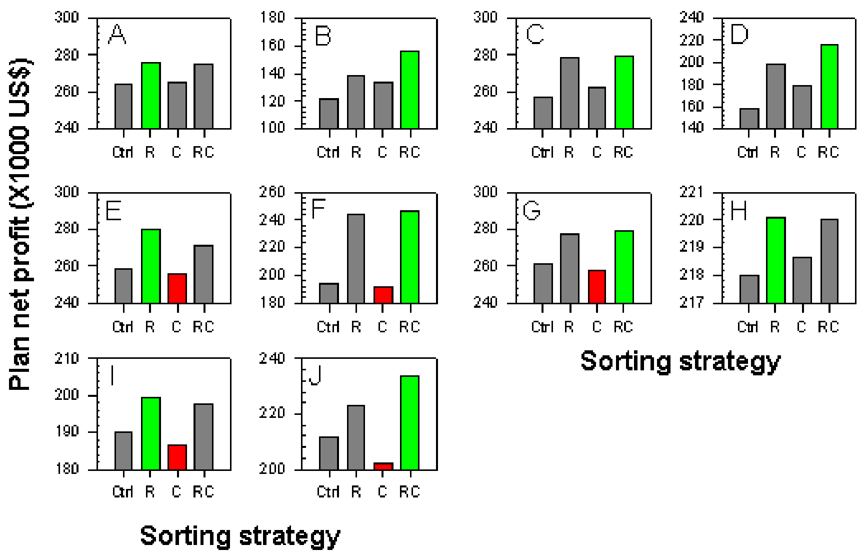

Figure 4-5 Plan net profit for different target crops and annual target acreages. Sorting strategies: (Ctrl) Control, (R) field domain, (C) target crop domain, (RC) both domains. Green bar indicates best (highest profit) strategy. Red bar indicates profit lower than the control. (A) Alfafa 3.6 target percent. (B) Alfalfa 49.6 target percent. (C) Corn grain 3.6 target percent. (D) Corn grain 49.7 target percent. (E) Orchard grass-red clover 3.6 target percent. (F) Orchard grass-red clover 49.7 target percent. (G) Soybeans 3.6 target percent. (H) Soybeans 49.7 target percent. (I) Corn silage 3.6 target percent. (J) Corn silage 49.7 target percent.

To determine the possible factors affecting the efficiency of sorting, the first plan found by each sorting strategy was selected and the distribution of the target crop across the fields and the planning years was examined. Two factors were found to affect the plan profit. The first factor is the high occurrence of non-target crops at low target acreage (3.6 target percent). The second factor is the pattern of occurrence of the target crop over the planning years in a field, named within-field distribution; this factor appeared at medium target acreage (49.7 target percent). Figure 4-6 shows where and when alfalfa was assigned in one plan in which the alfalfa target was 3.6 percent of the farm acreage. At this low target acreage, the difference in profit across the different sorting modes depends mostly on the effect of sorting on the non-target crop(s) that compose the plan. In these plans, corn silage is the most common non-target crop (Figure 4-6 red squares). The non-target crop(s) are the most profitable because when selecting rotations to fill in a plan, the planner chooses the rotations with the most profitable crops first (when sorted). So, the more frequent this non-target crop(s) within a rotation, the higher the profit.

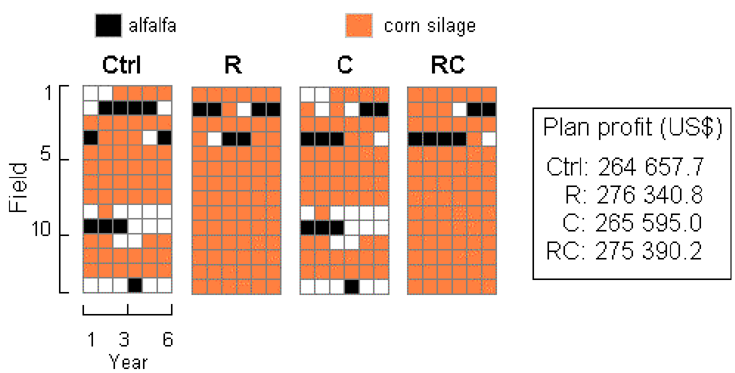

Figure 4-6 Distribution across fields and years of alfalfa (target crop) and corn silage (non-target crop) for different sorting strategies. (Ctrl) Control. (R) Field domain. (C) Target crop domain. (RC) Both domains. Annual target acreage is 3.6 percent of total farm area.

The second factor affecting the efficiency of sorting is the within-field distribution of the target crop. Figure 4-7 shows such distribution for alfalfa from plans generated using annual target acreage of 49.7 percent. Apparently, the within-field distribution of the target crop affects the frequency of the non-target crops, which determine the complementary profit.

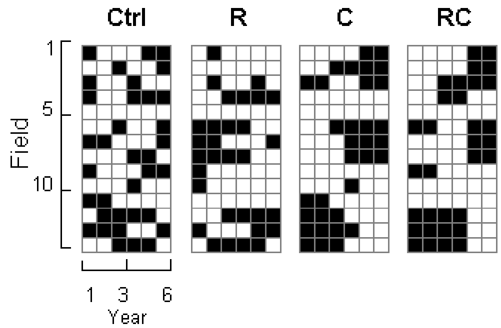

Figure 4-7 Distribution of alfalfa (49.7 target percent) for different sorting strategies. (Ctrl) Control. (R) Field domain. (C) Target crop domain. (RC) Both domains

The assumption that the within-field distribution affects the plan profit was tested by applying a test of ***runs***. A ***run*** is a set of elements that have the same feature in an ordered sequence. The ***runs*** test is a statistic proposed to measure aggregation of this kind (Wald and Wolfowitz, 1940; Madden et al. 1982; Sokal and Rohlf, 1995). The feature in this case is the occurrence or absence of the target crop in a given field and a given year. For example, field one in figure 4-7 (Ctrl) contains three ***runs***, the first has one element, corresponding to the target crop in the first year. The second ***run*** has three elements, corresponding to the next three years with non-target crops, and the third ***run*** has two elements, corresponding to the last two consecutive years with the target crop. Therefore, the test of ***runs*** measures the pattern of occurrence of the target crop in the planning years in a single field. A random pattern indicates that the target crop occurs at random years. An aggregated pattern (few ***runs***) indicates that the target crop occurs in consecutive years, while a uniform pattern (many ***runs***) indicates that the target crop occurs in alternate years. ***Runs*** were computed as described by Madden et al. (1982) and added across all the fields to obtain a total estimate. Exact *P*-values were computed using the distribution of ***runs*** (Wald and Wolfowitz, 1940) with a total sample size of 84, which is the total number of field-years (14 × 6). The test is a two-tail test. The null hypothesis is a random pattern and the alternative hypotheses are aggregation and uniform distributions. Table 4-5 presents the statistics of the test. The sorting strategies aggregated the target crop within fields. In other words, the target crop is planted in more consecutive years, and this aggregation is statistically significant.

The relationship between plan net profit and the number of ***runs*** for alfalfa (49.7 target percent) showed that the unsorted domains (control) had the largest number of ***runs*** and lowest profit, while sorting both domains (RC) had the lowest number of ***runs*** with the highest profit (Figure 4-8). The individual sorting strategies (R and C) have intermediate values. Therefore, the result supports the assumption that the more the target crop is planted continuously, the higher the profit.

Table 4-5 Target crop, sorting strategies, total number of years the target crop is planted (N1), *runs* number and *P*-value for the test of *runs* for alfalfa and corn silage plans.

Target crop	*Sorting strategy*	*N1*	***Runs***	*P-value (alternative hypothesis)*
Alfalfa 49.7 target percent	Control (Ctrl)	35	42	0.4417 NS[1] (uniform)
	Field domain (R)	31	32	0.0363* (aggregated)
	Target crop domain (C)	33	27	0.0008** (aggregated)
	Both domains (RC)	32	27	0.0011** (aggregated)
Corn silage 49.7 target percent	Control (Ctrl)	35	40	0.3805 NS (aggregated)
	Field domain (R)	49	29	0.0026** (aggregated)
	Target crop domain (C)	31	44	0.2111 NS (uniform)
	Both domains (RC)	44	17	0.0000** (aggregated)

[1]NS: nonsignificant, * significant at 0.05 type-I error level, ** significant at 0.01 type-I error level

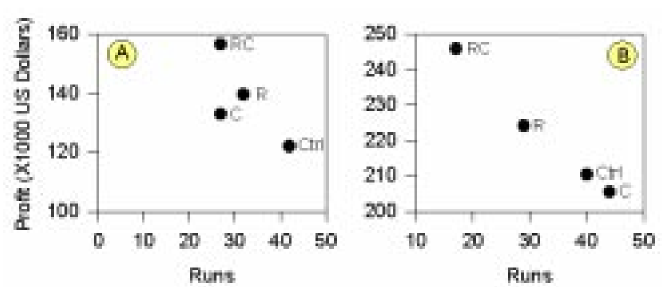

Figure 4-8 Relationship between plan net profit, *runs*, and sorting strategies. (A) Alfalfa (3.6 target percent). (B) Corn silage (49.7 target percent). (Ctrl) Control. (R) Field domain. (C) Target crop domain. (RC) Both domains.

Data for corn silage (49.7 target percent) was analyzed the same way to test the assumption with another crop. Figure 4-9 shows the distribution of the crop across the fields and planning years. In this case, sorting the crop domains reduced the profit (recall Figure 4-5J). The relationship between ***runs*** and profit suggests again that planting the target crop in consecutive years increased the profit (Figure 4-8B). The field combination sorting strategy had the lowest profit with the highest number of ***runs***, while sorting both domains had the highest profit with the fewer number of ***runs***.

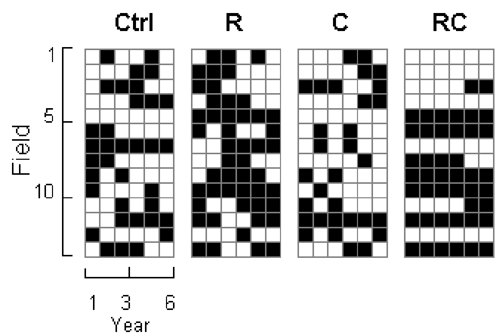

Figure 4-9 Distribution of corn silage (49.7 percent of farm area as target) for different sorting strategies. (Ctrl) Control. (R) Field domain. (C) Target crop domain. (RC) Both domains.

The previous results suggest that the efficiency of domain sorting can be affected by the occurrence of non-target crops and by the within-field distribution of the target crop. For low target acreages, non-target crops strongly influence plan profit while at higher target acreages, the distribution of the target crop within the fields plays a greater role. When the target crop is dispersed within the fields, then the planner is less likely to be able to insert highly profitable non-target crops in the plan. When the target crop is clustered in time and space, then the planner is better able to insert non-target crops with high profits.

Value ordering has been proposed to improve the performance of backtracking (Haralick and Elliot, 1980; Pearl and Korf, 1987; Kumar, 1992; Sadeh and Fox, 1995). In this research, the main purpose of value ordering was to obtain better plans in terms of profit. By putting the values with the highest profit at the beginning of the search, the likelihood of obtaining better plans in a limited time across different target crops was increased. The best sorting strategies of those tested were sorting the rotations and sorting both rotations and field combinations.

Effect of Target and Range Acreage

One of the objectives of the planner is to find plans that satisfy annual acreage requirements. The annual target acreage is established by the user as a range acreage for a given (target) crop. That is, target acreage varies between a minimum and a maximum value. Therefore, it is important to know how target and range acreage affects the planning process. Two objectives were pursued in this section. The first objective was to determine how the target acreage and acceptable range affects the time needed to obtain the first acceptable plan. The second objective was to determine the feasibility of planning for different target crops across different annual target acreages. Corn silage, alfalfa, soybeans, and millet were selected as target crops. The target acreage ranged from 3 percent to 97 percent of the total farm area, each target acreage was allowed to vary within a range of 2.02 and 4.05 ha. The response variables were the time to obtain the first acceptable plan and the number of consistency checks during the tree search. A consistency check is the comparison of a value against the constraint (the fewer the number of consistency checks the more efficient the search (Nadel, 1989)). Only consistency checks for binary constraints were counted.

The results are presented in Table 4-6 and 4-7 for the different target ranges. The only crop for which plans were found for all the target acreages was corn silage. For the other crops, the planner found plans only when target acreages were below 60 percent of the total acreage. This reflects the fact that corn silage is present in several rotations; these rotations contain corn silage from one year in six to six in six (Table A-3, Appendix A).

Table 4-6 Time to reach first plan (minutes) and consistency checks for different target crops. Range is 2.02 ha. Annual target acreage as percent of farm area. (Ed) empty domain.

Target acreage (%)	*Corn silage*	*Alfalfa*	*Soybeans*	*Millet*
4.9	0.0165 (7987) [1]	0.0164 (7114)	0.0228 (7046)	0.029 (5805)
18.1	0.0686 (58433)	0.119 (45690)	0.163 (81336)	0.131 (72782)
31.2	0.227 (259137)	0.399 (147752)	0.4173 (250092)	0.48 (239057)
44.4	0.45 (228167)	0.555 (228511)	0.6608 (390666)	Ed
57.6	0.436 (211473)	0.555 (263072)	Ed	Ed
70.8	0.228 (126280)	Ed	Ed	Ed
83.9	0.131 (87867)	Ed	Ed	Ed
97.1	0.0265 (7009)	Ed	Ed	Ed

[1]Consistency checks in parentheses

Table 4-7 Time to reach first plan (minutes) and consistency checks for different target crops. Range is 4.05 ha. Annual target acreage as percent of farm area. (Ed) empty domain.

Target acreage (%)	*Corn silage*	*Alfalfa*	*Soybeans*	*Millet*
3.6	0.0175 (6277) [1]	0.025 (7195)	0.029 (7304)	0.027 (8287)
16.4	0.110 (69337)	0.264 (67283)	0.179 (123150)	0.233 (103966)
29.6	0.531 (257415)	0.938 (267456)	0.83 (396008)	0.847 (507533)
42.8	1.1755 (427075)	1.433 (440561)	1.764 (847248)	Ed
55.9	1.321 (419554)	2.161 (508358)	Ed	Ed
69.1	0.7452 (300073)	Ed	Ed	Ed
82.3	0.326 (200007)	Ed	Ed	Ed
95.4	0.0613 (34616)	Ed	Ed	Ed

[1]Consistency checks in parentheses

The area that can be planted any given year with a target crop in a given rotation is the ratio between the number of years the target crop exists in that rotation and the rotation cycle. For example, if a target crop is present in a single year and the rotation cycle is six years, the area planted with this target crop using this rotation is just 1/6 of the total. Therefore, in general, the maximum potential area that can be planted in a given field with a target crop is restricted by the rotation with the highest number of years the target crop is present in that rotation. This maximum potential is further reduced by the field constraints that may eliminate one or more of the rotations that contain the target crop. Corn silage and soybeans are present in some rotations in all six years of the rotation cycle (Table A-3 Appendix A). However, soybean is present only in one rotation during the six years of planning. Should this rotation be eliminated in one or more fields, the potential target area would be reduced (Tables 4-6 and 4-7). In the case of alfalfa, it is contained in at most four years while millet is present in at most two years (Table A-3 Appendix A), thus restricting the potential target acreage.

There is a consistent pattern in the relationship between target acreage, acceptable acreage range, and time to reach the first plan (Figure 4-10). The pattern is determined by three relationships. As target acreage moves from extremes (zero and 100 percent) to 50 percent, the number of possible field combinations increases. The second relationship is that for any given target acreage, as the acceptable range increases, the number of field combinations increases. Finally, as the number of field combination increases, the planning time increases. The increase in the planning time is because the increase in the domain size (field combinations) increases the number of consistency checks (Figure 4-11).

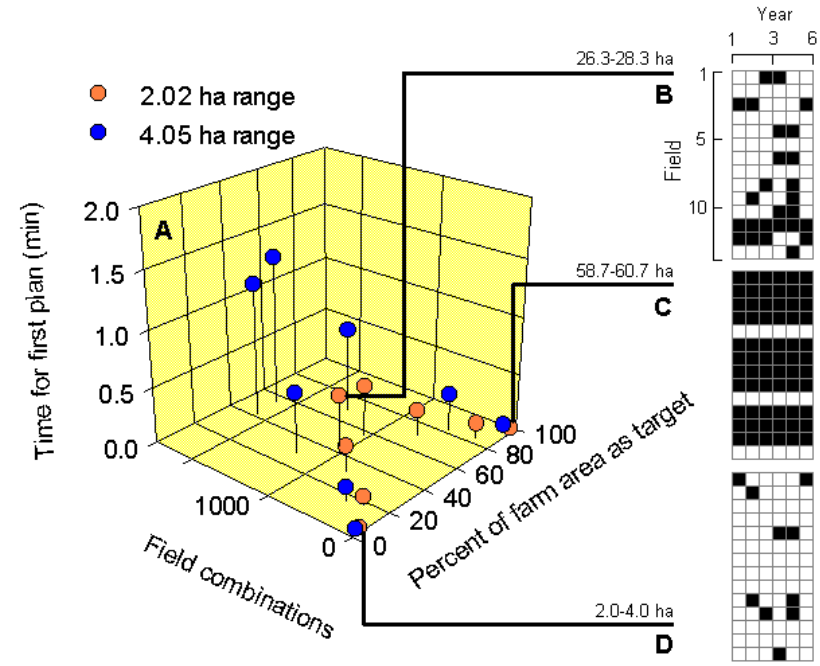

Figure 4-10 Relationship between annual target acreage, field combinations, and time to obtain the first plan (A). Distribution of target crop for selected plans (B-D). Target crop: corn silage.

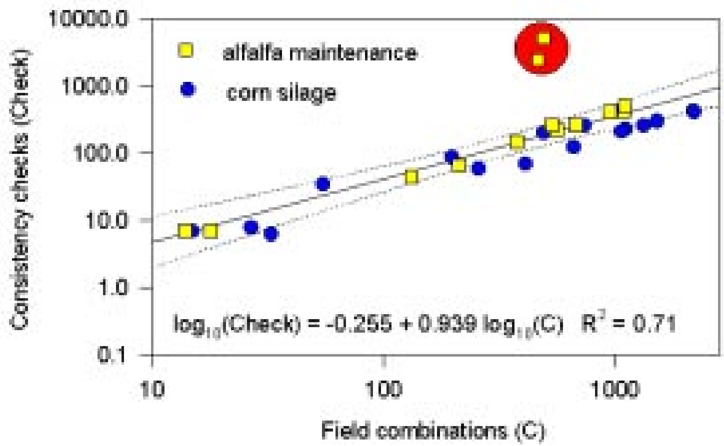

Figure 4-11 Regression between consistency checks and domain size (field combinations). Dotted lines indicate a 95 percent confidence interval. Outliers (encircled points) are explained in the text.

In the case of alfalfa, plans were obtained for up to 56 percent of the farm area. Additional runs were performed for target acreages between the last successful runs and the failed runs of tables 4-6 and 4-7 so plans were found near to 65 percent (Figure 4-12). The relationship between target acreage, acceptable acreage range, and time to reach the first plan for alfalfa is similar to the relationship for corn silage. However, there are two main differences. The first difference is that the planner did not find plans for all the target acreages. This problem was explained already. The second difference is that the time to reach the first plan increased again at target acreages close to 60 percent (Figure 4-12A). For example, for the target of 37.4-39.4 ha (62 percent), the planner required 43.5 minutes to reach the first plan compared with only 0.55 minutes for 26.3-28.3 ha (44 percent). The plan for the 62 percent target acreage required 9 fields to have rotations with alfalfa planted during 4 years (Figure 4-12B), while for the 44 percent target only 2 fields had rotations with 4 years (Figure 4-12C). The time increased because the planner required performing more checks to find rotations that provided the highest number of years with the target crop. The outliers in Figure 4-11 correspond to the points with the highest planning time in Figure 4-12A. These outliers show that the number of consistency checks and planning time depend not only on the number of field combinations, but also on the crop composition of the rotations. Millet had a behavior similar to alfalfa and soybeans (Tables 4-6 and 4-7).

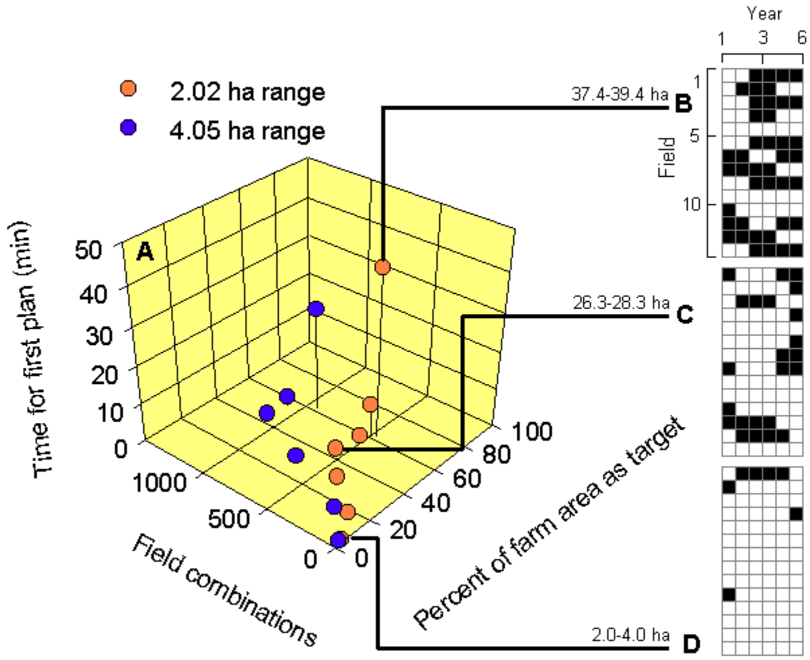

Figure 4-12 Relationship between annual target acreage, field combinations and time to obtain first plan (A). Distribution of target crop for selected plans (B-D). Target crop is alfalfa.

Finally, it is apparent by looking at Tables 4-6 and 4-7 that for a given target acreage, plans are found in the following ascending time order: corn, alfalfa, soybeans, and millet. The most probable explanation is that the frequency of the crops in the rotations affected the speed of search. Table A-3 shows that the occurrence of the crops in the rotations has the following ascending frequency order: soybeans, millet, alfalfa, and corn silage. However, millet is present in the rotations in at most 2 years while soybeans are present across the 6 years. The conclusion is that the more frequent a crop occurs in the rotations, the more likely it can reduce the planning time by reducing the number of checks.

The implications of the previous results have practical applications. For example, to reduce the planning time, one could decrease the domain size of the target crop(s). Such reduction can be obtained by implementing new constraints or, simply by eliminating the domain values at random. Figure 4-13 shows the original runs for corn silage and alfalfa (2.02 ha range) paired with new runs in which the field combination domain was reduced to 30. Times to reach the first plan were lower for both corn silage (Figure 4-13A) and alfalfa (Figure 4-13B). The reduction is greater where the original domain size was higher. However, over-constraining the problem can result in planning failure in which no solution is possible (Tsang, 1993) or can also eliminate good plans. Another practical application would be to increase the number of rotations with a high occurrence of the target crop when planning for high target acreage.

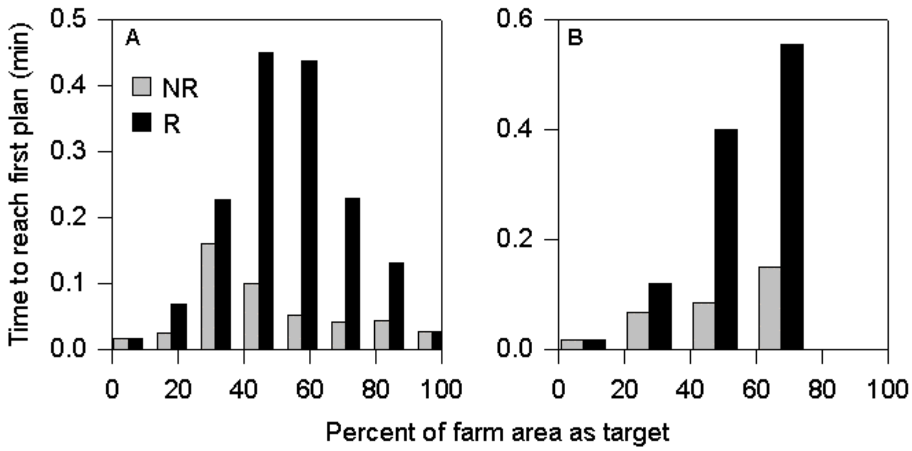

Figure 4-13 Relationship between annual target acreage and time to reach the first plan. (A) Corn silage. (B) Alfalfa. Target range is 2.02 ha. (NR) normal domain size reduction by consistency, (R) domain size reduced to 30 elements after node consistency.

Efficiency of Field Exclusion

Field exclusion is the removal of fields from planning when these fields are not present in any of the field combinations across all the target crops. Field exclusion comes after performing node and arc consistency and before the tree search. These fields can be removed because any rotation without the target crop(s) is equally satisfactory for the given field. Therefore, excluding these fields would speed up the search by reducing the number of levels in the tree search. In addition, the assignment of rotations to these fields can be done using some preference criterion, in this case the most profitable rotation. The current implementation of the planner uses field exclusion enabled by default. Because there is no way to know which fields are going to be excluded until after arc consistency, the examples provided here were selected from the domain ordering section with limited time. Target crops were checked for excluded fields. If fields were excluded, then the same configuration was used but with all the fields. Planning time was set to 10 minutes and the range target was set to 4.05 ha. Response variables were the time to reach the first plan and plan net profit. Comparison between the control (Ctrl, non-excluded fields) and the field exclusion mode (Exc) were performed using a t-test for paired differences using the sorting modes as blocks and a type-I error of 0.05 (Steel and Torrie, 1980).

