AN INTELLIGENT EIS WITH A DATA WAREHOUSE FACILITY

A THESIS

Presented to the Department of Information Systems

California State University, Long Beach

In Partial Fulfillment

of the Requirements for the

Master of Business Administration

By Steven Joseph Gramme

BS, 1990, San Diego State University

May 1996

UMI Number: 1380232

UMI
300 North Zeeb Road
Ann Arbor, MI 48103

WE, THE UNDERSIGNED MEMBERS OF THE COMMITTEE,

HAVE APPROVED THIS THESIS

AN INTELLIGENT EIS WITH A DATA WAREHOUSE FACILITY

By

Steven Joseph Gramme

COMMITTEE MEMBERS

Robert T. H. Chi, Ph