The planner excluded fields only for crops with annual target acreage of 3.6 percent of the farm area. The excluded fields were 1, 2, 4, 5, 9, 10, 11, and 14. For the annual target acreage of 49.7 percent, the planner excluded a field only for alfalfa. The excluded field in this case (field 12) was excluded by the consistency checks.

Time Performance

The time to reach the first plan was reduced across all the target crops (Figure 4-14). The difference was significant for all the crops except alfalfa (49.7 target percent) (Paired t-test, $P < 0.05$), but the direction of the response is the same. The times to reach the first plan are similar for all the crops except alfalfa (49.7 target percent) for which planning times are higher than the rest. The times are different for this crop because at this target acreage the domain size is larger, the number of consistency checks increases, and the planning time increases, regardless of the field exclusion alternative.

Although the times to reach the first plan were statistically different regarding the field exclusion alternative, the absolute values were small (less than a minute) and thus not practically different. However, by automatically excluding the fields that did not participate in the search, the overall performance of the system was improved. In these examples, the target acreages are very small (less than 5 percent of total acreage) and the planner excluded fields for a larger area and a single crop (alfalfa 47.9 percent of total acreage). Therefore, it is possible that field exclusion will help at larger target acreages. Having larger domain size (number of rotations) may be favored by this approach if considering that the times to reach the first plan strongly depend on the domain size.

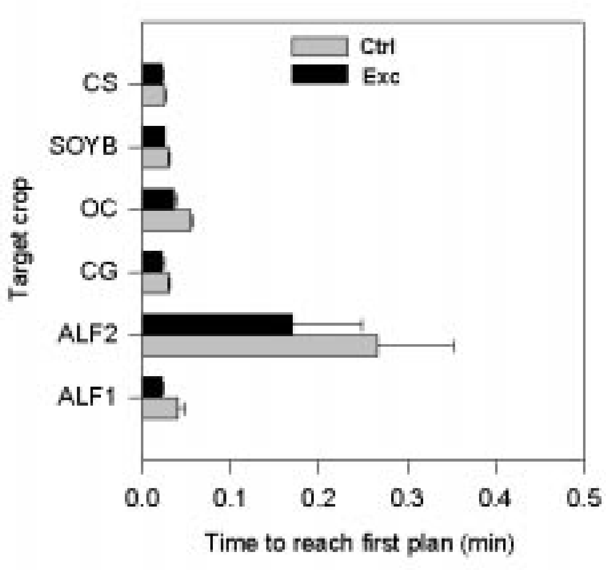

Figure 4-14 Time (mean + SE) to reach first plan by different target crops. CS Corn silage (3.6 target percent). SOYB Soybeans (3.6 target percent). OC Orchard grass-red clover (3.6 target percent). CG Corn grain (3.6 target percent). ALF2 Alfalfa (49.7 target percent). ALF1 Alfalfa (3.6 target percent). Exclusion modes: (Ctrl) Control, (Exc) Field exclusion.

Economic Performance

All the crops with the target of 3.6 percent have a higher profit when fields are excluded but the difference is significant only for alfalfa (Figure 4-15). For alfalfa (49.7 target percent), the profit is slightly higher for the control but the difference is not significant. The eight fields removed for the target of 3.6 percent account for 23.5 percent of the total farm area, while the single field excluded for alfalfa (49.7 target percent), account for 21.1 percent, similar to the previous reduction. Therefore, for the target crops with target percent of 3.6, the planner searched in only six fields (14 for the control), thus giving more opportunity to select high profit plans. In the case of the target of 49.7 percent, only one field was removed, and the search was done on 13 fields (14 for the control), therefore, the plan profit was similar (non-significant) to the control.

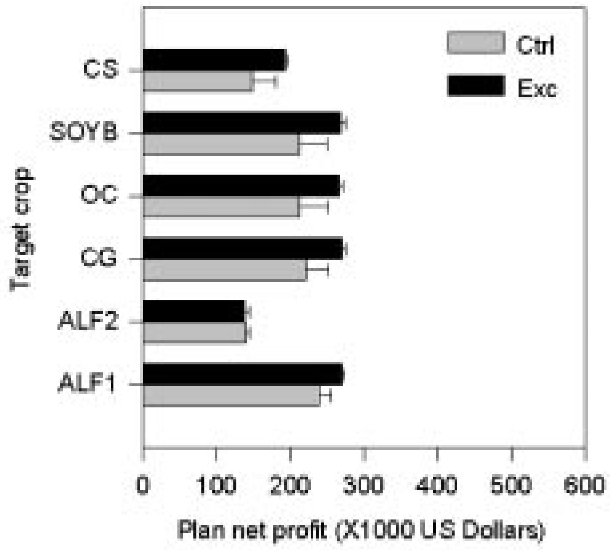

Figure 4-15 Plan net profit (mean + SE) by different target crops. CS Corn silage (3.6 target percent). SOYB Soybeans (3.6 target percent). OC Orchard grass-red clover (3.6 target percent). CG Corn grain (3.6 target percent). ALF2 Alfalfa (49.7 target percent). ALF1 Alfalfa (3.6 target percent). Exclusion modes: (Ctrl) Control, (Exc) Field exclusion.

Summary and Conclusions

The efficiency of the planning component was analyzed regarding domain ordering heuristic, target and range area, and field exclusion. Target and range acreage were evaluated in terms of time to obtain the first plan and number of consistency checks made by the planner. Field exclusion was tested across different target crops and annual target acreages.

The following conclusions were reached:

- Sorting the domains by profit allowed the planner to find high profit plans in less time than using unsorted (randomized) domains.
- High profit plans are found in clusters and the sorting strategies reduce the time to find such clusters.
- There was no absolute best sorting mode, but sorting field domains and both domains usually gave more profitable plans than the control when planning for limited time.
- The time to reach the first plan depended on domain size and crop composition of rotations. Domain size was affected by the target acreage. The higher the target, the higher the domain

size up to areas in the middle of the total farm area. At larger target acreages, the domain size reduces and so do the planning times.

- The occurrence of the target crop within the rotations affected the target acreage. The maximum target acreage was limited to the rotation with the highest number of years with the target crop.
- Range acreage also affected the time to reach the first plan. As the target range increased, the domain size also increased, thus increasing the time to reach the first plan.
- Field exclusion helped to reduce the time to reach the first plan by reducing the search space. In addition, field exclusion usually increases the profit by automatically assigning the highest profitable rotation to the excluded fields.
- In general, the domain sorting and field exclusion heuristics were appropriate in terms of reducing planning time and obtaining high profit plans and it is suggested to used them as default values in the planner.

Chapter 5 Implementing Preventive IPM Using Constraint Satisfaction: The Corn Rootworms CRW (*Diabrotica virgifera virgifera* and *D. barberi*) Control Model

Introduction

The economic importance as well as the particular biology of corn rootworms makes them ideal candidates to represent preventive IPM heuristics in a constraint-based control model. The control model would help to select plans with minimum risk of damage caused by these pests. Therefore, the objective of this chapter is to describe, implement, and evaluate a constraint-based model for preventive IPM control of the CRW. The chapter includes a description of the assumptions upon which the preventive model is based. A description of the constraints that compose the model is then presented. The model components were tested under different scenarios to determine their performance. Plans were generated for two representative farms from Virginia's Shenandoah Valley. To be successful, the planner needed to find plans that improved profits by reducing the risk of CRW outbreaks.

Model Assumptions

The CRW preventive control model is based on three assumptions. The first assumption is that the risk of CRW damage is proportional to the number of times 2nd-year corn is planted in a rotation cycle. This assumption is derived from the life-history of the pests. Second-year corn is more susceptible because if corn is planted in the first year, the arriving female adults will lay their eggs and the eggs will produce the damaging larvae for the next year. First-year corn (corn planted after a non-host crop) is less susceptible because the adults are not attracted to non-corn crops that occurred the year before (Levine and Oloumi-Sadeghi, 1991).

The second assumption is that the damage caused by larvae that come from prolonged diapause is negligible. Therefore, although prolonged diapause has been reported for both species of CRW, it is not included in the model. Furthermore, the occurrence of prolonged diapause does not necessarily mean that first-year corn will suffer damage (Steffey et al. 1992). Prolonged diapause occurrence also varies between species and by region (Levine and Oloumi-Sadeghi, 1991; Levine et al. 1992ab), thus a more complex model would be needed to account for these differences.

The third assumption is that fields can be classified into different CRW risk categories based only on biophysical properties, not on cultural practices. Some research supports this assumption. For example, Godfrey et al. (1995) reported that egg survival is affected by depth, soil temperature, and soil moisture, which depends on soil texture. In addition, laboratory studies have shown that survival of CRW larvae increases when the content of clay in soil increases (Turpin and Peters, 1971). It has also been reported that CRW causes reduced damage in sandy soils (Turpin et al. 1972).

To incorporate preventive IPM into the planner, based on these assumptions, I defined categories to identify the risk of rotations to CRW damage and rules to select rotations that were compatible with fields. The decision rules were encoded as field constraints. The strategy was to identify and eliminate rotations that were likely to promote CRW damage due to growing corn continuously or planting corn in high risk fields. As before, crop constraints also encoded the decision rules to ensure that other goals can be met, including annual target acreage. The model was thus composed of field and crop constraints and was integrated with the planning system described in Chapter 3.

The CRW Preventive Control Model

Field Constraint Description

The control of CRW can be approached in the same way as other field-level problems by selecting rotations that are most appropriate to control or mitigate the given problem. Rotations were ranked by risk according to how frequently 2nd-year corn occurred (Figure 5-1). For purposes of this discussion, I define “2nd-year corn” to be any corn planted in a rotation following corn the previous year. This means that third-year or four-year corn would be counted as 2nd-year corn. For example, in a six-year rotation, corn can be planted from zero to six years. The two extreme cases are rotations with no corn and rotations planted with corn in each of the six years. The first case does not involve CRW risk. The last case represents the highest risk because 2nd-year corn occurs six times. If we consider the rotation as a cycle, the first corn crop in the rotation is 2nd-year corn because corn was planted in the sixth year of the previous cycle. Rotations with just one-year corn have no 2nd-year corn; rotations with two corn crops in six years could have either no 2nd-year corn or one 2nd-year corn, etc. (Figure 5-1C).

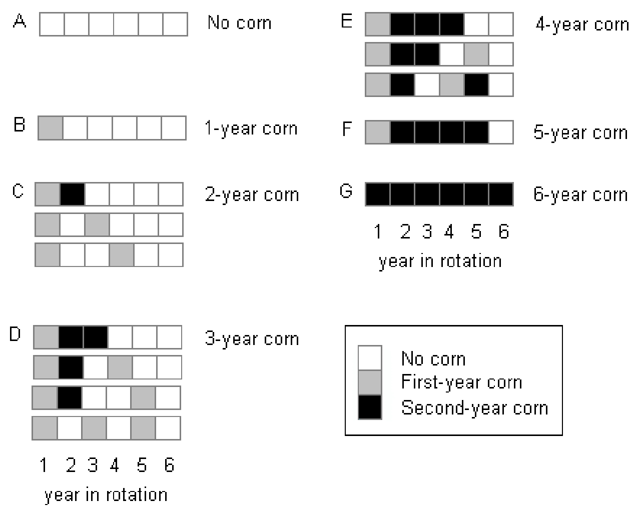

Figure 5-1 Number of 2nd-year corn crops found in six-year rotations involving corn. (A) No corn has zero 2nd-year corn. (B) Corn planted in one year has no 2nd-year corn. (C) Corn planted in two years has zero or one 2nd-year corn. (D) Corn planted in three years presents zero, one and two 2nd-year corn. (E) Corn planted in four years presents two and three 2nd-year corn. (F) Corn planted during five years presents four 2nd-year corn. (G) Corn planted during all six years has six 2nd-year corn.

Based on the number of 2nd-year corn crops it contains, a rotation can be classified into a risk category and an ordinal risk value can be assigned to this category to represent CRW damage risk. Table 5-1 shows an *ad-hoc* classification of all the rotations into three distinct risk categories: low, medium, and high, depending on the number of 2nd-year corn crops contained. As for rotations, the fields can also be classified into three categories: low, medium, and high, with similar ordinal values. Because what soil or field properties can be used to classify fields is not well defined, as compared to rotations, only the categories are proposed here. In addition, to account for this uncertainty, an empirical constant was included to compute a plan CRW risk index (see plan CRW risk description later) to make the analysis. Based on this classification, a decision table was constructed to select or reject rotations for specific fields according to their estimated risk. Table 5-2 shows the rules used here for acceptance and rejection of rotations according to its risk category. Three out of nine possible combinations are unacceptable because these would favor the occurrence of CRW.

Table 5-1 Risk categories and values to classify rotations to CRW damage.

CRW Risk		*Number of 2nd- year Corn crops in the rotation*	*Possible number of corn crops present out of six in a six year rotation*						
Category	*Value*		*0*	*1*	*2*	*3*	*4*	*5*	*6*
Low	1	0-1	X	X	X	X	-	-	-
Medium	2	2-3	-	-	-	X	X	-	-
High	3	4-6	-	-	-	-	-	X	X

Table 5-2 Decision rules to identify acceptable assignments of rotations to fields based on CRW risk categories.

Field CRW Risk category	*Rotation CRW Risk category*		
	Low	*Medium*	*High*
Low risk	Acceptable	Acceptable	Acceptable
Medium risk	Acceptable	Acceptable	**Unacceptable**
High risk	Acceptable	**Unacceptable**	**Unacceptable**

The decision rules in Table 5-2 were implemented in the planner as a unary field constraint. For each field in the farm, the planner checks the field constraint during the node consistency stage (refer to Chapter 3). Rotations that do not satisfy the rules in Table 5-2 are eliminated. If enforcing this constraint eliminates all rotations for a field, the constraint can be relaxed. Whether or not to allow constraint relaxation is an option for the user. In this analysis, relaxation was allowed.

Crop Constraint Description

The field combinations can be evaluated for CRW risk based on the individual risks of the fields that compose the given field combination. A CRW risk value for a field combination is computed as

$$R_c = \sum_{i=1}^{n} (\gamma_i a_i) / \sum_{i=1}^{n} a_i \qquad (5\text{-}1)$$

R_c is the risk for a field combination, γ_i is the CRW risk for the i^{th}-field, a_i is the area of the i^{th}-field, and n is the number of fields. As for rotations, the risk (R_c) can take values from 1 to 3 (low risk to high risk). To reduce CRW damage or control costs, the planner should eliminate the field combinations with the highest risk. In practice, the planner sorts the valid field combinations by CRW risk and retains a given number of the lowest-risk combinations. This constraint cannot be relaxed, but the number of combinations to retain can be adjusted.

Evaluating Model Performance

Experimental Conditions

The model was evaluated using a representative farm of Rockingham County, Virginia. The farm is described in Chapters 4 and 6 and called farm-1. In brief, farm-1 is composed of fourteen fields and has 62 ha of cropland. Field properties are listed in Table A-2 (Appendix A). The objective was to determine how CRW risk and profit respond to different levels of target acreage and high risk fields. The hypothesis is that adding these constraints to the planner results in plans that show lower risk of CRW. Ideally, the planner will be able to find profitable plans that meet target goals and help control CRW. To accomplish this objective, two factors were evaluated: target acreage (τ) as the percent of the total farm area, and the percent of farm area with high risk fields (ρ). Annual target acreage ranged from 10 to 90 percent of the farmland. High risk values were randomly assigned to fields from 10 to 90 percent of the total farm area. Within each high risk level the remaining fields were assigned to medium and low risk with a 0.50 probability. The planner was run for all the level combinations of target acreage and high risk field percentage.

Four runs were used at each combination of levels: control (Ctrl, no constraint enabled), field constraints enabled (R), crop constraint enabled (C), and both constraints enabled (RC, complete model). The planner was set to return 20 plans per run. Domains were sorted by profit after the node consistency stage. Twenty minutes for the tree search stage was allowed for each run. Corn silage was used as the target crop with a 4.05 ha range of acceptable values around the given target. Field constraints were allowed to relax if necessary.

For each of the runs, plans were evaluated to compute three dependent (response) variables: a plan CRW risk (R_p), corn silage net profit and the plan net profit. The plan CRW risk index (R_p), was computed as

$$R_p = \sum_{i=1}^{n} a_i (kr_i + \gamma_i) / \sum_{i=1}^{n} a_i \qquad (5\text{-}2)$$

R_p is the risk index for the pth-plan; a_i is the area of the ith-field, r_i is the rotation risk index, γ_i is the CRW risk for the ith-field, and k is an empirical dimensionless value. It represents the importance of or confidence in the rotation risk as compared to the field risk. For example, a value of k greater than one would indicate that the rotation risk is of more importance than the field risk. In the current implementation, k was set to two *a priori* to reflect more confidence in the estimates of rotation risk. The second dependent variable was net profit for corn silage (π) on a per yield unit (US$/ Mg), determined by dividing net profit in US$/ha by yield (Mg/ ha). This factors out profit differences due to planting corn in fields that are more productive or have different size. The third and last variable was plan net profit (Ψ). The plan net profit was computed as the difference of gross income minus production costs (variable and fixed costs) across all crop enterprises in the plan. Variable costs included an insecticide cost of US$37/ ha (Kuhar et al. 1997). Fields with a 2nd-year corn crop received a chemical application to control CRW with a probability of 0.41 regardless of the treatment (based on estimates from Virginia farms. C.A. Laub. 1999. Personal communication. Virginia Polytechnic Institute and State University. Blacksburg, Virginia).

Additionally, paired differences for some response variables were computed within each combination of levels of τ and ρ between the model components and the control. Differences indicate both magnitude and direction of the change in the response variable, as compared to the control. Negative differences indicate a reduction in the response variable and positive differences indicate an increase. For example, for the plan CRW risk, a negative difference between the field constraint and the control indicates a reduction in CRW damage.

The data were analyzed by fitting a second-degree polynomial $Z = \beta_0 + \beta_1\rho + \beta_2\tau + \beta_3\rho^2 + \beta_4\tau^2 + \beta_5\rho\tau$ to the data to represent a response surface. The variable Z could be a dependent variable or a difference. A second-degree polynomial was chosen because it usually gives enough information when performing empirical analysis of data for two independent variables (Draper and Smith, 1981). In the case of differences, statistical significance of the regression model indicates a significant effect of the CRW model components. CRW risk was analyzed directly (R_p) and as a difference (R_Δ), net profit for corn silage was analyzed as a difference (π_Δ), and plan net profit was analyzed as a difference (Ψ_Δ). A stepwise procedure was applied to select the regression models with the least number of significant parameters. The REG procedure (SAS Institute, 1990) was used to estimate the parameters and to select the models. For the stepwise model selection, a significance level of p=0.15 was used to test the entry and permanence of variables in the model. The regression R^2, ANOVA *P*-value, Mallow's *Cp* and studentized residuals were examined to determine significance of the statistical model and departure from the statistical model assumptions (Box and Draper, 1987; Draper and Smith, 1981; Sokal and Rohlf, 1995). The values of *Cp* are expected to be near the number of the parameters in the model (Draper and Smith, 1981). Studentized residuals should be less than the absolute value of two; an absolute value of two represents approximately a 95 percent level of confidence (Draper and Smith, 1981).

Results and Discussion

CRW Risk Analysis

The results for the plan CRW risk variable are in Table B-1 in Appendix B. The statistical indices R^2, ANOVA *P*-value, and Mallow's *Cp* showed no evidence to reject the regression model (Table 5-3). Few residuals were outside of the critical value of two (Figure 5-2), not undermining the assumptions of homogeneity of variance (Draper and Smith, 1981). The selection process yielded the model parameters in Table 5-4.

Table 5-3 Regression statistics for the model R_p = = $\beta_0 + \beta_1\rho + \beta_2\tau + \beta_3\rho^2 + \beta_4\tau^2 + \beta_5\rho\tau$. R_p is the plan CRW risk, ρ is the percent of farm area with high risk fields and τ is percent of farm area as target.

Statistic	*Control*	*Field Constraint*	*Crop Constraint*	*Both Constraints*
R^2	0.980	0.987	0.990	0.970
ANOVA *P*-value	0.0001	0.0001	0.0001	0.0001
Mallow's *Cp*	3.220	4.116	8.056	4.084

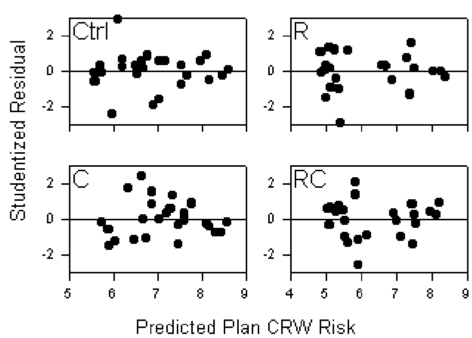

Figure 5-2 Studentized residuals versus expected plan CRW risk (R_p) for the model $R_p = \beta_0 + \beta_1\rho + \beta_2\tau + \beta_3\rho^2 + \beta_4\tau^2 + \beta_5\rho\tau$. R_p is the plan CRW risk, ρ is the percent of farm area with high risk fields and τ is percent of farm area as target. (Ctrl) Control. (R) Field constraint. (C) Crop constraint. (RC) Both constraints.

Table 5-4 Parameters for the model $R_p = \beta_0 + \beta_1\rho + \beta_2\tau + \beta_3\rho^2 + \beta_4\tau^2 + \beta_5\rho\tau$. R_p is the plan CRW risk, ρ is the percent of farm area with high risk fields and τ is percent of farm area as target.

Model Parameter	*Control*	*Field Constraint*	*Crop Constraint*	*Both Constraints*
β_0	5.54868	5.78108	5.47001	5.87078
β_1	0.01340	-0.00618	0.01392	-0.00924
β_2	-0.01277	-0.05257	0.01384	-0.02215
β_3	0	0	0	0
β_4	0.00037	0.00077	0.00008	0.00043
β_5	0	0.00020	0	0.00028

Response surfaces (Figure 5-3) visually show a general pattern, in which the risk increased as both τ and ρ increased but at a higher rate for τ. For the field constraint, the response indicates a reduction in risk for τ values less than 60 percent and across all the ρ values. At low target acreages, there was an increase of R_p that did not happen in the control; however, the absolute values of Pr were still lower than the control. In the case of the crop constraint, the risk response was similar to the control but with higher values. The full model (both constraints) presented a reduction in risk but less pronounced than the field constraint alone, indicating some interference by the crop constraint.

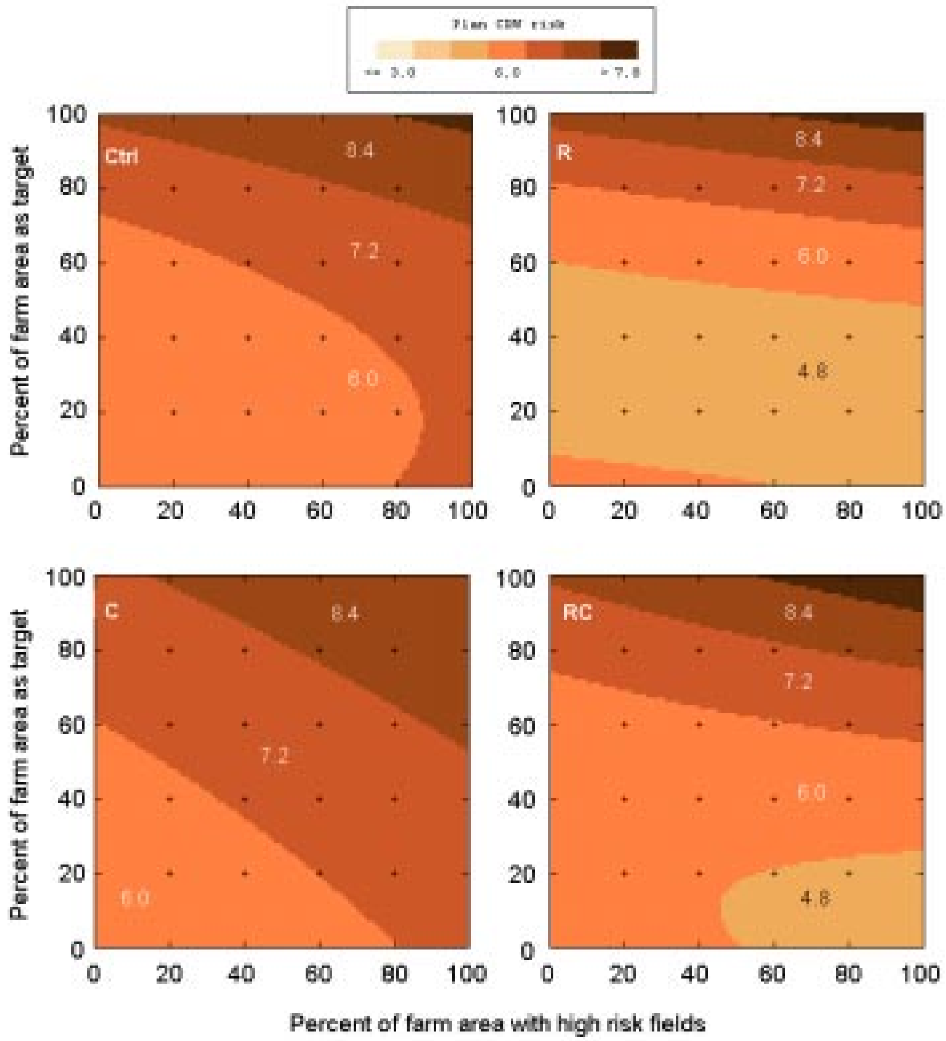

Figure 5-3 Plan CRW risk (R_p) response surface for the control and CRW model components. (Ctrl) Control. (R) Field constraint. (C) Crop constraint. (RC) Both constraints (full CRW model). Numbers in regions correspond to the mean value.

The analysis of the difference in R_p between the control and the constraint-enabled model gives more insight into where the reduction in R_p is happening (if any) and the magnitude of the reduction. Table 5-5 shows the parameters for such differences. There was no effect of the farm acreage with high risk fields on the crop constraint risk difference. Values of R^2, ANOVA *P*-value, and *Cp* did not show any irregularity (Table 5-6). Also, the residuals did not indicate any significant departure from the statistical assumption of homogeneity of variance (Figure 5-4).

Table 5-5 Parameter estimates for the model $R_\Delta = \beta_0 + \beta_1\rho + \beta_2\tau + \beta_3\rho^2 + \beta_4\tau^2 + \beta_5\rho\tau$. R_Δ is the plan CRW risk difference between the CRW model components and the control, ρ is percent of farm area with high risk fields and τ is percent of farm area as target. (R) Field constraint difference. (C) Crop constraint difference. (RC) Both constraints difference. (Cm) is the difference in R_p between both constraints and the field constraint.

Model Parameter	*R*	*C*	*RC*	*Cm*
β_0	0.23746	-0.05256	0.06644	0.07965
β_1	-0.01968	0	-0.02036	-0.00287
β_2	-0.03990	0.02662	0	0.03062
β_3	0	0	0	0
β_4	0.00040	-0.00029	0	-0.00034
β_5	0.00021	0	0.00016	0

Table 5-6 Regression statistics for the model $R_\Delta = \beta_0 + \beta_1\rho + \beta_2\tau + \beta_3\rho^2 + \beta_4\tau^2 + \beta_5\rho\tau$. R_Δ is the plan CRW risk difference between the CRW model components and the control, ρ is percent of farm area with high risk fields and τ is percent of farm area as target. (R) Field constraint difference. (C) Crop constraint difference. (RC) Both constraints difference. (Cm) is the difference in R_p between both both constraints and the field constraint.

Statistic	*R*	*C*	*RC*	*Cm*
R^2	0.926	0.773	0.864	0.668
ANOVA *P*-value	0.0001	0.0001	0.0001	0.0001
Mallow's *Cp*	4.351	2.330	3.858	2.359

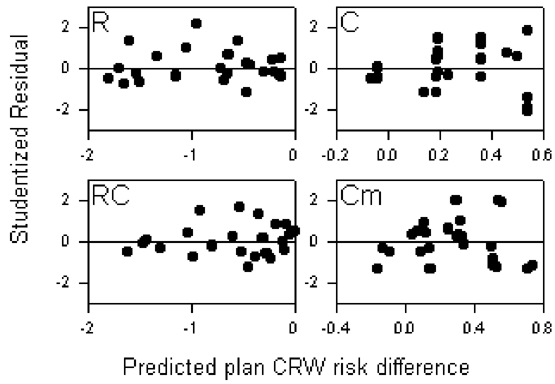

Figure 5-4 Studentized residuals versus expected plan CRW risk difference for the model $R_\Delta = \beta_0 + \beta_1\rho + \beta_2\tau + \beta_3\rho^2 + \beta_4\tau^2 + \beta_5\rho\tau$. R_Δ is the plan CRW risk difference between the CRW model components and the control, ρ is percent of farm area with high risk fields and τ is percent of farm area as target. (R) Field constraint difference. (C) Crop constraint difference. (RC) Both constraints difference. (Cm) Difference between both constraints and field constraint.

Looking at the differences in CRW risk (Figure 5-5), it is apparent that only the field constraint component and the full model presented a reduction in CRW risk. The crop constraint did not reduce the risk as expected, on the contrary, it increased the risk. For the field constraint, the highest reduction occurred when ρ was higher than 60 and τ was less than 60 percent. The reduction decreased as ρ decreased and τ either decreased or increased. The reduction in R_p was lower at high τ values because the field constraints had to be relaxed to comply with the target goal. When the target acreage was higher than 95, there was no gain in reduction. Instead, the R_p was slightly higher than the control. There was also a small increase in R_p as compared to the control for τ and ρ less than 10 percent.

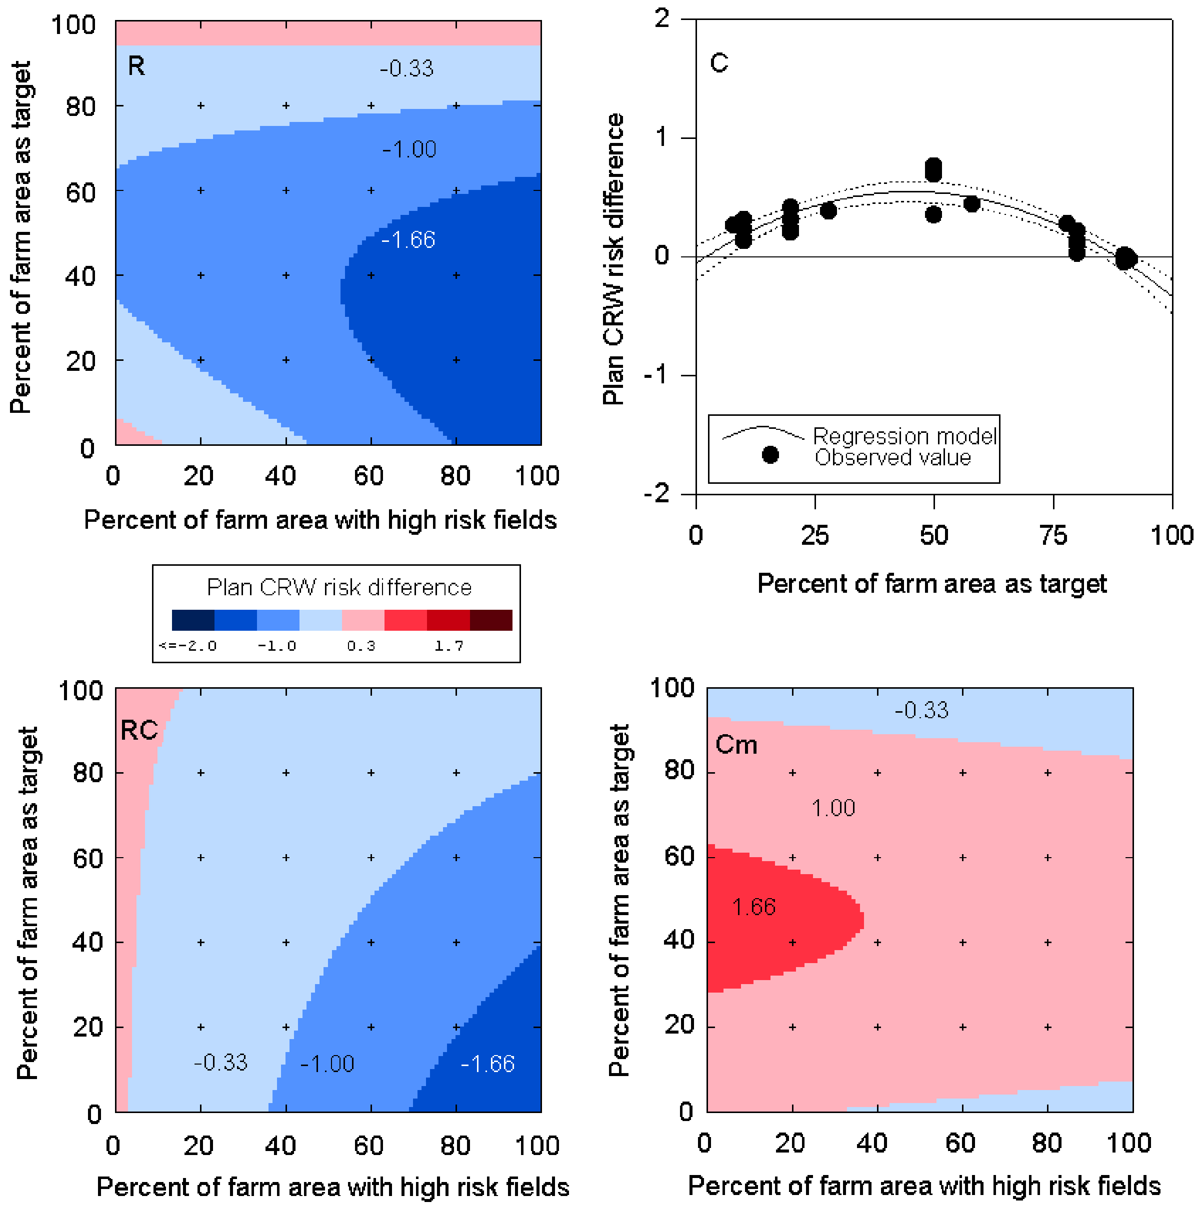

Figure 5-5 Response surface for the difference in plan CRW risk (R_{Δ}) against the control. Negative values indicate a reduction in CRW risk, as compared to the control. (R) Field constraint difference. (C) Crop constraint difference. Dotted lines indicate a 95 percent confidence interval. (RC) Both constraints (full CRW model) difference. (Cm) Difference between both constraints and field constraint. Numbers in regions indicate midvalue.

In the case of the crop constraint, there was no effect of ρ; therefore, R_p depended on τ only (Figure 5-5C and Table 5-5). There was no estimated reduction in R_p except when τ was less than five and τ higher than 85. In fact, the highest increase in R_p occurred in the middle (50 percent) of the target acreage across all values of ρ. Apparently, the crop constraint is not reducing the risk as expected. The differences for the full model (both constraints) show reductions lower than the field constraint alone, indicating that the crop constraint is overriding the effect of the field constraint. The marginal contribution of adding the crop constraint to the model when the field constraint is already present is estimated as the difference in R_p between both constraints and the field constraint (Table 5-5, Figure 5-5Cm). It is apparent that the addition of the crop constraint to the planner already applying the field constraint is ineffective. It actually increased risk across most of the response surface (Figure 5-5Cm). Only at extreme values of target acreage (τ) was there a positive effect.

Examining the distribution of corn (grain and silage) over time and across fields (Figures 5-6 and 5-7) helps explain the response in CRW risk for the different model components (constraints) and the control. For the control, there was no restriction based on CRW risk on which rotations were assigned to fields. For example, high risk rotations were assigned to three high risk fields (fields five, 11, and 12) for corn silage (Figure 5-6[Ctrl-S]) and medium risk rotations were assigned to two high risk fields (fields two and four) for corn grain (Figure 5-6[Ctrl-G]). The overall corn distribution for the control is in Figure 5-7[Ctrl-2] and has a risk rating of 6.9. With the field constraint applied (Figure 5-6[R-S]), rotations selected for high risk fields were lower risk rotations, except for fields five and 13 (in those fields, the constraint was relaxed as the planner attempted to comply with the goal). In addition, lower risk rotations for corn grain were selected compared to the control (Figure 5-6[R-G]). The whole corn sequence is in Figure 5-7[R-2] and has a risk rating value of 5.8, less than that of the control.

When the crop constraint was applied, fields 2, 11, and 12 were not permitted to be planted with the target crop, corn silage (Figure 5-6[C-S]). However, nothing prevented the planner from assigning corn grain to those fields, because the field constraint was not enforced in this treatment. As shown (Figure 5-6[C-G]) the planner did assign medium risk corn grain rotations to those fields, and the CRW risk was increased. Furthermore, because no corn silage could be planted in the high risk fields, the remaining fields were assigned rotations more intense in corn production to fulfill the target goal (Figure 5-6[C-S]). Thus, the net effect of the crop constraint was detrimental. It increased CRW risk to a value of 7.3, higher than for both the control and the field constraint.

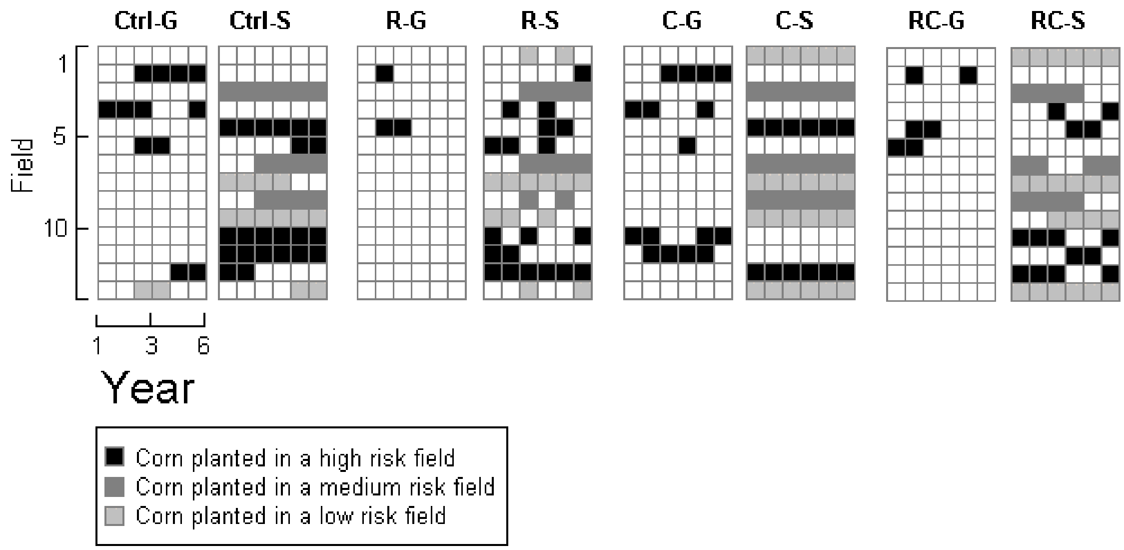

Figure 5-6 Corn distribution across fields and planning years for selected plans. (Ctrl) Control. (R) Field constraint. (C) Crop constraint. (RC) Both constraints. (1) Corn grain. (2) Corn silage (target). Annual target acreage was 58 percent and high risk field area was 55 percent of the farm's cropland.

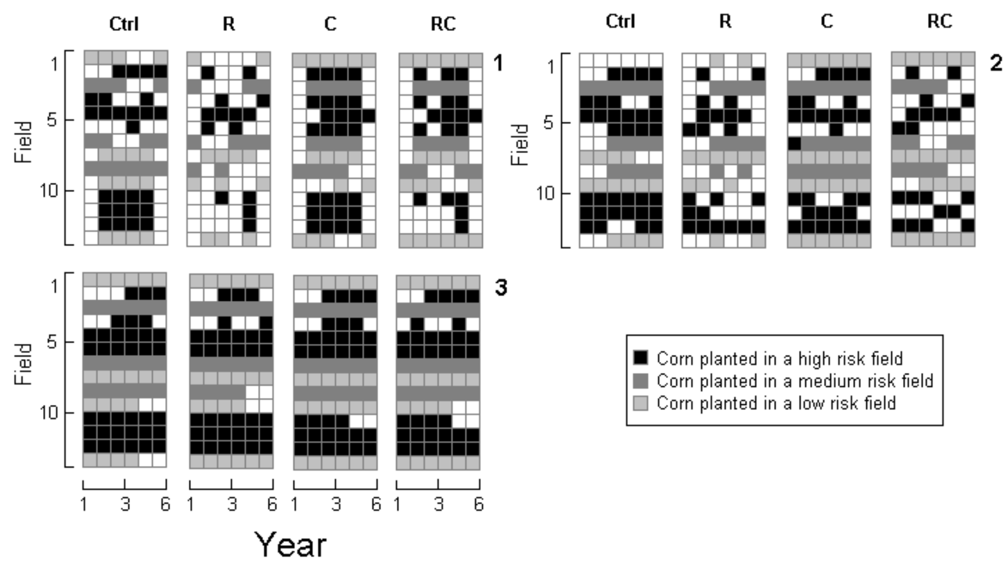

Figure 5-7 Corn (corn grain and corn silage combined) distribution across fields and planning years for selected plans and different annual target acreage. (1) Target acreage was 8 percent, (2) target acreage was 58 percent, (3) target acreage was 91 percent. (Ctrl) Control. (R) Field constraint. (C) Crop constraint. (RC) Both constraints. Percent of farm area with high risk fields was 55.

With both constraints enabled, the planner eliminated corn silage from fields two and six and assigned low risk corn grain rotations (Figure 5-6[RC-S]). However, three high risk fields (five, 11, and 13) ended up with medium risk corn grain rotations (Figure 5-6[RC-G]). The whole corn sequence is in Figure 5-7[RC-2] and has a risk rating of 6.0, intermediate between the control and field constraint.

The same relationship between treatments and the control just discussed also occurred at low target acreage (Figure 5-7[1]), but at high target acreage, the resulting plans were similar across all treatments (Figure 5-7[3]). This seems to be due to the planner having to relax its constraints to meet very high target acreages (Figure 5-8). Such relaxation permits the planner to select medium or high risk rotations for high risk fields, therefore increasing the plan CRW risk. At intermediate target acreages (about 50 percent), the planner does not relax the field constraints and was able to select rotations to match the fields, so that CRW risk is effectively reduced. At low target acreage, the limited amount of corn planted keeps the inherent risk low (Figure 5-7[1]) and limits any possible reduction.

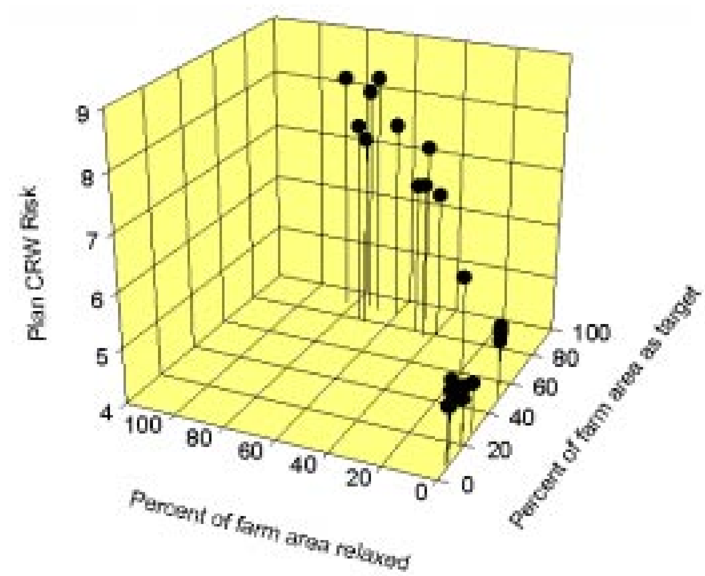

Figure 5-8 Plan CRW risk (R_p) as a function of relaxed field acreage and target acreage for the field constraint. Each point represents an average value of 20 plans.

Profit Analysis

The net profit for corn silage is in Table B-2 (Appendix B). All the statistical models used to represent the profit differences were significant at $P < 0.01$ except the full CRW control model, which was significant at $P = 0.0326$ (Table 5-7). The best fit corresponds to the crop constraint component ($R^2 = 0.83$). Examination of the residuals did not show any departure of the homogeneity of variance assumption (Figure 5-9). Because of the low R^2 value for the field constraint difference (0.58, Table 5-7), the response is presented as an interpolated grid, using the inverse method (Dewey, 1988). The regression model parameters are presented in Table 5-8; in the case of the full model (both constraints), there is no effect of the farm acreage with high risk fields on the response.

Table 5-7 Regression statistics for the model $\pi_\Delta = \beta_0 + \beta_1\rho + \beta_2\tau + \beta_3\rho^2 + \beta_4\tau^2 + \beta_5\rho\tau$. π_Δ is the corn silage net profit (US$/ Mg) difference against the control, ρ is percent of farm acreage with high risk fields and τ is percent of farm's cropland as target. (R) Field constraint difference. (C) Crop constraint difference. (RC) Both constraints difference.

Statistic	*R*	*C*	*RC*
R^2	0.580	0.831	0.183
ANOVA *P*-value	0.0001	0.0001	0.0326
Mallow's *Cp*	2.680	2.346	2.664

Table 5-8 Parameter estimates for the model $\pi_\Delta = \beta_0 + \beta_1\rho + \beta_2\tau + \beta_3\rho^2 + \beta_4\tau^2 + \beta_5\rho\tau$. π_Δ is the corn silage net profit (US$/ Mg) difference against the control. ρ is percent of farm area with high risk fields and τ is percent of farm area as target. (R) Field constraint difference. (C). Crop constraint difference. (RC) Both constraints difference.

Model Parameter	*R*	*C*	*RC*
β_0	0.74951	-0.01518	-0.28915
β_1	0	-0.01103	0
β_2	0	0	0.00386
β_3	0.00010	0	0
β_4	0	0	0
β_5	-0.00010	0.00012	0

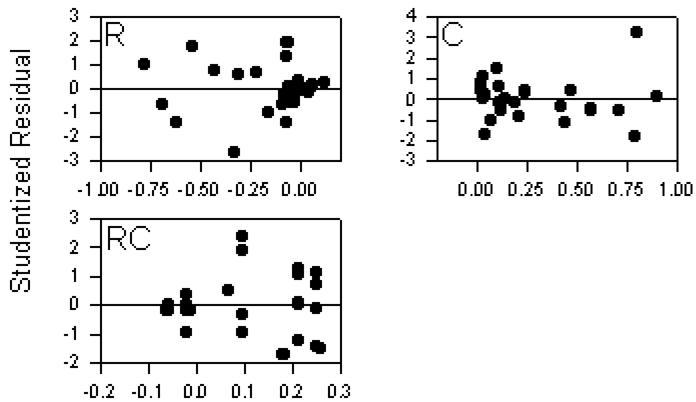

Figure 5-9 Studentized residuals versus expected corn silage net profit (US$/ Mg) difference for the model $\pi_{\Delta} = \beta_0 + \beta_1\rho + \beta_2\tau + \beta_3\rho^2 + \beta_4\tau^2 + \beta_5\rho\tau$. π_{Δ} is the corn silage net profit (US$/ Mg) difference against the control. ρ is percent of farm area with high risk fields and τ is percent of farm area as target. (R) Field constraint difference. (C) Crop constraint difference. (RC) Both constraints difference.

Including the field constraint had a generally positive effect on profit per yield unit (Figure 5-10R). Positive differences (increase in profit) correlate closely with the regions where the reduction of CRW risk is highest while reductions in profit correlate with increased CRW risk (Figure 5-5R). On the other hand, when the crop constraint was compared to the control, profits were lower for most of the experimental region (Figure 5-10C). As noted before, the elimination of some fields by this constraint forced the planner to plant more corn in the remaining fields while permitting non-target corn (i.e. corn grain) in the excluded fields. The overall effect is to increase the CRW risk and hence the number of insecticide applications. The increase in the use of insecticides should account for the increase in costs and the reduction of profit. With both constraints implemented, profit did not vary with field risk, varying only with target acreage (Table 5-8, Figure 5-10RC). In Figure 5-10RC, most of the points fall below the zero value, indicating a reduction in profit. This is probably due to the increase in CRW risk seen when implementing the crop constraint together with the field constraint as previously discussed (see Figure 5-5Cm).

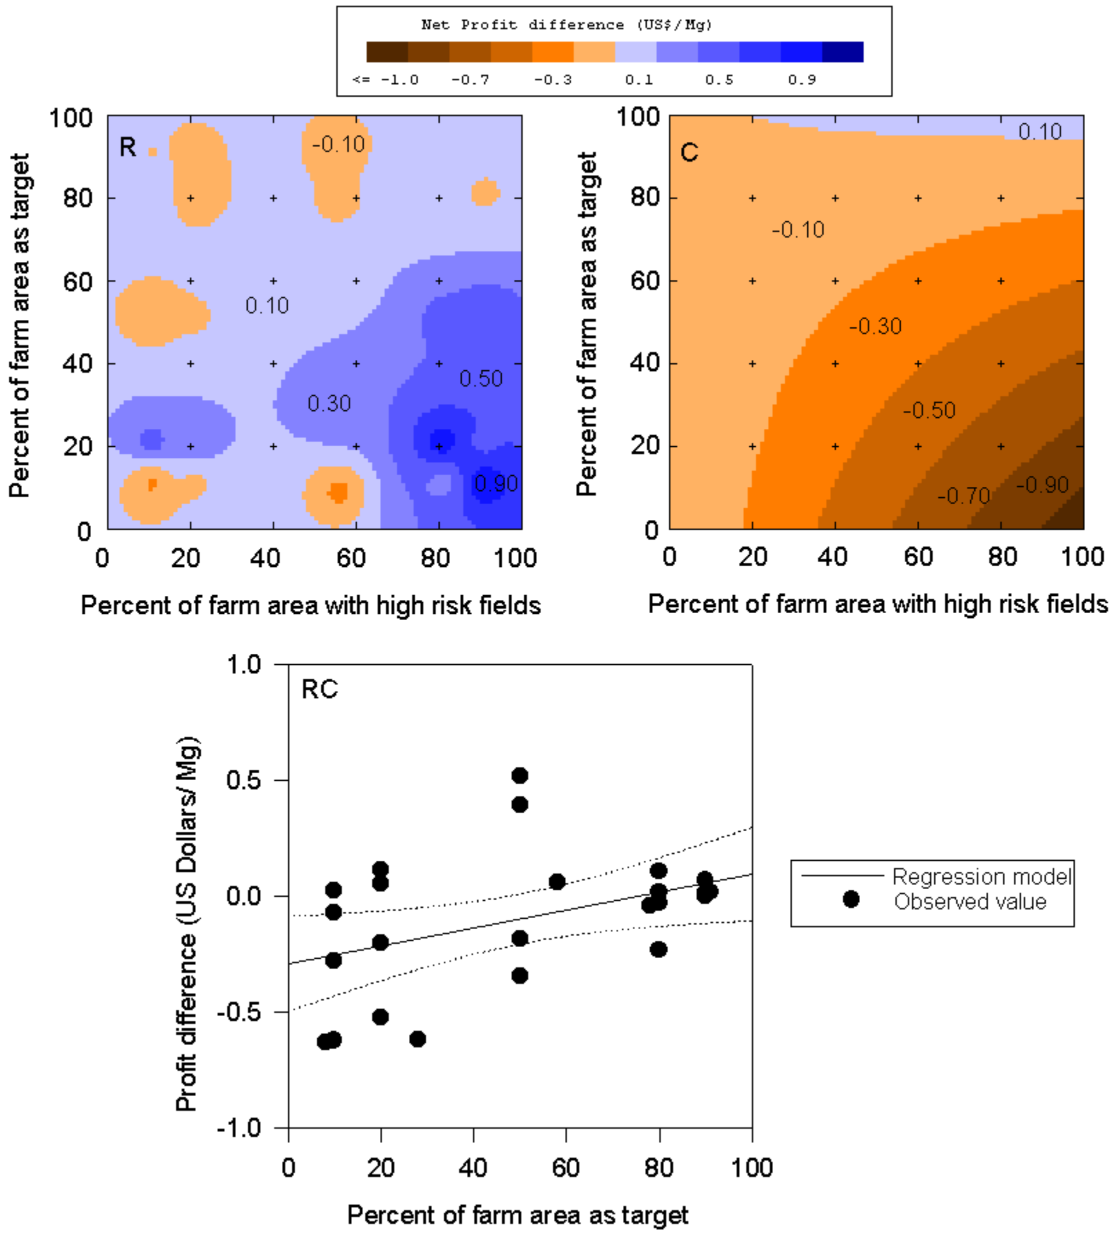

Figure 5-10 Response surface for the corn silage net profit difference (π_{Δ}, US$/ Mg). Positive values indicate an increase in profit, as compared to the control. (R) Field constraint difference. (C) Crop constraint difference. (RC) Full model difference. Dotted lines indicate a 95 percent confidence interval. The response surface for the field constraint difference is presented as an interpolated grid, using the inverse method (Dewey, 1988).

While increasing profit is a key goal, the best plans should also minimize environmental risks and risks of catastrophic pest losses by lowering CRW risk. The relationship between the corn silage profit and the plan CRW risk in presented in Figure 5-11 computed as observed differences against the control. The results indicate that adding the field constraint (R) increased the profit and reduced the risk. The crop constraint (C) increased the risk and reduced the profit. The effect of both constraints (RC) was to reduce the risk as the field constraint but did not increase the profit. This is probably because the corn silage occurred as 2nd-year corn more frequently than for the field constraint alone.

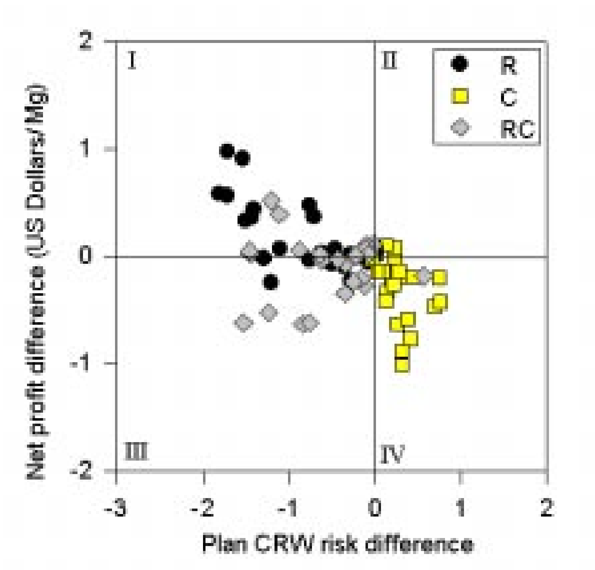

Figure 5-11 Corn silage net profit (US$/ Mg) difference as a function of the plan CRW risk difference. Differences measured against control values. Quadrants I and II indicate an increase in profit; Quadrants I and III indicate reduction in CRW risk. (R) Field constraint difference. (C) Crop constraint difference. (RC) Both constraints (full CRW model) difference.

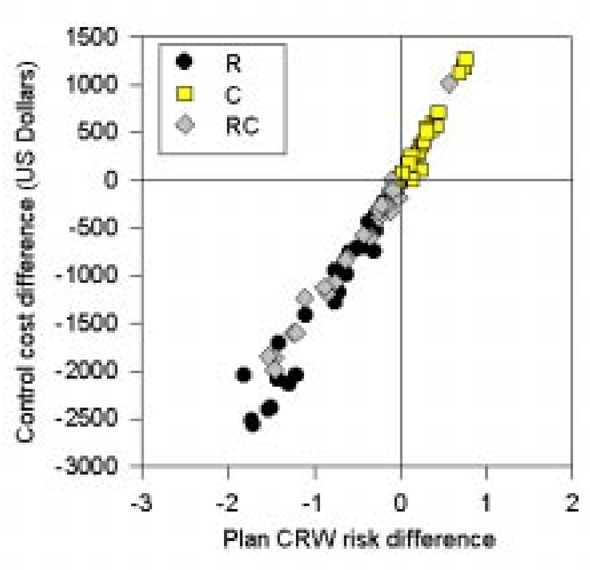

Figure 5-12 Relationship between the chemical control cost per plan and CRW risk. Differences against the control. (R) Field constraint difference. (C) Crop constraint difference. (RC) Both constraints (full CRW model) difference.

Since high CRW risk is associated with higher estimated pest control costs, profits should be correlated with CRW risk. In fact, this relationship did hold in the plans generated by all treatments (Figure 5-12). The relationship is clear: the greater the risk reduction the greater the reduction in control costs.

Looking at the overall plan net profits (rather than estimates per yield unit) shows some interesting differences. First, although the regression models were statistically significant ($P < 0.05$, Table 5-9) the residuals (Figure 5-13) showed a systematic pattern, suggesting the need of additional terms, and some points are outside of the critical value of 2 (Draper and Smith, 1981). In addition, the Mallow's *Cp* values were higher than the number of parameters in the model (Table 5-9). The only acceptable regression model was the difference in net profit between both constraints (full CRW model) and the control. Thus, except for the last model, the rest were considered unsuitable and the response surfaces are presented as an interpolated grid, using the inverse method again (Dewey, 1988).

Table 5-9 Regression statistics for the model $\Psi_\Delta = \beta_0 + \beta_1\rho + \beta_2\tau + \beta_3\rho^2 + \beta_4\tau^2 + \beta_5\rho\tau$. Ψ_Δ is the plan net profit difference against the control, ρ is percent of farm area with high risk fields and τ is percent of farm area as target. (R) Field constraint difference. (C) Crop constraint difference. (RC) Both constraints (full CRW model) difference.

Statistic	*R*	*C*	*RC*
R^2	0.401	0.215	0.772
ANOVA *P*-value	0.0036	0.0195	0.0001
Mallow's *Cp*	13.922	18.189	4.597

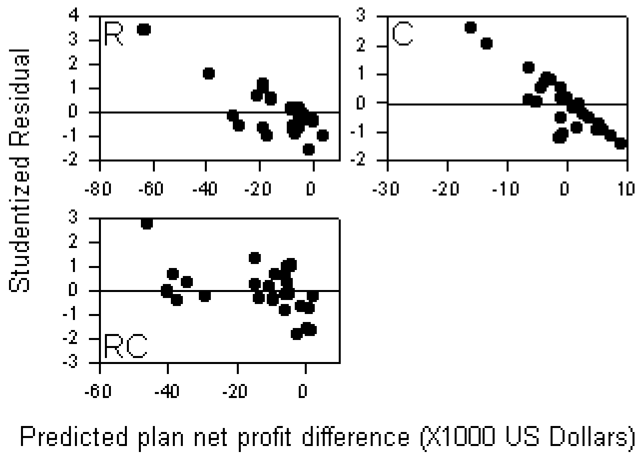

Figure 5-13 Studentized residuals versus expected plan net profit (Ψ_Δ) difference for the model $\Psi_\Delta = \beta_0 + \beta_1\rho + \beta_2\tau + \beta_3\rho^2 + \beta_4\tau^2 + \beta_5\rho\tau$. Ψ_Δ is the plan net profit difference against the control, ρ is percent of farm area with high risk fields and τ is percent of farm area as target. (R) Field constraint difference. (C) Crop constraint difference. (RC) Both constraints (full CRW model) difference

When the field constraint alone was applied in the planner, most plans generated showed a reduction in profit, except at low ρ values. The reduction in profit was highest for ρ higher than 80 percent and τ less than 50. As τ approaches the 90 percent value, the difference declined (Figure 5-14R). This negative effect of the field constraint on the plan profit could be due to the reduction in the number of crop rotations available after application of the constraint (Figure 5-15). This domain size reduction affects both CRW risk and plan profit. The plan CRW risk is reduced because of the elimination of high risk rotations; many of these are also rotations that contribute to higher net profits.

Implementing the field constraint often did reduce CRW risk, but usually at the expense of profit (Figure 5-16). Using just this constraint, there was an inverse relationship between control of CRW and the resulting profit.

With just the crop constraint implemented, again plan profits tended to decline compared to the control, though there were significant regions showing increased profit (Figure 5-14C). The effect of domain size on the profit and risk is presented in Figure 5-17. As in the case of the field constraint, applying the crop constraint also reduced the domain size (this time for field combinations) and the profit but increased the CRW risk. The counter-intuitive increase in the risk was explained earlier. The reduction in profit is probably due to increased control costs for CRW (Figure 5-12). Looking at the interaction between the crop constraints impacts on CRW risk and profit (Figure 5-16), we see that CRW risk was almost always increased, while profit was sometimes up and sometimes down. Never was profit increased while risk was reduced. Overall, this constraint performed poorly.

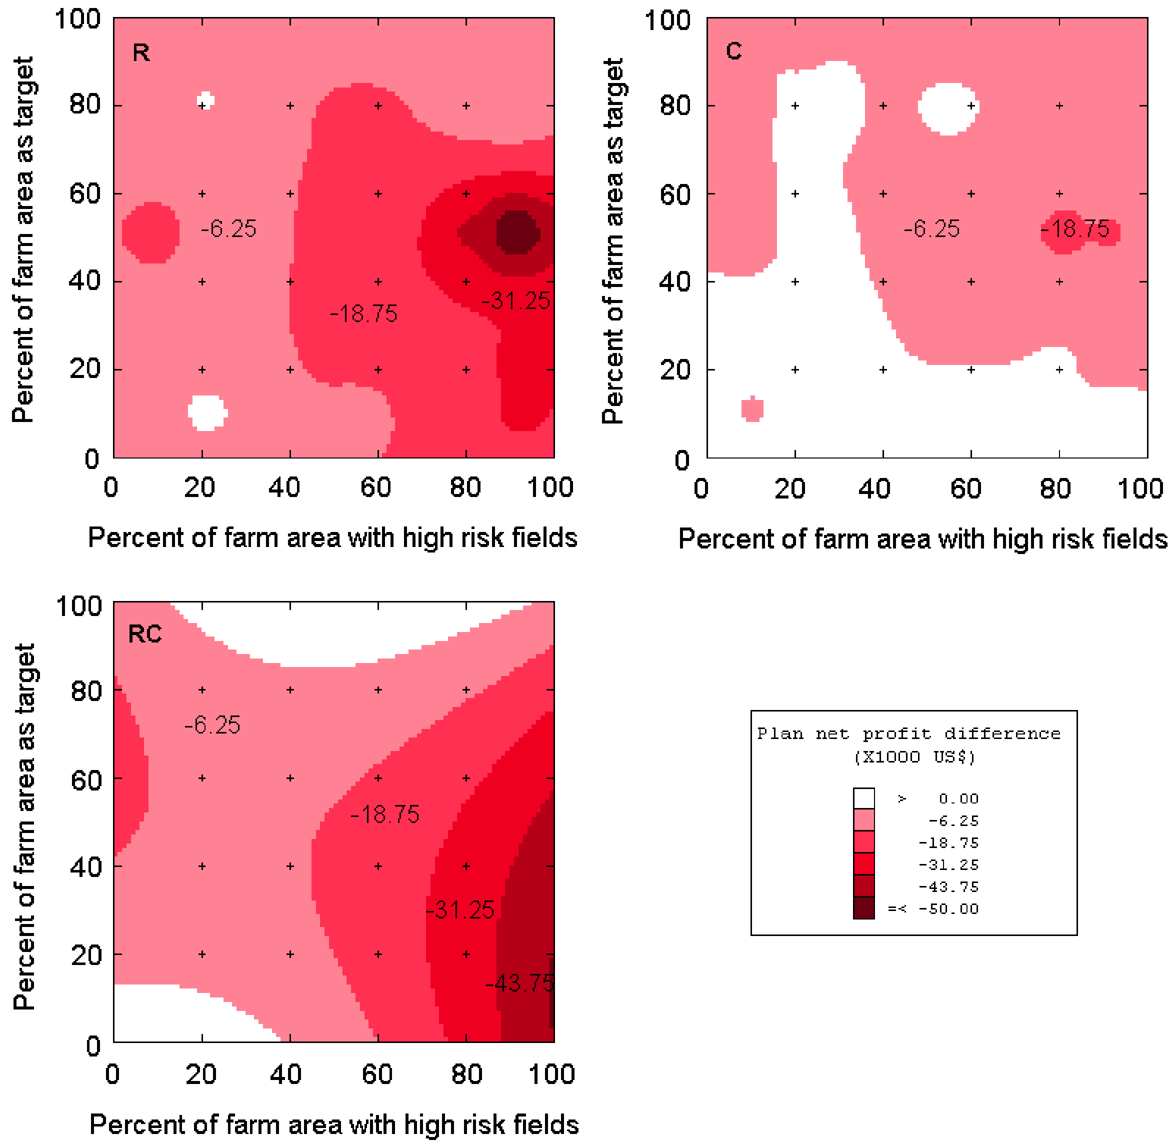

Figure 5-14 Response surface for the differences in plan net profit (Ψ_Δ) between the model components and the control. Negative values indicate a reduction in profit due to the application of the constraints. (R) Field constraint. (C) Crop constraint. (RC) Both constraints (full CRW model). Grids R and C were computed as interpolated values using the inverse method (Dewey, 1988). Grid RC was computed using the quadratic model: Ψ_Δ = 8846.2 -770.3τ - 5.9ρ^2 + 6.1τ^2 + 6.4$\rho\tau$. Numbers in regions indicate midvalue.

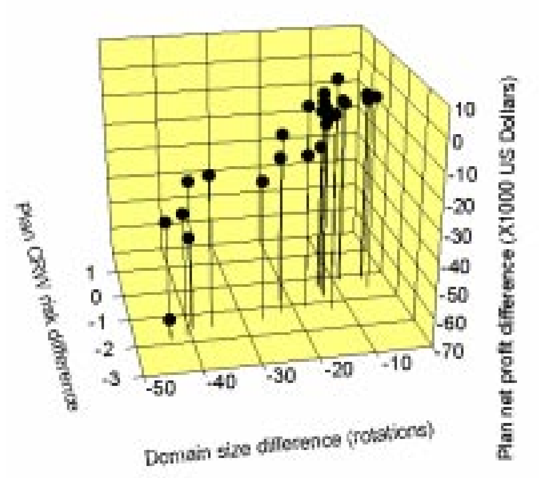

Figure 5-15 Relationships between domain size, plan CRW risk, and plan net profit. Values are shown as differences between the field constraint and the control.

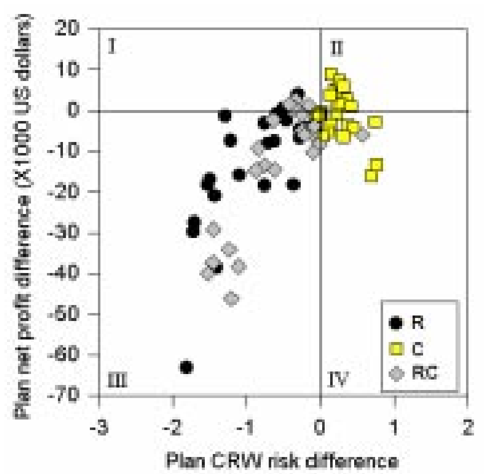

Figure 5-16 Relationship between plan CRW risk and plan net profit. Quadrants I and II indicate an increase in profit, quadrants I and III indicate a reduction in CRW risk, as compared to the control. Data reported are the observed differences between profit calculated for the three treatments: (R) field constraint, (C) crop constraint, (RC) both constraints, and the control.

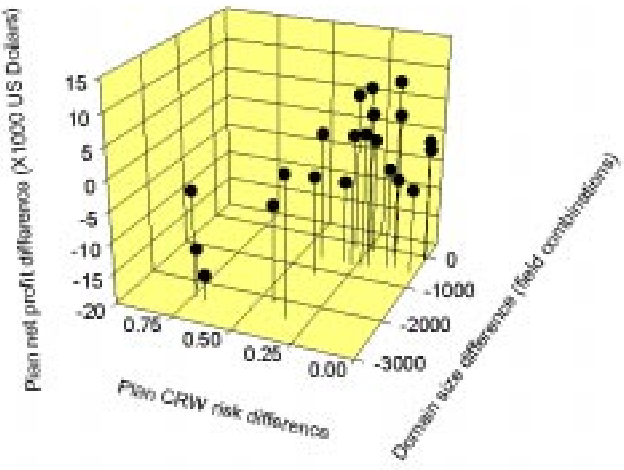

Figure 5-17 Relationship between domain size, plan CRW risk, and plan net profit. Observed differences between the crop constraint and the control.

When both constraints were applied, profit generally decreased (Figure 5-14RC). The pattern observed in the figure is a saddle shape (Box and Draper, 1987), showing only a slight gain in profit in two regions: when both τ and ρ have low values (τ less than 10 percent and ρ less than 40 percent) and when τ is high (τ higher than 80 percent) and ρ is near 50 percent. Profits were reduced for high risk situations, particularly for lower target acreages.

From Figure 5-16 it is difficult to distinguish the effects of the two constraints together from the effects of the field constraint alone, they show a similar trend. To see differences between the two, one can look at the marginal contribution of the crop constraint when the field constraint is already present. The effect was estimated as the difference between the field constraint and both constraints (Figure 5-18). From this figure, one can quickly see that the addition of the crop constraint to the planner already incorporating the field constraint almost always increased CRW risk, while having no definitive effect on profit. Combining this result with the results of the crop constraint alone, it is clear that this constraint is detrimental to the planning process and should not be used at all.

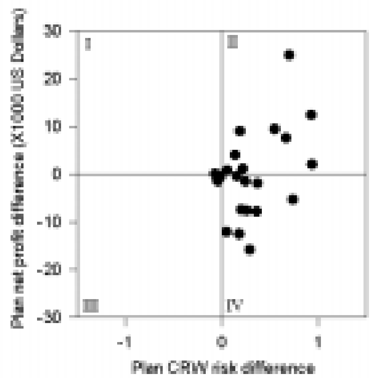

Figure 5-18 Relationship between plan CRW risk and plan net profit. Values indicate the marginal contribution of the crop constraint when the field constraint is already enforced. Quadrants I and II indicate an increase in profit and quadrants I and III indicate a reduction in CRW risk.

To this point it seems that implementing the field constraint reduces the CRW risk but decreases the profit at the plan level. However, the results at the plan level were also affected by factors related to the experimental design. For example, even within a given target acreage range there were differences in the amount of corn planted between the control and a treatment for plans of the same target class. Because corn is both risky to plant and highly profitable, the allowed acreage range of 4.05 ha led the planner to plant more corn in the control trials. In addition, it should be considered that the performance was evaluated using a single target crop. Results using other target crops could be different. One assumption of the planner (see Chapter 3) is that the system is supposed to be used to plan for crops to support dairy farm operations. Thus, besides corn silage, crops like alfalfa, pasture, and rye are usually considered for this kind of enterprise (Parsons, 1995). The next section, application of the model, tests the ability of the planner to find alternative plans for several target crops while limiting the risk of CRW outbreaks.

Model Application

<u>Experimental Conditions</u>

Ideally, the planner should be able to generate alternative plans that reduce CRW risk while maintaining income (plan net profit) and meeting production goals. A second farm was used to search for plans that satisfied production goals and reduced CRW risk. The analysis of this second farm helps both to verify that observations made previously were not dependent on a single farm's properties, and to test the planner's ability to produce useful plans that mitigate CRW risk. Farm-2 represented a 50-cow dairy, with about 20 ha of cropland and 22 ha of pasture and hay; in total, the farm has a size of 55.5 ha. The farm included 10 fields, two of which were designated HEL, divided into two tracts of five fields each. Soil composition is described in Chapter 6 and other field properties are listed in Table A-2 (Appendix A). The benchmark was done by hand to satisfy food requirements for the dairy enterprise. Four crops were selected as targets: corn silage, alfalfa, pasture, and rye silage; target acreage requirements are listed in Table 5-10. The crops and target acreages were based on a survey by Parsons (1995). The area of corn silage is about 19.8 ha or 35.7 percent for the benchmark plan (Table 5-10). Fields were randomly assigned high risk scores to account for 46.4 percent of the total farm area. This value was located where a medium reduction of CRW risk was found for the first farm (Figure 5-5R). Planning time was set to 10 minutes. The CRW field constraint was enabled for all fields. The plan net profit required was set to 100 percent of the benchmark, and ten plans were requested from the planner. Plan profit and control costs were computed as described in the section about evaluating model performance.

Table 5-10 Target crop, acceptable annual target acreage (ha), and benchmark values for farm-2.

Target crop	*Target acreage range (ha)*	*Benchmark values (ha)*
Corn silage	19.4-21.4	19.8
Alfalfa	7.3-9.3	7.8
Pasture	15.4-18.2	15.5
Rye silage	19.4-21.4	20.1

Results and Discussion

No plans were found until the field constraint was relaxed for fields one and eight. The planner then returned plans with an average plan CRW risk of 5.1 compared to 5.3 of the benchmark, a slight reduction in CRW risk. The crop sequences for corn grain and silage are presented in Figure 5-19 for the benchmark and one of the alternative plans. Note that the planner put a medium risk rotation into field eight and a high risk rotation in field one. As Table 5-2 indicates, these combinations are not acceptable but were allowed because the field constraints were relaxed. This result emphasizes the importance of controlling how the constraints are enforced.

The reduction in CRW risk in the ten plans generated was also reflected in the control costs. The control cost for the benchmark plan was US$1,830 compared to US$1,663 (± 1.23 SE) for the alternative plans (Figure 5-20A), a reduction of 9.1 percent. In addition, the plan net profit was 7.8 percent higher than the benchmark (Figure 5-20B) on average. The increase in profit was only partly due to the reduction in control costs.

In summary, the planner was able to produce alternative plans with higher profits and lower CRW risk. This result is analogous to the finding by Youngman et al. (1993) that rotating corn with sorghum provided similar yields compared to continuous corn while reducing CRW damage. Here, the corn acreage satisfied the target while providing better profit than the benchmark. It has also been pointed out by Lazarus and Swanson (1983) that crop rotation reduces pest damage and helps reduce risk by diversifying the farm. Therefore, the alternative plans not only reduce CRW risk but also may help to maintain economic stability within the farm, a desired attribute of sustainable agriculture (Smith and McDonald, 1998).

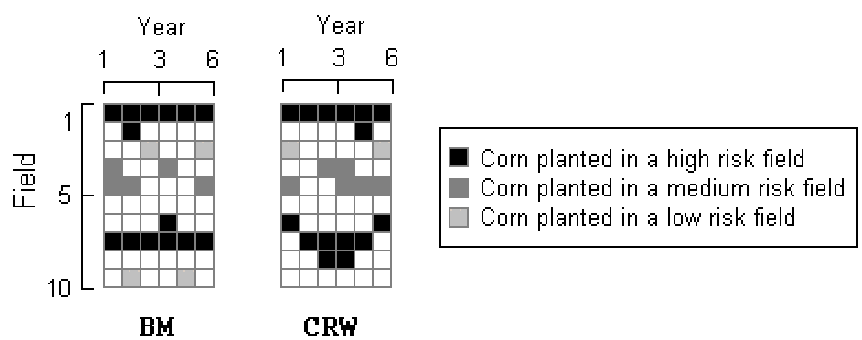

Figure 5-19 Corn distribution for the benchmark plan (BM) and a CRW control plan (CRW).

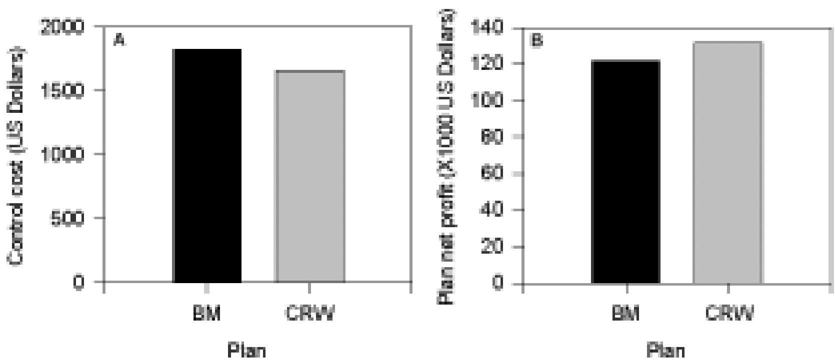

Figure 5-20 (A) CRW control cost for the benchmark plan and CRW control plans. (B) Plan net profit for the benchmark plan and CRW control plans.

The fact that profits increased for the second farm might seem contradictory to what was found in the performance experiments earlier. However, the experiments were different. This experiment specifically eliminated plans of inferior profit by enforcing a plan profit constraint. The performance evaluation disabled the plan profit constraint to assess overall trends in the planner's performance. Another complicating factor was that in the last application, the planner had to satisfy acreage requirements for four crops, thus increasing the restrictions on land use and availability of fields to plant corn. A single crop was used in the performance experiments. Finally, the sizes of the farms differed somewhat. The second farm included ten fields versus fourteen in the farm one used in the performance experiments. The ten fields had to be distributed over four target crops, limiting the ability of the planner to freely assign fields to corn. Thus, the two experiments can be considered complementary. The performance experiment was an internal evaluation to determine where reduction in CRW risk can be expected. The last study was a practical application of the system to determine if it could find alternative plans without affecting profit.

Summary and Conclusions

A preventive IPM model was designed and implemented as constraints within the planning subsystem. The corn rootworm (*Diabrotica virgifera and D. barberi*) complex was used as a case study. The model was based on three assumptions: first, that the crop rotation risk to CRW damage depends on the frequency a corn crop is planted in a given year when corn was planted the previous year; second, that prolonged diapause had a negligible effect and third, that fields have different inherent risks to CRW damage.

Based on the previous assumptions, two unary constraints were developed, a field constraint to select CRW-tolerant rotations and a crop constraint to select low risk field combinations. To apply the field constraint, rotations were classified into three categories: low, medium, and high risk, depending on the frequency a corn crop is planted in a given year when corn was planted the previous year. Decision rules were provided to select rotations for any given field based on both CRW rotation risk and CRW field risk. For the crop constraint, the field combinations were assigned a CRW risk based on the risk of the fields that compose the combination and weighed by the field acreage. The decision rule in this case was to eliminate as many field combinations as possible.

The model components (constraints) were tested by planning for a farm in which two factors were varied: the percent of farm's cropland with high risk fields and the percent of the farm's acreage set as a target for corn silage. Corn silage was the only target used. In addition, a second study was conducted on a second farm to determine the ability of the planner to find alternative plans for dairy operations while minimizing CRW risk and maintaining income.

The main conclusions of this study are:

- The crop constraint increased the risk of CRW by excluding high risk fields from planning. This caused the remaining fields to require rotations with a high frequency of corn planted after corn to satisfy the target acreage. This model component was not recommended for implementation.

- The field constraint component effectively reduced the CRW risk across most of the experimental region of target acreage and field risk intensity.
- The net profit of the target crop, corn silage, increased with the reduction of CRW risk. The increase occurred on a per yield unit basis. The increase in net profit was due to the reduction of the number of control applications.
- The reduction in CRW risk reduced the plan net profit by removing high profit rotations from planning.
- The application study demonstrates that it is possible to obtain reduction in CRW while maintaining production and economic goals.

Chapter 6 Reactive Planning Strategies for Nonpoint Source Pollution Control

Introduction

Incorporating any constraint into the planner requires that an evaluator also be included to assess whether the constraint is met for a particular case. However, because the planner must evaluate many thousands of options in any given run, these evaluators must be simplistic; at least, they must be very fast. Evaluating environmental and economic risks associated with implementing a whole farm plan, however, cannot always be done well using quick and simplistic models. To compensate, Stone et al. (1992) proposed using simulation models to evaluate those few plans generated by the planner as a final check. The challenge remains, however, of automating this kind of post-planning analysis and using the simulation results to improve the plans generated.

Here, I explore how a nonpoint source (NPS) pollution simulation model can be integrated into the planning process to improve the resulting plans. The NPS model used was ANSWERS, and the pollutant was soil detached from agricultural fields by rain and transported to waterways. ANSWERS can also simulate nutrient losses, but for simplicity, only soil erosion was considered here. ANSWERS represents a landscape as a grid of square, homogeneous cells, each containing its own set of parameter values required by the model. The size of the grid and the width of the cells can be varied to simulate large or small areas with fine or coarse resolution. This allows the inclusion of spatial variability such as crop cover, and soil and topographic properties. Like FARMSCALE (Wolfe et al. 1995), WFP is a decision support system that includes ANSWERS as the simulation model to evaluate management plans at the farm level. However, instead of creating plans manually, the user can create plans by using the planning module (CROPS-LT). The planning module, like CROPS, still uses an empirical model, the Universal Soil Erosion Equation (USLE), to obtain estimates of soil loss (Buick et al. 1992) and to estimate the risk of soil loss inherent in a crop rotation (Buick et al. 1992; Stone 1995).

Framework for Reactive Planning

A classic feedback control loop was used as the framework for turning the planning system into a reactive planner (Figure 6-1A). The classic control loop includes an autonomous agent that runs a planner, one or more sensors that monitor changes to the environment, and controllers within the agent to respond to changes perceived (Michie and Chambers, 1968; Wilkins, 1988). Feedback control systems are widely used in manufacturing (Hopgood, 1993) and by artificially intelligent robots to help them react appropriately in changing environments (Linney, 1991; Kabanza et al. 1997). Feedback control systems are also commonly found in natural systems with examples from human physiology to community ecology (Holland, 1992).

As applied to the whole farm planning problem, the control loop was implemented with a manual agent, the investigator/ user. The user runs the model, generating a set of farm plans to be used in place of a current or "benchmark" plan. The monitor was the simulation model, ANSWERS, used to simulate real-world effects of implementing a farm plan on the farm's soils. Results from the ANSWERS model were assessed by the user, and changes to the control variables in the planner were made if necessary. A repair-based version of the feedback control loop (Figure 6-1B) was also implemented. In this case, the monitor can identify a subset of satisfactory fields within an unsatisfactory plan. The planner keeps the partial plan and replans only for the remaining unsatisfactory fields. There is one potentially appealing feature of this repair-based approach. Removing part of the plan should decrease the replanning time by excluding some fields (see section on field exclusion in Chapter 4). This approach has been used in other problem domains such as scheduling of manufacturing operations (Linney, 1991).

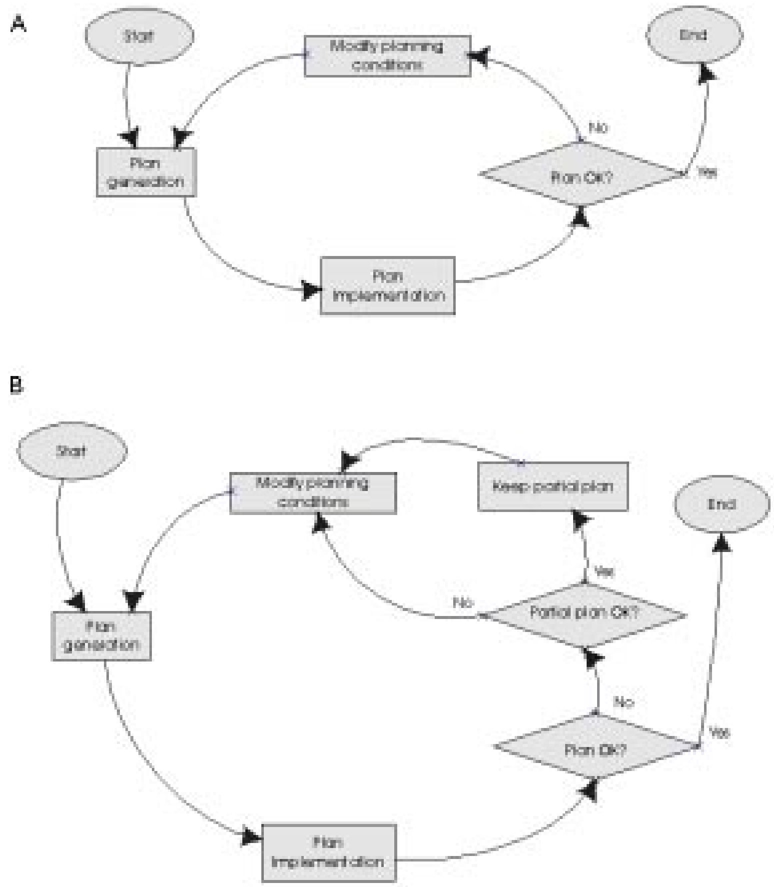

Figure 6-1 (A) Reactive planning procedure applied to farm planning. The planner generates plans that are implemented by an NPS simulation model. The pollutant values are then compared against benchmark values. If the proposed plans meet the expectations the search ends, otherwise it continues. (B) Repair-based reactive planning procedure. In this case, those fields in a plan that are satisfactory are kept, and the planner replans just for the remaining fields.

The goal of the reactive planning system was to reduce soil losses from the farm due to crop production. The relevant variable monitored from the ANSWERS simulations was the amount of soil leaving the farm. This measure of soil loss is different from what is calculated by the USLE model, which estimates soil detachment within a field. Because the most serious environmental consequences of soil erosion result from transport of sediment off site (Crutchfield et al. 1993), and because soil detached in one field may be deposited in another field, reducing the total soil leaving the farm was considered to be a more appropriate goal.

Experimental Conditions

Farm Description and Data Sources

Two artificial but representative farms were assembled from agricultural lands in Rockingham County, Virginia. Farm resources were assigned based on descriptions by Parsons (1995) to approach typical crop-livestock farms in the region. Farm-1 represented a typical 70-cow dairy, with about 30 ha of cropland and 30 ha of pasture and hay (Table 6-1). In all, the farm comprised 14 fields; each delineated spatially using field boundaries from a field boundary map of current farms obtained from the National Resources Conservation Service (NRCS) field office in Harrisonburg, Virginia. Four of the fields were considered “highly erodible land” (HEL) by NRCS. The fields were grouped into three distinct tracts of contiguous fields (Figure 6-2). Soils were composed mainly of Frederick silt loam and Allegheny fine sandy loam soil series; with slopes ranging from zero to 30 percent (Figure 6-3). These soils are typical for dairy farms in the Shenandoah Valley, representing a mix of some prime farmland and other soil classes (see Table 6-2 for soil description).

Table 6-1 Target crops, acceptable annual target range acreage (ha) and benchmark (BM) values for farms 1 and 2.

Target crop	*Farm-1*		*Farm-2*	
	Target range	*BM values*	*Target range*	*BM values*
Corn silage	28-32	30.0	19.4-21.4	19.8
Alfalfa	8-12	8.1	7.3-9.3	7.8
Pasture	18-22	19.2	15.4-18.2	15.5
Rye silage	28-32	30.0	19.4-21.4	20.1

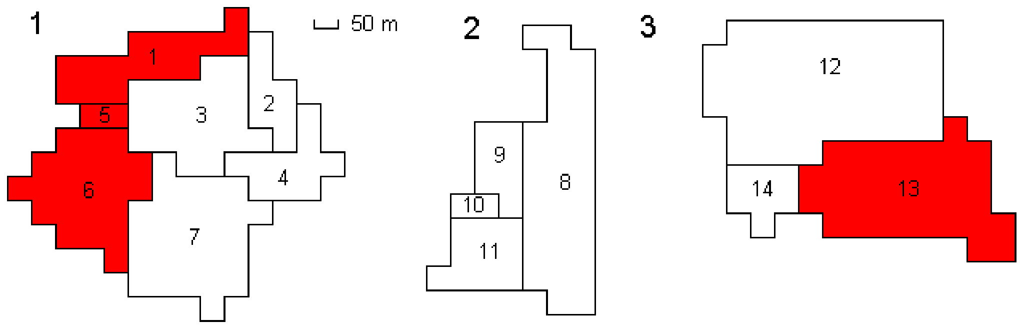

Figure 6-2 Field distribution for farm-1. The farm is composed of 14 fields distributed in three tracts. Shaded areas indicate “highly erodible land” (HEL) fields.

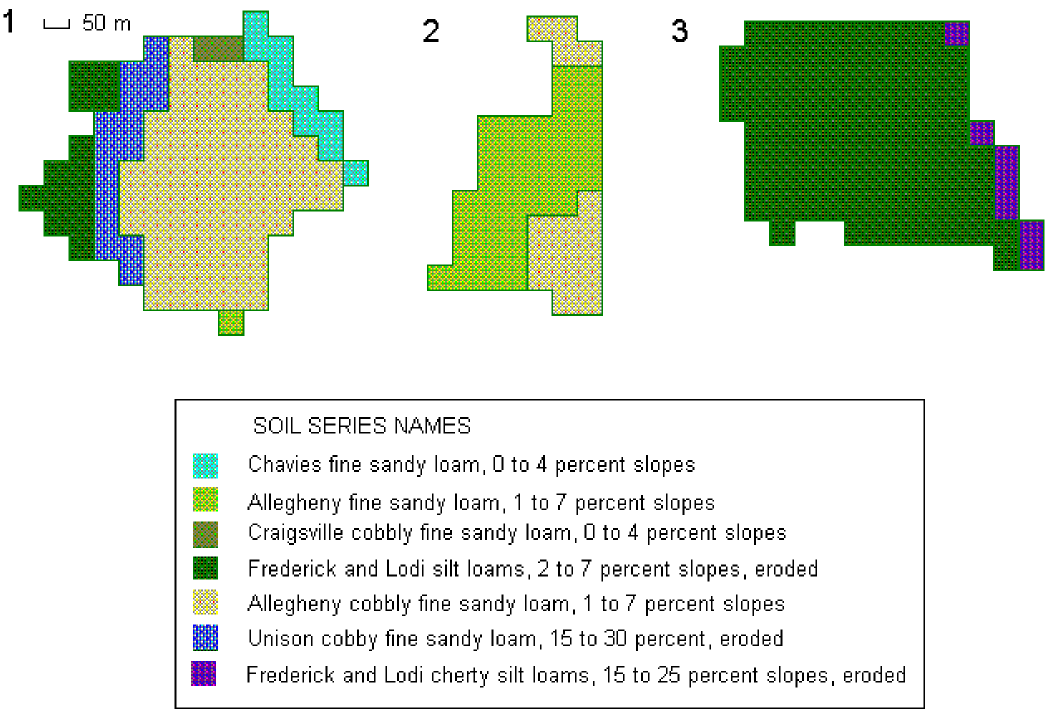

Figure 6-3 Soil type distribution for farm-1. Soil series are described in Table 6-2.

Table 6-2 Description of soils series from farms 1 and 2 (abridged from SCS, 1982).

Soil series name	*Description*
Allegheny fine sandy loam, 1 to 7 percent slopes	Appropriate to cultivated crops, pasture, and hay, with slopes less than 10 percent
Allegheny cobbly fine sandy loam, 1 to 7 percent slopes	Appropriate to cultivated crops, pasture, and hay, with slopes less than 10 percent
Buckton loam, 0 to 4 percent slopes	Prime farmland, best suited to produce food, feed, forage, fiber, and oilseed crops
Chavies fine sandy loam, 0 to 4 percent slopes	Prime farmland, best suited to produce food, feed, forage, fiber, and oilseed crops
Cotaco fine sandy loam, 0 to 7 percent slopes	Prime farmland, best suited to produce food, feed, forage, fiber, and oilseed crops
Craigsville cobbly fine sandy loam, 1 to 7 percent slopes	Good for small grains, pasture, and hay, and moderately appropriate for row crops
Edom silty clay loam, 2 to 7 percent slopes, eroded	Prime farmland, best suited to produce food, feed, forage, fiber, and oilseed crops
Frederick and Lodi silt loams, 2 to 7 percent slopes, eroded	Appropriate for cultivated crops, pasture and hay; it has slopes from two to seven percent
Frederick and Lodi cherty silt loams, 2 to 7 percent, eroded	Prime farmland, best suited to produce food, feed, forage, fiber, and oilseed crops
Frederick and Lodi cherty silt loams, 15 to 25 percent slopes, eroded	Not well suited for row crops, but appropriate for pasture and hay, most of the area is farmed; slopes higher than 10 percent
Guernsey silt loam, 2 to 7 percent slopes	Prime farmland, best suited to produce food, feed, forage, fiber, and oilseed crops
Unison cobbly fine sandy loam, 15 to 30 percent, eroded	Not very appropriate for cultivated crops but appropriate for pasture and hay; part of the area is farmed; slopes higher than 10 percent

Farm-2 represented a smaller, 50-cow dairy, with about 20 ha of cropland and 22 ha of pasture and hay (Table 6-1). The farm included 10 fields, two of which were designated HEL, divided into two tracts of five fields each (Figure 6-4). Soils were composed mainly of prime farmland (78 percent of total), and slopes ranged from zero to seven percent (Figure 6-5).

Benchmark plans were developed by hand to satisfy the acceptable annual target acreage ranges listed in Table 6-1 for both farms. The crops selected as targets were corn silage, alfalfa, pasture, and rye silage. The crops and appropriate target acreages were based on a survey by Parsons (1995) and adjusted to fit the dimensions of the current farms. Benchmark target values are in Table 6-1.

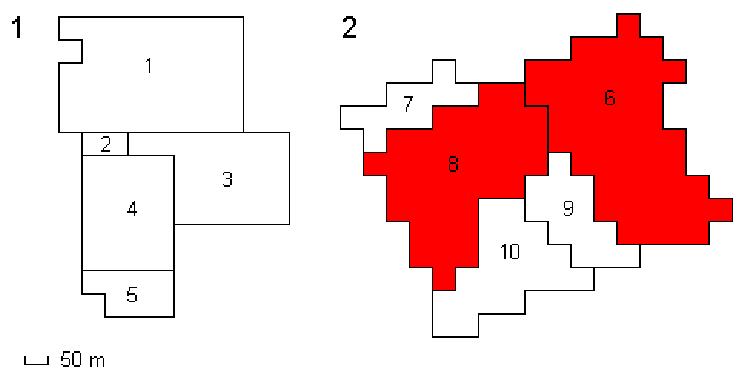

Figure 6-4 Field distribution for farm-2 consisting of ten fields in two tracts. Field numbers are as shown, shaded areas indicate HEL fields.

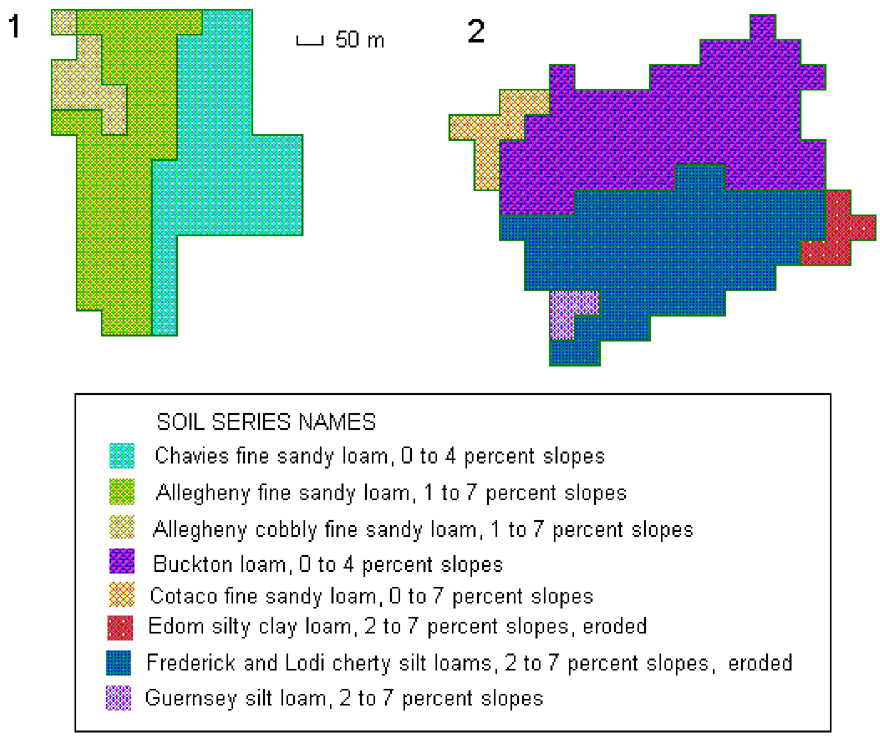

Figure 6-5 Soil type distribution for farm-2. Soil series are described in Table 6-2.

ANSWERS was used to simulate soil erosion for the benchmark plans on both farms. A single run on each tract was completed and the results pooled. Outputs from the simulation were the net change in soil (soil leaving minus soil deposited) within each basin and summed for each basin on the farm. Cell size was set to 50 m; this value has been suggested by Bouraoui (1994).

The data required by the planner and the ANSWERS model were taken from diverse sources. Much of the biophysical data describing the farm fields were extracted from GIS databases using the FARMSCALE system (Wolfe et al. 1995). Soil data were retrieved from the SOILS5 database (Goran, 1983) and the SCS (1982). Yields for crops were calculated for each farm field based on yield estimates for Virginia soil series obtained from the VALUES system (Simpson et al. 1993). Field yield was calculated as the average yield for all the soil types in the field, weighed by the proportion of the field containing each soil. Environmental risk values for fields were computed according to Buick et al. (1992) and are listed in Table A-2 of Appendix A. The GLEAMS manual provided the crop parameter data required by ANSWERS (Knisel et al. 1992).

Crop budget information was obtained from the Virginia Cooperative Extension (VCE, 1997). Meteorological data to run the simulations was generated by CLIGEN (Nicks et al. 1995) using data from the Monterey, Virginia, weather station.

Planning Configuration and Experimental Design

Two experiments were conducted to apply and evaluate the reactive approaches for planning, and a third experiment was conducted to evaluate the simultaneous control of corn rootworms and soil erosion. The first experiment consisted of applying the classic reactive approach, using a sequence of runs or iterations, based on the procedure shown in Figure 6-1A. For each run, the planner ran until 10 plans were found (except as noted) or the planner depleted the allotted time without finding any plan. If plans were found, the best plan was used as input to a simulation run of the ANSWERS model, with all other inputs identical to those used in the corresponding benchmark simulation. The soil losses were then compared with the benchmark values. If the plan generated soil losses higher than the benchmark, field-specific constraints or crop-specific constraints were modified. The iterations finished after the plan or plans found yielded soil losses lower than the benchmark. In all runs, total profit for alternative plans was constrained to be equal to or higher than the benchmark plan. The second experiment was similar to the first but followed the procedure shown in Figure 6-1B. Plans to repair were selected from plans generated from the first experiment. Portions of each plan were designated to be repaired; then the planner was run on those subsets of the farm. As for the first experiment, the iterations were stopped when the alternative plans had soil losses lower than the benchmark. Target crops were corn silage, alfalfa, pasture, and rye silage. The acceptable annual target acreage ranges were those listed in Table 6-1 for both experiments. The percent of benchmark profit for the partial plans to repair were adjusted to satisfy the farm-level goals. This is because the fields excluded containing the satisfactory part of the plan were already contributing to the profit. For example, if the satisfactory partial plan contributed with 30 percent of the total benchmark profit, then the portion of the plan to repair should contribute with 70 percent or more.

The third experiment tested how plans generated to control corn rootworms (see Chapter 5) performed in controlling soil erosion. Plans were generated for farm-2 using only the field constraint to select low CRW risk rotations; the crop constraint was not used because it was unsatisfactory. The best plan obtained was simulated using ANSWERS. The same weather data generated by CLIGEN was used across all runs for each experiment to make valid comparisons. The use of common random numbers as a variance-reduction technique has been suggested to compare different runs (Law and Kelton, 1991). In all cases, soil loss was measured as the total weight of soil leaving the farm, i.e., transported out the outlet of each of the basins that composed the flow network. Another variable analyzed was the ratio (R) of soil loss by T, which is the rate of soil formation (SCS, 1982). Values of R higher than one indicates that soil loss occur at rates higher than the rate of soil formation and thus will affect productivity in the long term. The T value was computed as a weighed average based on the soil composition of the farms. Two soil constraints were included in the planner to reduce soil loss, a field-level USLE-C constraint and a target crop-level HEL percent constraint. The USLE-C constraint is applied by the planner to select low erosion risk rotations by fields (Buick et al. 1992). The HEL percent constraint is applied by the planner to select field combinations with high proportion of HEL fields by low erosion target crops (like pasture or alfalfa). The HEL percent constraint is described in Chapter 3 in the section about the planning engine.

Results and Discussion

Classic Reactive Planning

The soil losses for the benchmark plan of farm-1 show an uneven distribution (Figure 6-6). The highest losses occurred in tract three, followed by tract one. The total soil loss was 1671.7 Mg or 27.2 Mg/ ha. In addition, the farm-level value of R was 0.54, indicating that soil erosion was lower than the rate of soil formation. However, looking at tract-level, soil losses were considered unacceptable in tract three (R= 1.3) and acceptable in tracts one and two (Figure 6-6). The approach to control soil erosion was thus aimed to try to reduce soil losses for tract three while maintaining the soil losses of the other tracts at similar levels.

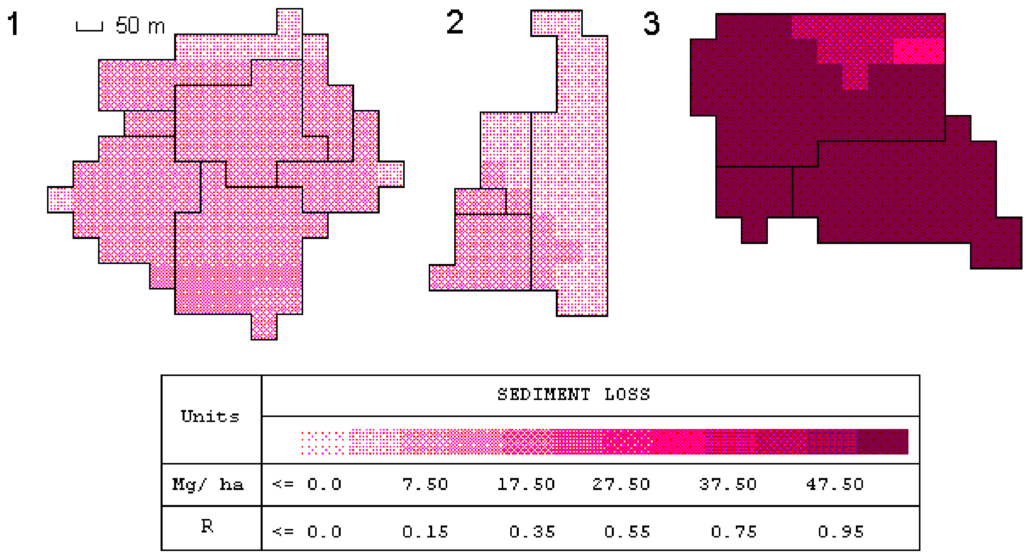

Figure 6-6 Soil loss (Mg/ ha) per basin for the benchmark plan of farm-1. R is the ratio of soil loss by the rate of soil formation.

For farm-1, the reactive planner required four iterations to obtain soil losses lower than the benchmark (Table 6-3, Figure 6-7). The third iteration plan showed essentially the same level of soil loss, an increase of 0.4 percent. In the fourth iteration, a large decrease in soil loss in tract three more than offset slight increases in the remaining tracts (Figure 6-8). In addition, in all iterations, the estimated plan net profit was equal to or higher than the benchmark (Figure 6-9). Also, because tightening the HEL constraint reduced field combinations, the planning time was reduced (Table 6-3). The farm-level value of R changed from 0.54 to 0.37, that is, although the benchmark plan had total soil losses lower than the rate of soil formation, it was unacceptable in tract three. The final plan reduced soil losses in all three tracts to be less than the soil formation rate (Figure 6-8).

Table 6-3 Iteration number, planning time PT (h), percent reduction in soil loss, result and action for the next iteration. Results for farm-1. The first iteration was run using the planner's default constraints (see Chapter 3).

Iteration	*PT*	*Soil loss change (%)*[1]	*Result/ action for next iteration*
1	2.5	9.6 %	Soil loss increased. HEL constraint enabled to assign at least 50 percent of HEL fields to pasture
2	12.0	4.3%	Soil loss increased. Manually enabled HEL constraint for fields 1, 5, and 6, essentially assigning them to pasture
3	5.5	0.4%	Soil loss increased. HEL constraint enabled to assign all HEL fields to pasture
4	0.3	-31.9%	Soil loss reduction. No further planning

[1] negative values indicate a percent reduction, positive values indicate an increase, as compared to the benchmark

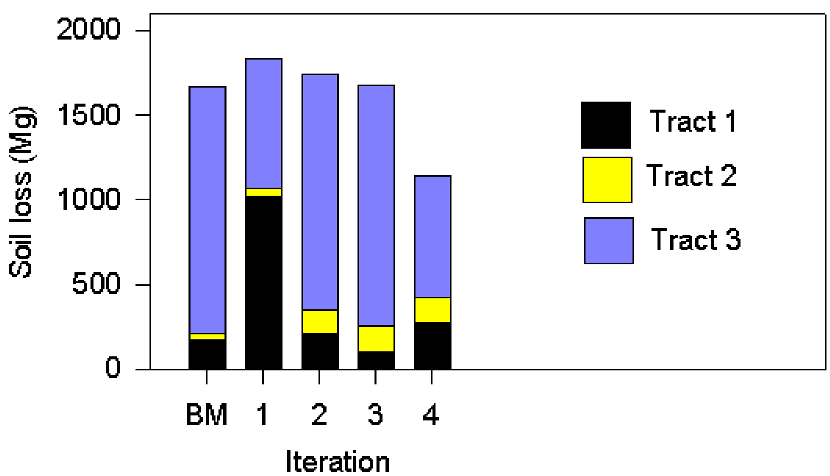

Figure 6-7 Soil losses (Mg) for different iteration plans using a reactive planning approach. BM benchmark plan. Results for farm-1.

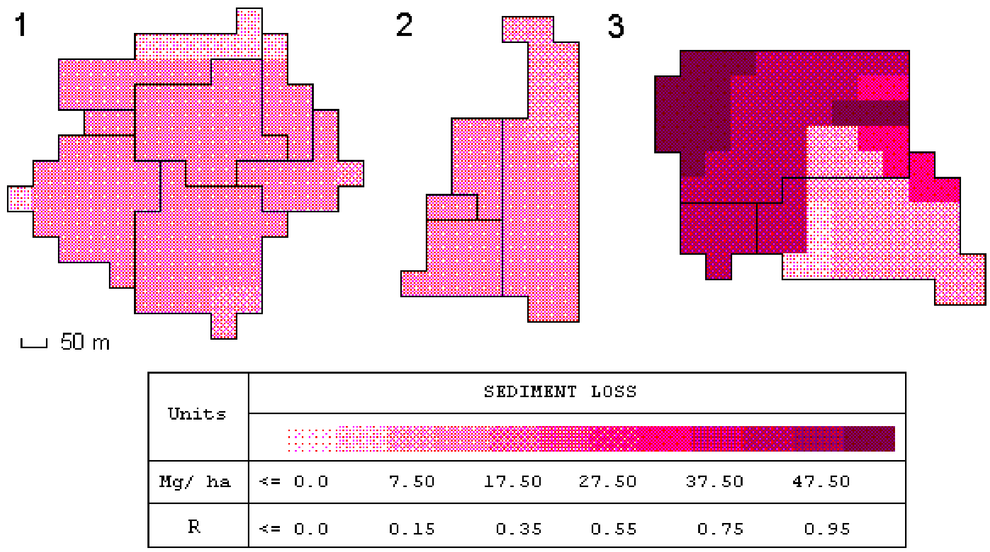

Figure 6-8 Soil loss (Mg/ ha) per basin for a fourth iteration plan. HEL constraint enabled to assign all HEL fields to pasture. Results for farm-1. R is the ratio of soil loss by the rate of soil formation.

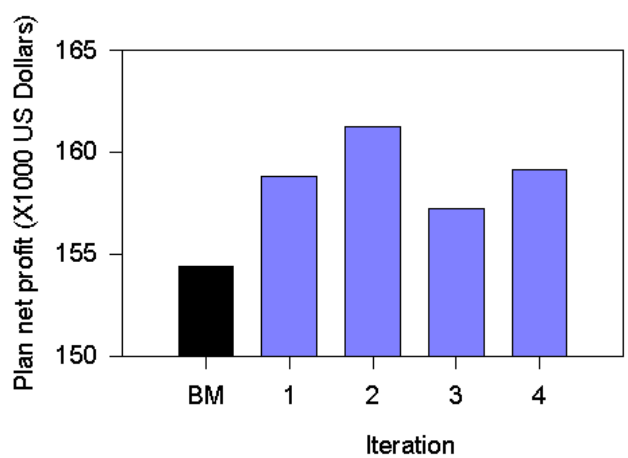

Figure 6-9 Plan net profit for the iteration plans (1-4) compared with the benchmark (BM) plan. Results for farm-1.

For farm-1, until all HEL fields were constrained to be pastures, plans failed to reduce erosion significantly (Table 6-3). How the planner selects rotations and potential plans in the final tree-search explains why the planner did not succeed during the first iterations in farm-1. As with any

CSP solver, if the planner has more than one valid value in the domain of any variable (field or target crop), more than one solution may be possible (Kumar, 1992). The rotation with the lowest USLE-C may be eliminated based on another criterion, or a plan using that rotation may be removed from the best ten plans due to returning lower profits. The fact that partially enabling the HEL constraint did not solve the problem indicates that the (ANSWERS) simulation's estimates for soil loss with some rotations were higher than the USLE estimates used by the planner. Only when the planner was forced to put those fields in pasture did it resolve the problem.

The benchmark plan for the second farm produced less soil erosion in tract one than in tract two (Figure 6-10). The total soil loss was 1648.7 Mg or 29.7 Mg/ ha, similar to farm-1. The farm-level value of R was 0.61, slightly higher than the benchmark for farm-1 but still acceptable. The R value for both tracts was below one, indicating an even distribution of soil losses, as compared to the farm-1. Based on the experience with farm-1, the initial run was done with the HEL constraint set to compel the planner to assign all the HEL fields to pasture.

Farm-2 required five iterations to obtain a reduction in soil loss (Table 6-4, Figure 6-1). Note that soil erosion was reduced in tract one but increased slightly in tract two (compare Figures 6-10 and 6-12). Also note that achievement of this reduction required more restrictive constraints that had been needed for farm-1, though only for one field. The planning time was about 10 minutes in the two runs to find plans, lower than for the first farm and probably due to a smaller number of fields (10 compared to 14 of farm-1). The farm-level R value was reduced from 0.61 to 0.53, thus achieving a better protection of soil across all the fields in the farm.

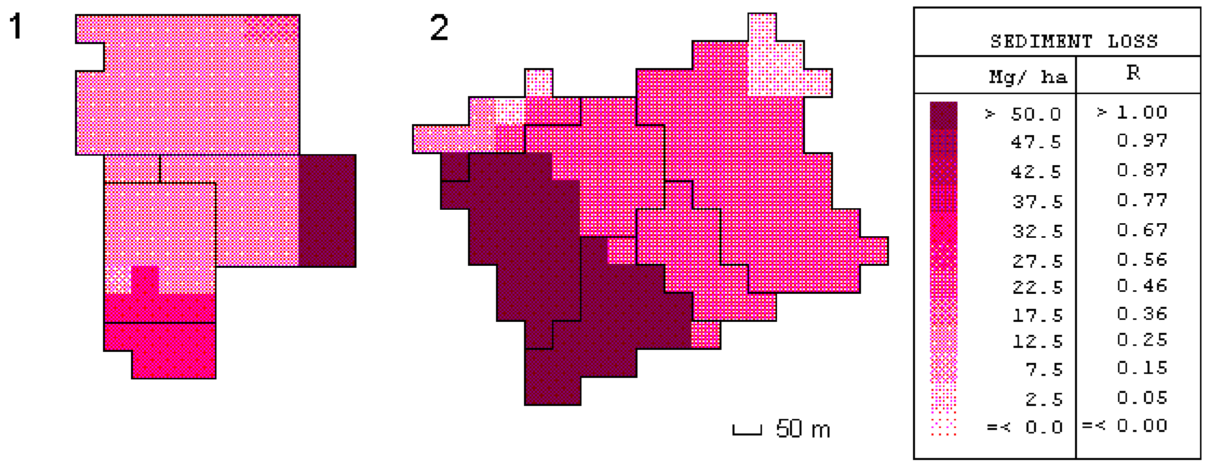

Figure 6-10 Soil loss (Mg/ ha) per basin for the benchmark plan of farm-2. R is the ratio of soil loss by the rate of soil formation.

Table 6-4 Iteration number, planning time PT (h), percent reduction in soil loss, result and action for the next iteration. Results for farm-2. The first iteration was run with the HEL constraint activated.

Iteration	*PT*	*Soil loss change (%)*	*Result/ action for next iteration*
1	0.16	-0.4%	Soil loss reduction. Tighten USLE-C constraint for all fields
2	N/A	N/A	No plans. Domains were emptied in arc consistency stage/ Tighten fields one to five, relax fields six to ten
3	N/A	N/A	No plans. Domains were emptied in arc consistency stage/ Tighten fields six to ten, relax fields one to five
4	N/A	N/A	No plans. Planner exhausted the search space/ Tighten field one only, relax rest of fields
5	0.16	-13.6%	Soil loss reduction. No further planning

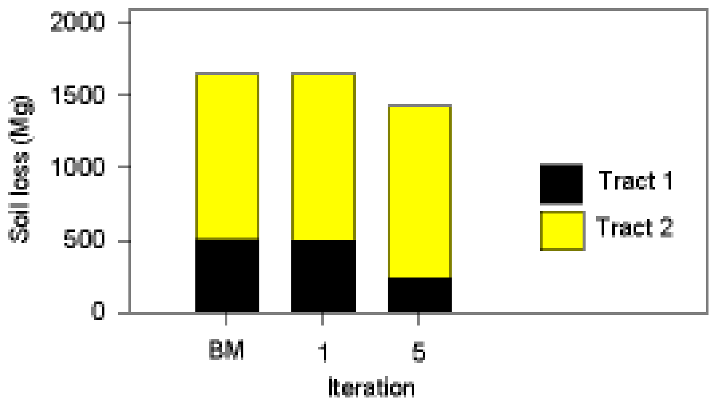

Figure 6-11 Soil losses for the iteration plans of farm-2. BM benchmark plan. Other iterations failed to produce plans.

With farm-2, the basic strategy was to find an appropriate level of enforcement of the USLE-C constraint through repeated tightening and relaxing, intermediate iterations failed because tightening the USLE-C field constraint across all the fields or groups of fields eliminated all the solutions. The USLE-C constraint (Buick et al. 1992) removes rotations that have USLE-C values incompatible with a maximum USLE-C value calculated for a given field based on an acceptable level of erosion.

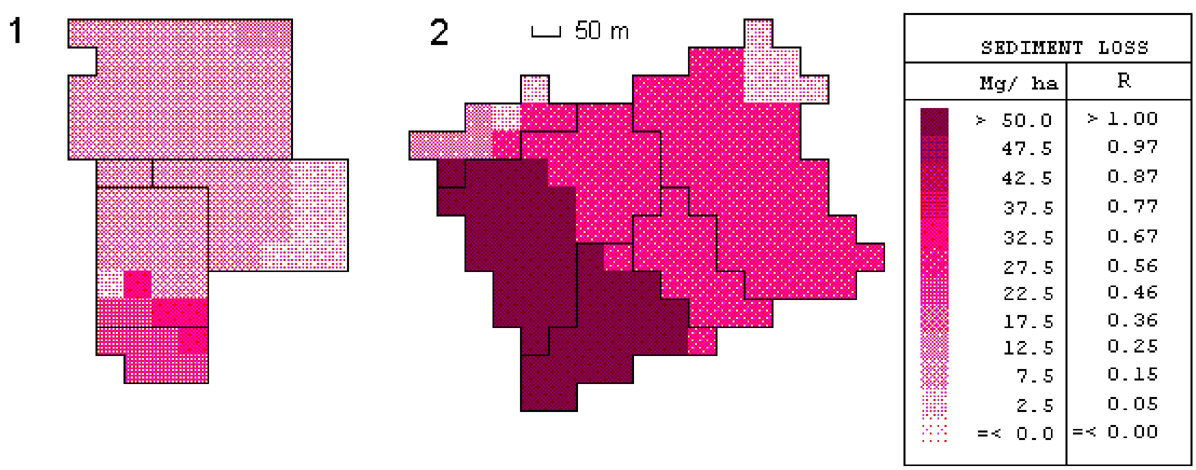

Figure 6-12 Soil loss (Mg/ ha) per basin for a plan from the final iteration of farm-2. R is the ratio of soil loss by the rate of soil formation.

For both farms, the search for better plans was successful in terms of reducing the farm soil losses. Because the target crops and the target acreages were within the same ranges for the different iterations, the difference occurred in the spatial distribution of the crops. For example, in the benchmark plan for farm-1, fields one, five, and six, (HEL fields) were assigned to pasture, but another HEL field (field 13) was not. In the alternative plan, all the HEL fields (one, five, six, and 13) were assigned to pasture by the planner, thus effectively reducing soil losses (Figure 6-13).

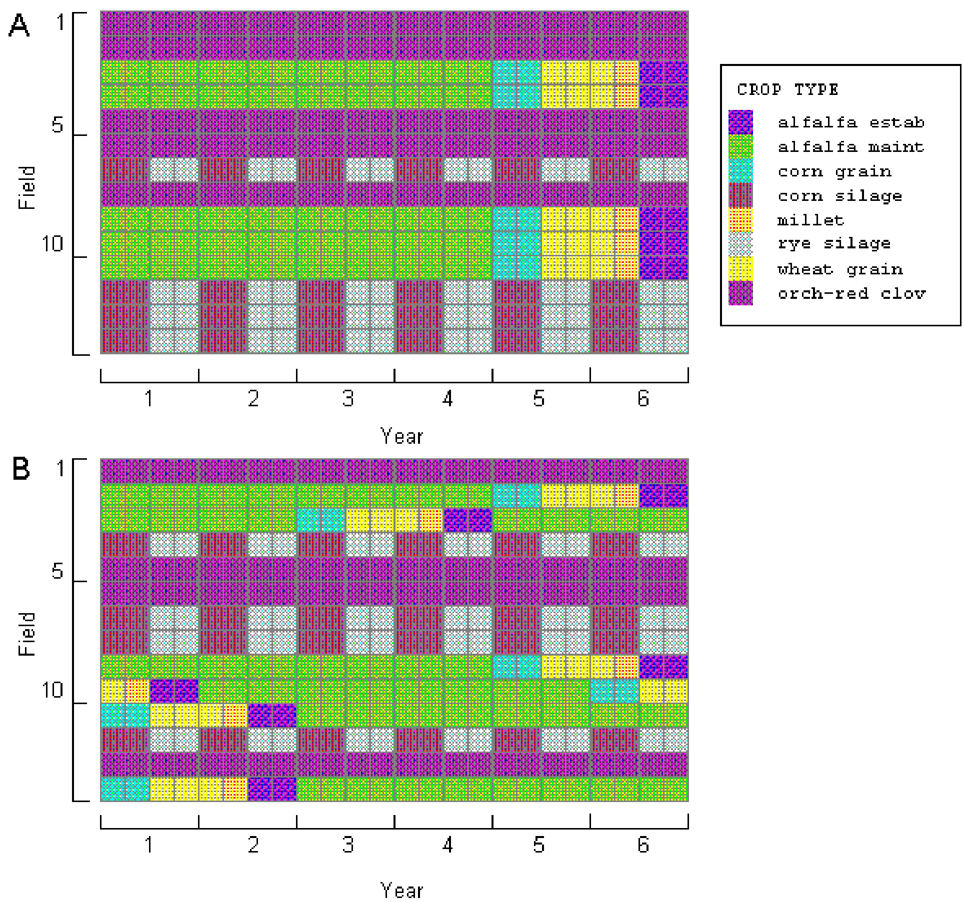

Figure 6-13 Crop distribution across fields and years. (A) Benchmark plan. (B) First plan of last iteration. Farm-1 data.

In both cases, reducing soil erosion was achieved by selectively enforcing one of two soil erosion constraints. The strategy was similar to one used to solve a job-shop scheduling problem using reactive scheduling, as described by Fox (1990). In Fox's system, solutions were obtained by selectively relaxing constraints when the automatic constraint propagation procedure could not find solutions. In this case, the feedback from ANSWERS identified the areas with high erosion rates and guided the selection of constraints to tighten or to relax. The default parameter values for the constraints did not always satisfy the goal of reducing soil erosion, but by individually manipulating the constraints, it was possible to find solutions that best satisfied all the constraints.

Repair-based Reactive Planning

As the basis for the repair-based approach, two plans were selected from the first iteration run of the reactive planning experiment for farm-1. Farm-2 was not used in this experiment. In both cases, the candidate plans had decreased soil loss in the third tract relative to the benchmark plan. The repair, therefore, focused on tracts one and two only (Figure 6-14). For the first plan to be repaired, preliminary attempts failed to find plans, by either failing final economic constraint or the target acreage (Table 6-5). Therefore, erosion constraints were tightened and target acreage for alfalfa relaxed until the planner returned two acceptable plans. In the case of the second candidate plan, and based on the results of the first repaired plan, the search started by relaxing the target acreage for alfalfa and tightening the HEL percent constraint, obtaining plans in the first iteration in 27 minutes. For both cases, overall soil erosion was lowered significantly in tract one while erosion in tract two increased (Figures 6-14 and 6-15). In both cases, overall soil loss was reduced by over 20 percent compared to the benchmark and over 25 percent compared to the original plans (Figure 6-14 and Table 6-6). The farm-level value of R indicated a reduction from 0.60 (original plan) to 0.42 (repaired plan) for the first selected plan and from 0.56 (original plan) to 0.40 (repaired plan) for the second selected plan. Compare these values with the benchmark value of R= 0.54.

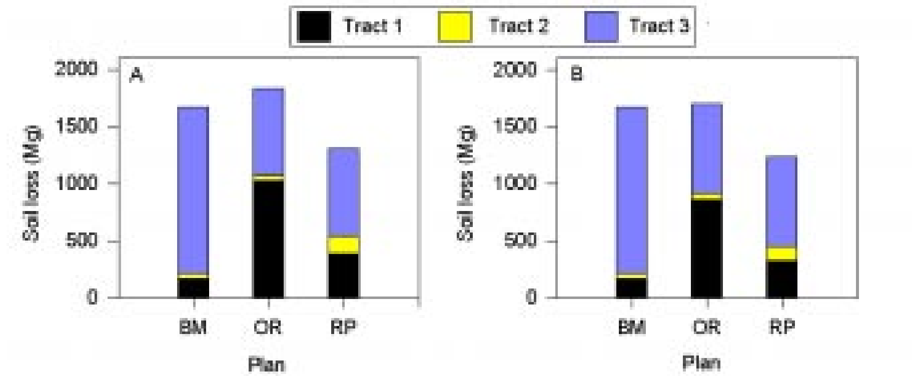

Figure 6-14 Soil losses for candidate plans to repair, taken from first iteration of the classic reactive approach experiment. (A) First repaired plan. (B) Second repaired plan. (BM) benchmark plan, (OR) original plan, (RP) repaired plan.

Table 6-5 Iteration number, planning time PT (h), percent reduction in soil loss, result and action for the next iteration. Results for farm-1 using a repair-based approach. The first iteration was run using the HEL percent constraint to assign at least 50 percent of HEL fields to pasture.

Iteration	*PT*	*Soil loss change (%)*	*Result/ Action for next iteration*
1	0.5	N/A	Plans did not satisfy final economic constraint/ Relax target acreage range for alfalfa. Assign at least 50 percent of HEL fields to pasture
2	6.5	N/A	No plans found in allotted time/ Relax target acreage range for alfalfa. Assign all HEL fields to pasture
3	2.6	-21.9	Found two plans only, soil loss reduction/ No further planning

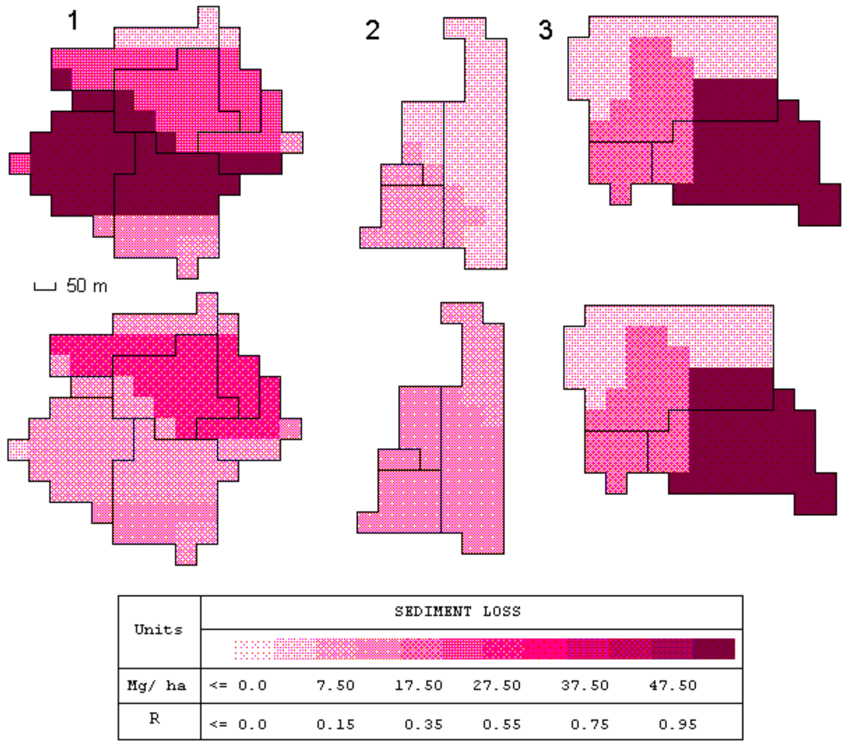

Figure 6-15 Soil loss (Mg/ ha) per basin for the first repaired plan. Tracts on the top show soil losses for the original plan. Tracts on the bottom show soil losses for the repaired plan. Soil loss reduction occurred in tract one and increased in tract two. The soil loss for tract three in the repaired plan is the same as in the original plan because that part of the plan is the same. R is the ratio of soil loss by the rate of soil formation.

Table 6-6 Estimated soil loss (Mg) using a repair-based approach. Results for farm-1. Plans selected from first iteration of the classic reactive approach. Losses for tract three between the original and the repaired plan are the same because that tract contained the same partial plan.

Tract	***Plan 1***		***Plan 2***	
	Original	*Repaired*	*Original*	*Repaired*
1	1022.8	390.2	864.9	325.0
2	47.8	151.7	52.8	130.7
3	762.7	762.7	781.5	781.5
Total	1833.3	1304.6	1699.2	1237.2

Although the repair approach was successful, the final plans included a trade-off: in both cases, the planner had to expand the acceptable target acreage for alfalfa, one of the target crops (Figure 6-16). If alfalfa were a crop for which the farmer had to comply with a minimum acreage to participate in a federal program, then the plans would have to be rejected. On the other hand, if the crop is planted only to provide food to feed livestock, other forage or purchased food could probably be used to compensate for such reduction. Over six years, the farmer has time to plan for loans to buy alternative food or keep excess hay from previous years. Ultimately, a more thorough economic analysis, such as cash flow (Kay, 1986) would probably shed more information on the feasibility of accepting these alternative plans.

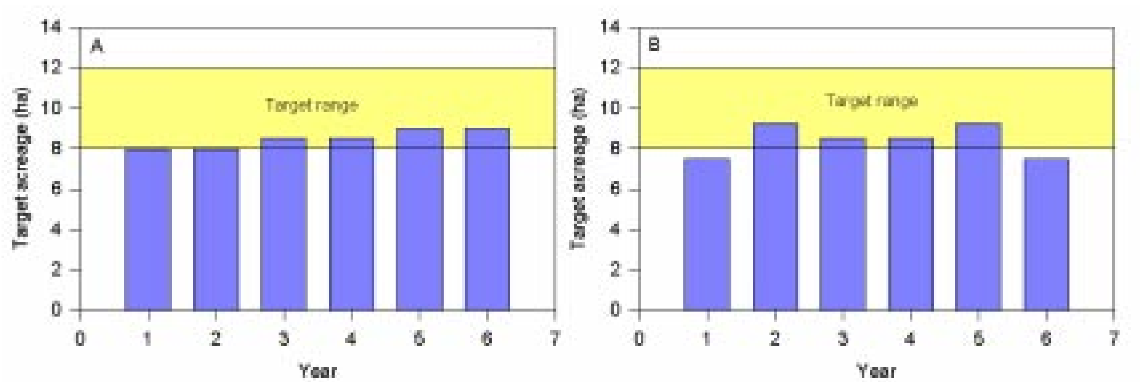

Figure 6-16 Expected target range acreage and actual values for alfalfa. (A) Repaired plan 1. (B) Repaired plan 2.

Combining Corn Rootworm Control and NPS Control

The true power of a whole-farm planning system is in its ability to plan simultaneously for several resources. As a simple verification, a single plan (the best) from Chapter 5 related to corn rootworm (CRW) control was tested here for its effect on soil erosion. Ideally, plans generated to control CRW also reduce soil erosion to acceptable levels. To test this, a single CRW plan was used as input to ANSWERS and resulting soil loss was compared to the benchmark. As shown in Figure 6-17, the CRW plan reduced soil loss by 10.8 percent, as compared to the benchmark, though not as much as was achieved by the reactive planner (13.6 percent reduction, Table 6-4). Soil loss was reduced from R= 0.61 to R= 0.54. The result highlights the trade-off when planning to address a specific problem. In this last experiment, the objective was to reduce CRW damage;

nothing was done to specifically improve its erosion mitigation. Consequently, the erosion results were inferior to those obtained for the farm when specifically targeting soil erosion. Nevertheless, the result tentatively suggests that it is possible to simultaneously reduce multiple environmental problems without jeopardizing the production and economic goals of farmers.

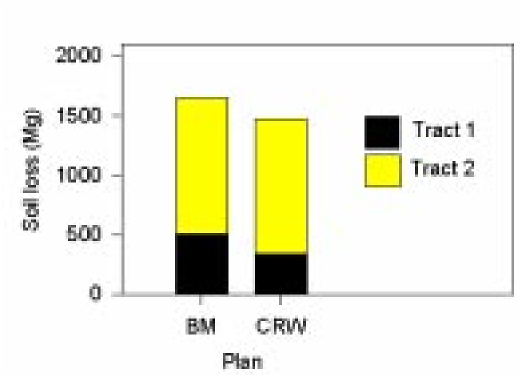

Figure 6-17 Soil losses estimates for the benchmark plan (BM) and a corn rootworm (CRW) plan. Results for Farm-2. The CRW plan reduced soil erosion by 10.8 percent.

Summary

A classic reactive approach was explored to reduce NPS pollution. The objective was to reduce sediment as a representative NPS pollutant. In this approach, the planner obtained a candidate farm plan and simulated its effects on soil erosion using ANSWERS, an NPS distributed-parameter model. The plan was evaluated by comparing its simulated soil losses against values simulated using a benchmark (control) plan. If the candidate plan's soil loss was lower than soil loss from the benchmark, it was accepted. Otherwise, it was discarded.

To find improved plans, constraint parameters were modified to focus more severely on sensitive areas, new plans were generated and their effects simulated and compared. The reactive approach was tested on two artificially composed farms in Rockingham Co., Virginia, to represent small and medium dairy farms. Overall, the reactive planner reduced soil erosion by 32 percent, compared to a benchmark plan. In the most highly erodible tract, the planner reduced losses from 1.3T to 0.6T, a net gain of 5.3 Mg/ ha•yr of topsoil. Values of T higher than one indicate that soil erosion affects crop productivity in the long term (SCS, 1982).

A second approach, repairing partially acceptable plans, reduced soil losses equivalent to 0.4T, but the acreage for a target crop had to be relaxed. To determine the compatibility between corn rootworm control and soil erosion control, a third experiment evaluated a plan generated specifically to control corn rootworms. The result showed a reduction in soil erosion from 0.6T to 0.5T, suggesting that IPM and soil conservation goals can be achieved simultaneously. An important observation from the reactive planning experiments was that manipulating individual soil erosion constraints was necessary to allow the planner to obtain acceptable plans.

Chapter 7 Summary and Conclusions

The aim of this study was to design, implement, and evaluate a whole-farm planning decision support system able to deal with the complexity of farm planning. That is, simultaneously consider most of the factors that affect farm planning, such as limited resources (land, crops available, and machinery), and goals and preferences of the farmer (income, target acreages for selected (target) crops, control of soil erosion, control of nitrate leaching, and control of insect pests). A second objective was to evaluate some sorting heuristics included into the planner to find high-profit plans quickly. The third objective was to evaluate the inclusion of preventive IPM heuristics to reduce damage caused by some insect pests. The final objective was to use a nonpoint source (NPS) pollution model to evaluate proposed plans in terms of soil erosion control and direct the search to obtain better plans in terms of soil erosion while maintaining profit.

The resultant system includes five modules: a planning module, a nonpoint source (NPS) pollution module, a farm-level management module, a database module and utility programs. The planning module is based on the CROPS planning engine (Stone, 1995). It uses a constraint satisfaction approach (CSP) to find farm plans that satisfy goals and preferences of the farmer, for example, income and target acreage for selected crops. The planner also restricts the negative impacts of farm activities (for example, soil erosion, and nitrate and pesticide leaching). The NPS module includes a distributed-parameter model (ANSWERS, Bouraoui, 1994) and a weather generator (CLIGEN, Nicks et al. 1995). The farm-level management module manages the data required by the system, displays farm data, creates the input files for the NPS module and performs economic analysis (crop budgets). The database module includes databases and a database engine. Its tasks are data storage and processing of queries. The utility programs module allows the user to communicate between the programs across different machines.

Because an exhaustive search of all the possible combinations of potential solutions is impractical, a planner should be able to find near-optimal plans in limited time. Sorting heuristics were added to the planner to achieve this goal. The experiments in Chapter 4 demonstrated that sorting rotations effectively reduced the time to obtain plans, as compared to an unsorted (random) strategy. In addition, the frequency of occurrence of the target crop in the available rotations restricts the maximum potential acreage planted with that crop. Finally, the exclusion of fields not assigned to the target crops usually reduced the planning time and increased the profit.

Preventive IPM was included in the planning system to control corn rootworms CRWs (*Diabrotica virgifera virgifera* and *D. barberi*). CRWs were selected because their life cycle can be disrupted and damage can be avoided by planting a non-host crop (a crop other than corn) the next year after planting corn, thus preventing the occurrence of the pests. Rules were developed to select rotations and field combinations with low-risk to CRW attack. The evaluation and application of these strategies demonstrated that selecting low-risk rotations effectively reduced CRW damage, reduced control costs, and increased profit.

The last objective was achieved by using the planner with the NPS model in a feedback control loop. Plans found by the planner were evaluated with ANSWERS and compared with values from a benchmark (control) plan. If the alternative plan failed, planning continued until alternative plans had soil losses lower than the control. The approach proved effective, reducing overall soil losses by 32 percent. In the most highly erodible tract, it reduced losses from 1.3T (a net loss of 10.4 Mg/ ha•yr of topsoil), to 0.6T (a net gain of 5.3 Mg/ ha•yr of topsoil). Soil losses higher than one T indicate that soil erosion affects crop productivity in the long term (SCS, 1982). In addition, a repair-based version approach was tested. In this case, an alternative plan could reduce soil losses in parts of the farm, so these fields and their rotations were retained and planning was done for the remaining fields. The search stopped when soil losses were lower than the benchmark as well. This approach reduced soil losses overall by 25 percent, and in the most highly erodible tract, from 1.3T (a net loss of 10.4 Mg/ ha•yr of topsoil) to 0.7T (a net gain of 4.9 Mg/ ha•yr of topsoil). However, relaxation of the target acreage for one target crop was necessary to obtain plans. Finally, the evaluation of a plan that controlled CRW showed reduced soil losses, as compared to a benchmark plan, thus suggesting the feasibility of simultaneously control soil erosion and insect pests.

The development, implementation, and application of this decision support system required organizing data from multiple sources, and capturing and applying knowledge from several domains, like simulation modeling, pest and farm management, and computer science. This approach is necessary to address the complexity of whole-farm planning. This study demonstrated that it is feasible to develop a whole-farm planning decision support system to deal with the complexity of farm planning, include sorting heuristics to obtain good plans quickly, include preventive IPM heuristics to avoid pest damage, and use simulation model's results to improve the planning process. Finally, the development and dissemination of systems like this one would help farmers or planners to better make decisions regarding farm management while protecting natural resources.

Chapter 8 Recommendations for Future Research

The whole-farm planning decision support system was designed to be easy to use, based on available components if possible, and potentially upgradeable. Although no formal tests were performed, I consider that all these objectives were met to a high degree. By using a rapid application development tool, the construction of the visual interface, planning engine, and database management was accomplished efficiently and in a modular fashion well suited to upgrades and modifications. The database subsystems (FLAME and database modules) greatly reduced the effort of storing and retrieving the required information. For example, instead of writing routines to obtain averages, using appropriate SQL calls, the database engine computed such values. In addition, the automation of storage and display of simulation results made it easy to analyze the output from ANSWERS. For example, all figures in Chapter 6 that show the spatial distribution of soil loss were generated by the system on a per tract basis and then grouped by farm. Integrating the ANSWERS and CLIGEN programs was a more challenging task because they run on different computer platforms. The FARMSCALE system (Wolfe et al. 1995), which includes versions of ANSWERS and CLIGEN, runs under Unix. CLIGEN was ported and recompiled under MS DOS. An attempt to do the same with ANSWERS was abandoned due to incompatibilities in the code. Better integration could be achieved by converting all the components to a single language like C++. In fact, the CROPS planner was coded in C++ (Stone, 1995).

ANSWERS could be improved by including a module to simulate pesticide dynamics as in models like EPIC, GLEAMS, and AGNPS (USEPA, 1997b). This would allow the whole-farm planner to adjust planning goals based on the fate of pesticides, in a way similar to how it can react to the transport of sediments. Further automation and simplification of input data for ANSWERS is needed, e.g., the whole input file could be generated in a single step from the databases with further programming. In addition, adding some of the functionality of FARMSCALE would help in the initial stages to retrieve biophysical data about the farm from GIS layers. GIS display capabilities like those in the FARMSCALE system would also improve data visualization. However, I would not recommend a full GIS program because of the large overhead, instead I would recommend using an OCX component like MapObjects (ESRI, 1998) that could be embedded within the system.

In the case of the CSP planner, the ability to change constraint parameters manually greatly facilitated planning to control both soil erosion and CRW. The CROPS planner has an automated constraint relaxation scheme (Stone, 1995) that does not allow the user to modify constraint properties. Testing CROPS against the current version would help to address strengths and weaknesses of the automatic relaxation approach versus having manual control of the constraints.

The planner can be improved by including a more complete economic analysis. At this time, the planner considers plan net profit based on crop budgets; however, other indices and methodologies can be included. That is, a full analysis should include medium and long-term assets, such as land, buildings, and machinery. Methods like cash flow budgeting, balance sheet, income statement, and total budget analysis could be used to evaluate a farm plan and obtain a better assessment to make a decision (Harsh et al. 1981; Kay, 1986). For example, to assess farm profitability, net farm income is a preferred index. Net farm income is computed from the income statement and measures the profitability of the business as return to unpaid family job, operator's labor, equity capital, and management (Boehjl and Eidman, 1984). These economic analysis tools would improve the decision making process by allowing the farmer to select and evaluate proposed plans, that could, for example, allow the expansion of crop enterprises.

The proposed CRW control model proved that it is possible to organize, synthesize, and include heuristics within a whole-farm planning system while maintaining profit and reducing pesticide applications. However, further work is needed to determine whether including prolonged diapause can improve the model. The current implementation does not take this into account. Prolonged diapause is becoming more common and may reduce the effectiveness of crop rotation as a preventive strategy (Levine et al. 1992ab). The model could be modified to account for prolonged diapause by assigning probabilities of CRW occurrence for 3rd-year corn (as it currently does for 2nd-year corn) and determining the effect on corn profit. Another improvement would be to incorporate spatial heuristics related to CRW movement. For example, cornfields adjacent to 2nd-year corn cornfields might have higher risk of CRW damage due to movement of CRW adults to contiguous cornfields. This could be incorporated in the planner through field combination objects. For example, each corn field combination could include a “patch” property representing the number of sets of contiguous fields in corn within the field combination. A patch constraint could select field combinations to reduce the risk of having contiguous cornfields or to minimize overall risk including a component based on the clumping of corn. The path constraint was designed and implemented (see Table 3-2) but it was not evaluated in terms of CRW control because of time limitations.

Regarding the control of soil erosion, the feedback control framework helped to direct the search to find better plans. In addition, the design and implementation of the HEL constraint helped by finding better field combinations. An improvement would be to perform probability-based risk assessment. The analyses performed in Chapter 6 were based on single simulation runs. By using multiple runs based on random weather scenarios, one could generate probability distributions of soil loss replicates to adjust probabilistic models. Models like the normal, lognormal, Erlang, or Weibull could then be fit to the simulated data (Law and Kelton, 1991; Sokal and Rohlf, 1995) and a probabilistic risk assessment generated, giving the likelihood of obtaining high losses during a given time frame.

Using the methods developed here with appropriate modifications, other NPS pollutants can be controlled during planning. For example, ANSWERS simulates the dynamics of the plant nutrients nitrogen and phosphorus. ANSWERS would simulate their fate in the soil and would estimate leaching and runoff. The planner would have to be able to modify application timing and rates based on constraints.

Finally, more research integrating insect pest management with other resource management (soils, nutrients, animal and plant health, etc.) strategies would be useful. Here, we saw that insect pest management and soil conservation can be combined successfully, but a formal analysis is required to determine the robustness of combining both control methods. In addition, updating and expanding the crop network to include alternative crops or new agronomic practices would be helpful. For example, alternative silage crops like sorghum silage can be used in a dairy system to reduce corn acreage and the need to plant continuous corn (Youngman et al. 1993). Until that crop is included in the planner, however, the system will not take advantage of the possibility.

References

Allee, L.L., and P.M. Davis. 1996. Effect of manure on maize tolerance to western corn rootworm (Coleoptera: Chrysomelidae). J. Econ. Entomol. 89: 1608-1620.

Ambler, S.W. 1998. Building object applications that work: Your step-by-step handbook for developing robust systems with object technology. SIGS Books. New York.

Beasley, D.B., L.F. Huggins, and E.J. Monke. 1980. ANSWERS: A model for watershed planning. Transactions of the ASAE. 23: 938-944.

Bhogaraju, P. 1996. A case-based reasoner for evaluating crop rotations in whole-farm planning. Master of Science Thesis. Virginia Polytechnic Institute and State University. Blacksburg, Virginia.

Boehjl, M., and V.R. Eidman. 1984. Farm management. Wiley. New York.

Boetel, M.A., D.D. Walgenbach, G.L. Hein, B.W. Fuller, and M.E. Gray. 1992. Oviposition site selection of the northern corn rootworm (Coleoptera: Chrysomelidae). J. Econ. Entomol. 85: 246-249.

Booch, G. 1994. Object-oriented analysis and design. The Benjamin/Cummings Publishing Company, Inc. Redwood City, California.

Borland International. 1995. User's guide. Borland Delphi for Windows 95 and Windows NT. Borland International. Scotts Valley, California.

Bouchard, D.C., M.K. Williams, and R.Y. Surampalli. 1992. Nitrate contamination of groundwater: sources and potential health effects. J. Am. Waterworks Assoc. 84: 85-90.

Bouraoui, F. 1994. Development of a continuous, physically-based, distributed parameter, nonpoint source model. Ph.D. Dissertation. Virginia Polytechnic Institute and State University. Blacksburg, Virginia.

Box, G.E.P., and N.R. Draper. 1987. Empirical model-building and response surfaces. Wiley. New York.

Bridge, G. 1993. Is whole-farm planning the answer? J. Soil and Water Cons. 48: 95-98.

Brust, G.E., and L.R. King. 1994. Effects of crop rotation and reduced chemical inputs on pests and predators in maize ecosystems. Agric. Ecosyst. and Environ. 48: 77-89.

Buick, R.D., N.D. Stone, R.K. Scheckler, and J.W. Roach. 1992. CROPS: A whole-farm crop rotation planning system to implement sustainable agriculture. AI Applications. 6: 29-50.

Buick, R.D., R.R. Youngman, and N.D. Stone. 1993. VICE-Corn: An expert system for insect pest management in corn. Virginia Agricultural Experiment Station. Information Series 93-2. Virginia Polytechnic Institute and State University. Blacksburg, Virginia.

Bullock, D.G. 1992. Crop rotation. Crit. Rev. Plant Sci.11: 309-326.

Chiang, H.C. 1973. Bionomics of the northern and western corn rootworms. Ann. Rev. Entomol. 18: 47-72.

Crutchfield, S., L. Hansen, and M. Ribaudo. 1993. Agricultural and water-quality conflicts. Economic dimensions of the problem. USDA-ERS. Agric. Inf. Bull. No. 676. Washington, D.C.

Davenport, T.E. 1994. EPA’s perspective- you need to protect water quality. J. Soil and Water

Cons. 49: 14-15.

Dent, D. 1991. Insect pest management. CAB International. Wallingsford. United Kingdom.

Dewey, B.R. 1988. Computer graphics for engineers. Harper and Row. New York.

Dillaha, T. 1990. Role of BMPs in restoring the health of the Chesapeake Bay: Assessment of effectiveness. In: Perspectives on the Chesapeake Bay, 1990: Advances in Estuarine Sciences. Chesapeake Bay Program. United States Environmental Protection Agency. Washington D.C.

Dillaha, T. 1996. Nonpoint source pollution. Class notes. Virginia Polytechnic Institute and State University. Blacksburg, Virginia.

Doyle, C.J. 1990. Applications of systems theory to farm planning and control: Modelling resource allocation. In: Systems theory applied to agriculture and the food chain. Jones, J.G.W., and P.R. Street (eds.). Elsevier. London.

Draper, N.R., and H. Smith. 1981. Applied regression analysis. Wiley. New York.

Elliot, N.C., J.J. Jackson, and R.D. Gustin. 1990. Predicting western corn rootworm beetle (Coleoptera: Chysomelidae) emergence from the soil using soil or air temperature. Can. Ent. 122: 1079-1091.

ESRI. 1997. Understanding GIS. The ARC/INFO method. GeoInformation International. Cambridge, United Kingdom.

ESRI. 1998. MapObjects. Environmental Systems Research Institute, Inc. Redlands, California.

Feather, P.M., and J. Cooper. 1995. Voluntary incentives for reducing agricultural nonpoint source water pollution. USDA-ARS. Agric. Inf. Bull. 716. Washington D.C.

Flanagan, D.C., J.C. Ascough II, A.D. Nicks, M.A. Nearing, and J.M. Laflen. 1995. Overview of the WEPP erosion prediction model. In: USDA- Water Erosion Prediction Project: Hillslope profile and watershed model documentation. NSERL Report 10. USDA-ARS-NSERL. West Lafayette, Indiana.

Foster, R.E., J.J. Tollefson, J.P. Nyrop, and G.L. Hein. 1986. Value of adult corn rootworm (Coleoptera: Chrysomelidae) population estimates in pest management decision making. J. Econ. Entomol. 79: 303-310.

Fox, M.S. 1990. Constraint-guided scheduling- A short history of research at CMU. Computers in Industry. 14: 79-88.

Frere, M.H. 1982. Modeling the quality of water from agricultural land. In: Hydrologic modeling of small watersheds. Haan, C.T., H.P. Johnson, and D.L. Brakensiek (eds.). ASAE Monograph No. 5. American Society of Agricultural Engineers. St. Joseph, Michigan.

Georgeff, M.P. 1987. Planning. Ann. Rev. Comp. Sci. 2: 359-400.

Godfrey, L.D., L.J. Meinke, R.J. Wright, and G.L. Hein. 1995. Environmental and edaphic effects on western corn rootworm (Coleoptera: Chrysomelidae) overwintering egg survival. J. Econ. Entomol. 88: 1445-1454.

Goran, W.D. 1983. An interactive soils information system users manual. U.S. Construction Engineering Research Laboratory Technical Report N-163.

Gray, M.E., and J.J. Tollefson. 1987. Influence of tillage and western corn rootworm (Coleoptera: Chrysomelidae) egg populations on larval populations and root damage. J. Econ. Entomol. 80: 911-915.

Gray, M.E., and W.H. Luckman. 1994. Integrating the cropping system for corn insect pest management. In: Introduction to insect pest management. Metcalf, R.L., and W.H. Luckman (eds.). Wiley. New York.

Gruber, M. 1993. SQL instant reference. Sybex. Alameda, California.

Haralick, R.M., and G.L. Elliot. 1980. Increasing tree search efficiency for constraint satisfaction problems. Artificial Intelligence. 14: 263-313.

Harsh, S.B., L.J. Connor, and G.D. Schwab. 1981. Managing the farm business. Prentice-Hall. Englewood Cliffs, New Jersey.

Heinemann, P.H., D.D. Calvin, J. Ayers, J.M. Carson, W.S. Curran, V. Eby, R.L. Hartzler, J.G. Kelley, J. McClure, G. Roth, and J. Tollefson. 1992. MAIZE: a decision support system for management of field corn. Appl. Eng. Agric. 8: 407-413.

Holland, J.H. 1992. Adaptation in natural and artificial systems. The MIT Press. Cambridge, Massachusetts.

Hopgood, A.A. 1993. Knowledge-based systems for engineers and scientists. CRC Press. Boca Raton, Florida.

Houstis, E., E. Gallopolous, R. Bramley, and J.R. Rice. 1997. Problem solving environments for computational science. IEEE Computational Science and Engineering. 4: 18-21.

Jackson, J.J., and N.C. Elliot. 1988. Temperature-dependent development of immature stages of the western corn rootworm, *Diabrotica virgifera virgifera* (Coleoptera: Chrysomelidae). Environ. Entomol. 17: 166-171.

Jordan, T.E., D.L. Correll, and D.E. Weller. 1997. Effects of agriculture discharges of nutrients from coastal plain watersheds of Chesapeake Bay. J. Environ. Qual. 26: 836-848.

Kabanza, F., M. Barbeau, and R. St-Denis. 1997. Planning control rules for reactive agents. Artificial Intelligence. 95: 67-113.

Kahler, A.L., A.E. Olnes, G.R. Sutter, C.D. Dybing, and O.J. Devine. 1985. Root damage by western corn rootworm and nutrient content in maize. Agron. J. 77: 769-774.

Kay, R.D. 1986. Farm management. McGraw-Hill. New York.

Knight, A.L., and G.W. Norton. 1989. Economics of agricultural pesticide resistance in arthropods. Ann Rev. Entomol. 34: 293-313.

Knisel, W.G. (ed.) 1980. CREAMS: A field-scale model for chemicals, runoff, and erosion from agricultural management systems. U.S. Department of Agriculture, Science, and Education Administration, Conservation report No. 26.

Knisel, W.G., F.M. Davis, R.A. Leonard, and A.D. Nicks. 1992. GLEAMS version 2.10 Part III: User manual. Coastal Plain Experiment Station. University of Georgia. Tifton, Georgia.

Kuhar, T., R.R. Youngman, and C.A. Laub. 1997. Risk of western corn rootworm (Coleoptera: Chrysomelidae) damage to continuous corn in Virginia. J. Entomol. Sci. 32: 281-289.

Kumar, V. 1992. Algorithms for constraint satisfaction problems: A survey. AI Magazine. 13: 32-44.

Laflen, J.M. 1997. WEPP- Erosion prediction for the next millennium. Agric. Res. 45: 2.

Law, A.M., and W.D. Kelton. 1991. Simulation modeling and analysis. McGraw-Hill. New

York.

Lazarus, S.S., and G. White. 1984. Economic impact of introducing rotations on Long Island potato farms. NorthEast. J. Agric. and Res. Econ. 13: 221-228.

Lazarus, W.F., and E.R. Swanson. 1983. Insecticide use and crop rotation under risk: Rootworm control in corn. Am. J. Agric. Econ. 65: 738-747.

Lee, J.G., J.R. Ellis, and R.D. Lacewell. 1988. Crop selection and implications for profit and wind erosion in a semi-arid environment. Agric. Econ. 2: 319-334.

Leonard, R.A., W.G. Knisel, and D.A. Still. 1987. GLEAMS: Groundwater loading effects of agricultural management systems. Transactions of the ASAE. 30: 1403-1418.

Levine, E., and H. Oloumi-Sadeghi. 1991. Management of diabroticite rootworms in corn. Ann. Rev. Entomol. 36: 229-255.

Levine, E., H. Oloumi-Sadeghi, and J.R. Fisher. 1992a. Discovery of multiyear diapause in Illinois and South Dakota northern corn rootworm (Coleoptera: Chrysomelidae) eggs and incidence of the prolonged diapause trait in Illinois. J. Econ. Entomol. 85: 262-267.

Levine, E., H. Oloumi-Sadeghi, and C.R. Ellis. 1992b. Thermal requirements, hatching patterns, and prolonged diapause in western corn rootworm (Coleoptera: Chrysomelidae) eggs. J. Econ. Entomol. 85: 2425-2432.

Linney, J. 1991. Review of reactive planning literature. Department of Computer Science. Technical Report 560. University of London. United Kingdom.

Mackworth, A.K. 1977. Consistency in network of relations. Artificial Intelligence. 8: 99-118.

Mackworth, A.K. 1987. Constraint satisfaction. In: Encyclopedia of artificial intelligence. Shapiro, S.C., and D. Eckroth (eds.). Vol 1. Wiley. New York.

Mackworth, A.K., and E.C. Freuder. 1985. The complexity of some polynomial network consistency algorithms for constraint satisfaction problems. Artificial Intelligence. 25: 65-74.

Madden, L.V., R. Louie, J.J. Abt, and J.K. Knoke. 1982. Evaluation of tests for randomness of infected plants. Phytopathology. 72: 195-198.

Marse, K., and S.D. Roberts. 1983. Implementing a portable FORTRAN uniform (0,1) generator. Simulation. 41: 135-139.

Metcalf, R.L.1980. Changing role of insecticides in crop protection. Ann Rev. Entomol. 25: 219-256.

Metcalf, R.L., and R.A. Metcalf. 1993. Destructive and useful insects. Their habits and control. McGraw-Hill. New York.

Meinke, L.J., B.D. Siegfried, R.J. Wright, and L.D. Chandler. 1998. Adult susceptibility of Nebraska western corn rootworm (Coleoptera: Chrysomelidae) populations to selected insecticides. J. Econ. Entomol. 91: 594-600.

Michie, D., and R.A. Chambers. 1968. Boxes: An experiment in adaptive control. In: Machine intelligence 2. Dale, E., and D. Michie (eds.). American Elsevier. New York.

Mittra, S.S. 1991. Principles of relational database systems. Prentice-Hall. Englewood Cliffs, New Jersey.

Morgan, R.P.C., J.N. Quinton, and R.J. Rickson. 1994. Modelling methodology for soil erosion asssessment and soil conservation design: The EUROSEM approach. Outlook in Agric. 23: 5-9.

Morgan, R.P.C., J.N. Quinton, R.E. Smith, G. Govers, J.W.A. Poesen, K. Auerswald, G. Chisci, D. Torri, M.E. Styczen, and A.J.V. Folly. 1998. The European soil erosion model (EUROSEM): documentation and user guide. Silsoe College, Cranfield University. United Kingdom.

Nadel, B.A. 1988. Tree search and arc consistency in constraint satisfaction algorithms. In: Search in Artificial Intelligence. L. Kamal and V. Kumar (eds.). Springer-Verlag. New York.

Nadel, B.A. 1989. Constraint satisfaction algorithms. Comput. Intell. 5: 188-224.

NASS. 1998. Agricultural statistics 1998. United States Department of Agriculture. National Agriculture Statistics Service. U.S. Government Printing office. Washington D.C.

Nevo, A., R. Oad, and T.H. Podmore. 1994. An integrated expert system for optimal crop planning. Agric. Syst. 45: 73-92.

Nicks, A.D., L.J. Lane, and G.A. Gander. 1995. Weather generator. In: Water erosion prediction project. Hillslope profile and watershed model documentation. U.S. Department of Agriculture. NSERL Report No. 10. West Lafayette, IN 47907-1196.

Novotny, V. and H. Olem. 1994. Water quality. Prevention, identification, and management of diffuse pollution. Van Nostrand. New York.

NRCS. 1996a. Soil quality resource concerns: sediment deposition on cropland. National Resource Conservation Service. USDA. Soil quality information sheet. Washington D.C.

NRCS. 1996b. Conservation crop rotation effects on soil quality. National Resource Conservation Service. USDA. Soil quality-agronomy technical note No. 2. Washington D.C.

NRCS. 1998. Effects of soil erosion on soil productivity and soil quality. National Resource Conservation Service. USDA. Soil quality-agronomy technical note No. 7. Washington D.C.

NSERL. 1994. Water erosion prediction model. Version 94.7 User Summary. NSERL Report No. 9. USDA-Agricultural Research Service. National Soil Erosion Research Laboratory. West Lafayette, Indiana.

Osmond, D.L., R.W. Gannon, J.A. Gale, D.E. Line, C.B. Knott, K.A. Phillips, M.H. Turner, M.A. Foster, D.E. Lehning, S.W Coffey, and J. Spooner. 1997. WATERSHEDSS: A decision support system for watershed-scale nonpoint source water quality problems. J. Am.Water Res. Assoc. 33: 327-341.

Pearl, J., and R.E. Korf. 1987. Search techniques. Ann. Rev. Comput. Sci. 2: 451-467.

Parsons, R.L. 1995. Financial costs and economic tradeoffs of alternative manure management policies on dairy and dairy/poultry farms in Rockingham County, Va. Ph.D. Dissertation. Virginia Polytechnic Institute and State University. Blacksburg, Virginia.

Pacheco, X., and S. Teixeira. 1996. Delphi 2 developer's guide. Sams Publishing. Indianapolis, Indiana.

Pedigo, L.P. 1989. Entomology and pest management. Macmillan Publishing. New York.

Pimentel, D., L. Loughlin, A. Zepp, B. Lakitan, T. Krauss, P. Kleinman, F. Mancini, W.J. Roach,

E. Graap, W.S. Keeton, and G. Selig. 1993. Environmental and economic effects of reducing pesticide use in agriculture. Agric. Ecosyst. and Environ. 46: 273-288.

Plant, R.E., and N.D. Stone. 1991. Knowledge-based systems in agriculture. McGraw-Hill. New York.

Rellier, J.P., and S. Chedru. 1992. An artificial- intelligence based software for designing crop management plans. Computers and Elect. Agric. 6: 273-294.

Richardson, C.R., and D.A. Wright. 1984. WGEN: A model for generating daily weather variables. U.S. Department of Agriculture. Agricultural Research Service. ARS-8.

Richardson, C.W., and A.D.Nicks. 1990. Weather generator description. In: EPIC-Erosion/ Productivity Impact Calculator. 1. Model documentation. Sharpley, A.N., and J.R. Williams (eds.). USDA-ARS. Technical bulletin 1768.

Riedell, W.E., R.D. Gustin, D.L. Beck, and D.G. Hanson. 1991. Western corn rootworm damage: Effect of tillage on plant response and grain yield. Crop Sci. 31: 1293-1297.

Roth, G.W., D.D. Calvin, and S.M. Lueloff. 1995. Tillage, nitrogen timing, and planting date effects on western corn rootworm injury to corn. Agron. J. 87: 189-193.

Russell, S.J., and P. Norvig. 1995. Artificial Intelligence. Prentice-Hall. Englewood Cliffs, New Jersey.

Sadeh, N.M., and M.S. Fox. 1995. Variable and value order heuristics for the job shop scheduling constraint satisfaction problem. Technical Report CMU-Ri-Tr-95-39. Carnegie Mellon University. Pittsburgh, Pennsylvania.

SAS Institute. 1990. SAS/STAT user's guide. Vol. 2. GLM-VARCOMP. SAS Institute Inc. Cary

North Carolina.

SCS. 1982. Soil survey of Rockingham County Virginia. United States Department of Agriculture. Soil Conservation Service. U.S. Printing office. Washington, D.C.

Schaafsma, A.W., G.H. Whitefield, and C.R. Ellis. 1991. A temperature-dependent model of egg development of the western corn rootworm, *Diabrotica virgifera virgifera* LeConte (Coleoptera: Chrysomelidae). Can. Ent. 123: 1183-1197.

Schaller, N. 1993. The concept of agricultural sustainability. Agric. Ecosyst. and Environ. 46: 89-97.

Schroder, R.F.W. 1998. Corn rootworms get juiced. Agric. Res. 46: 11.

Sharifi, M.A., and H. Van Keulen. 1994. A decision support system for land use planning at farm enterprise level. Agric. Syst. 45: 239-257.

Simpson, T.W., S.J. Donohue, G.W. Hawkins, M.M. Monnett, and J.C. Baker. 1993. The development and implementation of the Virginia Agronomic Land Evaluation System (VALUES). Department of Crop and Soil Environmental Sciences. Virginia Polytechnic Institute and State University. Blacksburg, Virginia.

Smith, R.E., D.C. Goodrich, D.A. Woolhiser, and C.L. Unkrich. 1995. KINEROS- A kinematic runoff and erosion model. In: Computer models of watershed hydrology. V. P. Singh (ed.). Water Resources Publications. Highlands Ranch. Colorado.

Smith, C.S., and G.T. McDonald. 1998. Assessing the sustainability of agriculture at the planning stage. J. Environ. Manag. 52: 15-37.

Sokal, R.R., and F.J. Rohlf. 1995. Biometry. Freeman. New York.

Spike, B.P., and J.J. Tollefson. 1991. Yield response of corn subjected to western corn rootworm (Coleoptera: Chrysomelidae) infestation and lodging. J. Econ. Entomol. 84: 1585-1590.

Srinivasan, R., and B.A. Engel. 1994. A spatial decision support system for assessing agricultural nonpoint source pollution. Water Res. Bull. 30: 441-452.

Steel, R.G.D., and J.H. Torrie. 1980. Principles and procedures of statistics. McGraw-Hill. New York.

Steffey, K.L., M.E. Gray, and D.E. Kulhman. 1992. Extent of corn rootworm (Coleoptera: Chrysomelidae) larval damage in corn after soybeans: Search for the expression of the prolonged diapause trait in Illinois. J. Econ. Entomol. 85: 268-275.

Stern, V.M., R.F. Smith, R. van den Bosch, and K.S. Hagen. 1959. The integrated control concept. Hilgardia. 29. 81-101.

Stone, N.D. 1995. Object-oriented constraint satisfaction planing for whole-farm management. AI Applications. 9: 61-69.

Stone, N.D., R.D. Buick, J.W. Roach, R.K. Scheckler, and R. Rupani. 1992. The planning problem in agriculture: Farm-level crop rotation planning as an example. AI Applications. 6: 59-74.

Sumner, D.R. 1982. Crop rotation and plant productivity. In: Handbook of agricultural productivity. Vol. I. Plant productivity. M. Rechlig Jr. (ed.). CRC Press. Boca Raton, Florida.

Tsai, Y.J., J.W. Mishoe, and J.W. Jones. 1987. Optimizing multiple cropping systems: Simulation studies. Agric. Syst. 25: 165-176.

Tsang, E. 1993. Foundations of constraint satisfaction. Academic Press. London.

Turban, W. 1995. Decision support and expert systems. Prentice-Hall. Englewood Cliffs. New Jersey.

Turpin, F.T., and D.C. Peters. 1971. Survival of southern and western corn rootworm larvae in relation to soil texture. J. Econ. Entomol. 64: 1448-1451.

Turpin. F.T., L.C. Dumenil, and D.C. Peters. 1972. Edaphic and agronomic characters that affect potential for rootworm damage to corn in Iowa. J. Econ. Entomol. 65: 1615-1619.

USDA. 1994. USDA's integrated pest management (IPM) initiative. United States Department of Agriculture. Office of communications. News backgrounder. Release No. 0942.94. Washington D.C.

USEPA. 1997a. Chesapeake Bay: Nutrient reduction progress and future directions. Summary report. United States Environmental Protection Agency. Publication No. EPA903-R-97-030. Washington, D.C.

USEPA. 1997b. Compendium of tools for watershed assessment and TMDL development. United States Environmental Protection Agency. Publication No. EPA841-B-97-006. Washington, D.C.

Varshney, P., U.S. Tim, and C.E. Anderson. 1993. Risk-based evaluation of ground-water contamination by agricultural pesticides. Ground Water. 31: 356-362.

VASS. 1998. Virginia agricultural statistics. Bulletin No. 70. Virginia Agricultural Statistics Service. Richmond, Virginia.

VCE. 1993. Pest management guide for field crops 1993. Virginia Cooperative Extension. Publication 456-016. Virginia Polytechnic Institute and State University. Blacksburg, Virginia.

VCE. 1997. Virginia farm management crop and livestock enterprise budgets 1997. Virginia Cooperative Extension Publication 446-047. Virginia Polytechnic Institute and State University. Blacksburg, Virginia.

Wald, A., and J. Wolfowitz. 1940. On a test whether two samples are from the same population. Math. Stat. 11: 147-162.

Weiss, M.A. 1994. Data structures and algorithm analysis in C++. The Benjamin/Cummins Publishing Co. Redwood City, California.

Wilkins, D.E. 1984. Domain-independent planning: representation and plan generation. Artificial Intelligence. 22: 269-301.

Wilkins, D.E. 1988. Practical planning. Morgan Kauffmann Publishers. San Mateo, California.

Williams, J.R., K.G. Renard, and P.T. Dyke. 1983. A new method for assessing the effect of soil erosion on productivity. J. Soil and Water Cons. 38: 381-383.

Williams, J.R., C.A. Jones, and P.T. Dyke. 1990. The EPIC model. In: EPIC-Erosion/ Productivity Impact Calculator. 1. Model documentation. Sharpley, A.N. and J.R. Williams (eds.). USDA-ARS. Technical bulletin 1768.

Wischmeier, W.H., and D.D. Smith. 1978. Predicting erosion soil loses- A guide to conservation planning. U.S. Department of Agriculture. Agriculture Handbook No. 537. Washington D.C.

Wolfe, M.L., W.D. Batchelor, T.A. Dillaha, S. Mostaghimi, and C.D. Heatwole. 1995. A farm scale water quality planning system for evaluating best management practices. Proceedings of the International Symposium on water quality modeling. American Society of Agricultural Engineers. Orlando, Florida.

Woodson, W.D., and J. J. Jackson. 1996. Developmental rate as a function of temperature in northern corn rootworm (Coleoptera: Chrysomelidae). Ann. Entomol. Soc. Am. 89: 226-230.

Young, R.A., C.A. Onstad, D.D. Bosch, and W.P. Anderson. 1989. AGNPS: A nonpoint-source pollution model for evaluating agricultural watersheds. J. Soil and Water Conserv. 44: 168-173.

Youngman, R., K. Nixon, and C. Laub. 1993. Scouting and cultural control methods for managing western corn rootworm on corn in Virginia. In: Virginia integrated pest management. Laub, C., and W. Ravlin (eds.). Virginia Polytechnic Institute and State University. Blacksburg, Virginia.

Zalom, F.G. 1993. Reorganizing to facilitate the development and use of integrated pest management. Agric. Ecosyst. Environ. 46: 245-256.

Appendix A Crop, Farm and Rotation Information

Table A- 1 Example crop budget for a specific field.

Economic Budget for Crop: CORN SILAGE HY-CT ID: 29

Farm_id: 1	Tract_id: 1	Field_id: 1	
Acres for this field:	9.26	Expected yield for this Field:	16.67
Expected Crop Sale Price:	30.00	Expected Gross Return:	4630.93
APR:	0.09		
-------------- Pre Harvest Costs -----------			
Cost of seed:	235.82		
Cost of N:	444.94		
Cost of P:	103.71		
Cost of K:	194.46		
Cost of Fertilizer Application:	50.93		
Cost of Chemical Application:	55.56		
Chemical Product Costs			
Cost of Herbicides:	214.82		
Machinery Costs			
Cost of Fuel and Oil:	62.13		
Cost of Repair:	104.27		
Cost of labor:	77.23		
Other preharvest costs:	104.04		
Total Preharvest Costs:	1647.91		
Preharvest Input Taxes:	74.16		
------------- Harvest Costs ---------------			
Cost of Fuel and Oil:	76.49		
Cost of Machinery Repair:	269.10		
Cost of Labor:	95.01		
Total Harvest Costs:	440.59		
Machinery Fixed Costs:	802.93		
Total Costs:	2965.59		
Net Profit:	1665.33		

Table A- 2 Relevant information for the farms used in this research. Highly Erodible Land (HEL), (Cmax) is the maximum USLE-C value, (N Risk) nitrate leaching risk, (Pst Risk) pesticide leaching risk, (SL Risk) surface runoff risk. Acreage and slope calculated from GIS data (see Chapter 6). Risk and HEL values calculated according to Buick et al (1992).

Tract	*Field*	*Acreage (ha)*	*Slope (%)*	*HEL (T/F)*	*Cmax*	*N Risk*	*Pst Risk*	*SL Risk*
				Farm-1				
1	1	3.75	7.3	T	0.16	6.0	2.0	2.5
	2	1.75	2.5	F	0.69	6.0	2.0	2.0
	3	5.00	3.5	F	0.40	6.0	2.0	2.1
	4	2.50	2.7	F	0.58	6.0	2.0	2.0
	5	0.50	19.4	T	0.03	6.0	2.0	3.0
	6	5.50	9.9	T	0.09	6.0	2.0	2.8
	7	6.75	2.6	F	0.51	6.0	2.0	2.0
2	8	7.75	2.5	F	0.57	6.0	2.0	2.0
	9	1.75	2.4	F	0.58	6.0	2.0	2.0
	10	0.50	2.7	F	0.54	6.0	2.0	2.0
	11	2.50	3.7	F	0.39	6.0	2.0	2.0
3	12	13.0	3.1	F	0.53	6.0	2.0	2.9
	13	8.50	5.9	T	0.23	6.0	2.0	2.8
	14	1.75	1.8	F	0.72	6.0	2.0	3.0
				Farm-2				
1	1	9.75	3.8	F	0.41	6.0	2.0	2.0
	2	0.50	3.6	F	0.41	6.0	2.0	2.0
	3	5.50	4.2	F	0.44	6.0	2.0	2.0
	4	5.00	3.3	F	0.52	6.0	2.0	2.0
	5	1.75	3.3	F	0.52	6.0	2.0	2.0
2	6	12.25	5.0	T	0.18	5.7	2.2	2.2
	7	2.50	2.7	F	0.35	4.2	2.0	2.6
	8	9.75	6.9	T	0.13	5.8	2.3	2.4
	9	3.25	4.7	F	0.34	6.0	3.0	3.0
	10	5.25	3.0	F	0.60	5.7	2.9	2.9

Table A- 3 Crop composition and rotation frequency per crop occurrence through the years of planning. Total number of rotations is 100.

Crop name	*Number of rotations with crop present in the given planning years*						
	1	2	3	4	5	6	total
Alfalfa establishment fall	60	1	0	0	0	0	61
Alfalfa (maintenance)	43	11	4	3	0	0	61
Barley silage	37	3	1	0	0	0	41
Corn grain	35	21	5	4	4	2	71
Corn silage	28	29	12	9	5	3	86
Millet	60	1	0	0	0	0	61
Rye	8	32	18	3	7	27	95
Sorghum silage	37	3	1	0	0	0	41
Wheat grain	60	1	0	0	0	0	61
Soybeans full season	17	6	4	1	1	1	30
Orchard grass-red clover (maintenance)	0	0	0	0	0	1	1

Appendix B Data for the Corn Rootworm Control Model Experiments

Table B-1 Plan CRW risk for different level combination of annual target acreage, high risk fields and different CRW control treatments. (Control) No CRW constraints enforced by the planner. (Field constraint) Planner selected low CRW risk rotations. (Crop constraint) Planner selected low CRW risk field combinations. (Both constraints) Planner selected low CRW risk values for rotations and field combinations. Average of 20 plans obtained during 20 minutes of planning.

Target acreage (ha)[1]	*Percent of farm acreage with high risk fields*				
	10	*20*	*55*	*80*	*90*
			Control		
10	5.51	5.71	6.23	6.51	6.68
20	5.56	5.75	6.30	6.55	6.72
50	5.66	6.47	6.90	6.66	6.84
80	7.10	7.24	7.58	8.04	8.22
90	7.43	7.63	8.12	8.43	8.61
			Field Constraint		
10	5.23	5.41	5.02	5.01	4.97
20	4.80	5.04	4.87	5.02	5.00
50	5.30	5.18	5.80	5.25	5.03
80	6.64	6.74	6.82	7.42	7.63
90	7.23	7.54	8.05	8.25	8.33
			Crop constraint		
10	5.74	5.85	6.49	6.64	6.99
20	5.79	5.95	6.68	6.95	7.03
50	6.40	6.82	7.34	7.35	7.60
80	7.23	7.45	7.85	8.14	8.25
90	7.44	7.63	8.10	8.38	8.57
			Both constraints		
10	5.39	5.37	5.39	5.06	5.15
20	5.54	5.59	5.54	5.31	5.26
50	6.23	6.12	6.02	5.45	5.73
80	7.00	6.99	6.96	7.61	7.58
90	7.22	7.52	8.04	8.18	8.39

[1] annual target acreages were 8%, 28%, 58%, 78%, and 91% respectively for the level of 55% of farm area with high risk fields.

Table B-2 Corn silage net profit (US$/ Mg) for different levels of annual target acreage, percent of farm area with high-risk fields and different CRW control treatments. (Control) No CRW constraints enforced by the planner. (Field constraint) Planner selected low CRW risk rotations. (Crop constraint) Planner selected low CRW risk field combinations. (Both constraints) Planner selected low CRW risk values for rotations and field combinations. Average of 20 plans obtained during 20 minutes of planning.

Target	*Percent of farm area with high risk fields*				
acreage (ha)[1]	*10*	*20*	*55*	*80*	*90*
			Control		
10	12.608	12.608	12.606	12.290	12.290
20	12.584	12.584	12.611	12.584	12.584
50	12.632	12.587	12.600	12.632	12.632
80	12.346	12.475	12.584	12.475	12.475
90	12.292	12.297	12.292	12.292	12.301
			Field Constraint		
10	12.381	12.569	12.363	12.620	13.262
20	13.062	12.949	12.974	13.496	13.150
50	12.540	12.569	12.676	13.077	13.216
80	12.420	12.406	12.545	12.507	12.444
90	12.292	12.276	12.230	12.308	12.320
			Crop constraint		
10	12.567	12.288	11.975	11.873	11.411
20	12.667	12.373	12.027	11.814	11.571
50	12.433	12.398	12.409	12.170	12.213
80	12.455	12.212	12.451	12.342	12.336
90	12.324	12.377	12.318	12.258	12.296
			Both constraints		
10	12.332	12.539	11.975	12.316	11.670
20	12.698	12.384	11.992	12.062	12.638
50	12.448	12.246	12.659	13.152	13.026
80	12.457	12.242	12.545	12.447	12.494
90	12.295	12.366	12.313	12.299	12.315

[1] annual target acreages were 8%, 28%, 58%, 78%, and 91% respectively for the level of 55% of farm area with high risk fields.

Table B-3 Plan net profit (US$) for different level combinations of annual target acreage, high risk fields and different CRW control treatments. (Control) No CRW constraints enforced by the planner. (Field constraint) Planner selected low CRW risk rotations. (Crop constraint) Planner selected low CRW risk field combinations. (Both constraints) Planner selected low CRW risk values for rotations and field combinations. Average of 20 plans obtained during 20 minutes of planning.

Target acreage (ha)[1]	*Percent of farm area with high risk fields*				
	10	*20*	*55*	*80*	*90*
			Control		
10	199628.20	199628.20	198708.53	200363.50	200363.50
20	210627.80	210627.80	217844.65	210269.28	210269.28
50	237156.44	235691.61	244308.79	237156.44	237156.44
80	267421.14	264416.00	260330.28	264006.22	264416.00
90	278050.19	278573.53	278528.07	278050.19	278649.98
			Field Constraint		
10	194945.67	203635.82	191281.59	183357.65	172830.87
20	207441.03	202460.75	196791.02	192008.33	180531.75
50	218792.90	234298.10	228465.75	198521.39	173991.25
80	264878.77	264949.89	241902.41	256508.50	263455.85
90	273583.60	273646.53	275167.34	273087.54	271970.61
			Crop constraint		
10	198704.78	208610.03	206022.63	204117.13	206251.61
20	212601.81	215892.99	214222.99	211270.52	203908.02
50	234416.70	238366.34	240046.93	221205.16	223813.60
80	263986.97	269319.12	262082.33	258835.90	258206.91
90	277959.75	277515.29	277580.40	277355.33	277477.17
			Both constraints		
10	194547.35	202177.22	189402.48	171221.74	160303.36
20	202191.24	211891.39	204332.50	176115.61	172896.62
50	231231.67	236331.84	229600.96	190995.72	198858.82
80	257124.88	263347.27	245842.81	265531.18	262092.71
90	273235.89	272841.36	274741.16	273092.91	272709.63

[1] annual target acreages were 8%, 28%, 58%, 78%, and 91% respectively for the level of 55% of farm area with high risk fields

Vita

Jose Lopez-Collado

Jose “Pepe” was born to Galdina Collado Garcia and Francisco Lopez Sanchez in the city of Veracruz, Mexico in 1961. He spent his childhood and youth there, until he moved to Chapingo, México. There, he attended college from 1981 to 1986 at Universidad Autonoma Chapingo. He obtained his degree of Ingeniero Agrónomo Parasitólogo in 1987. He was employed at Colegio de Postgraduados, Montecillo, México, where he studied population ecology of alfalfa insect pests. From 1989 to 1991, he enrolled as a Master’s student at Colegio de Postgraduados. Once he got his degree in entomology, he taught graduate courses on statistics and quantitative insect ecology. In 1994, he married Josefina Solis Cruz, who gave him 2 children, Egil and Jazmin. Ever since they have enlightened his spare time. Living in Blacksburg for about 5 years taught him a different view of the world.

After finishing his Ph.D., he is planning to return to Mexico and rejoin his job at Colegio de Postgraduados where he expects to teach and research the application of information systems to farm and pest management.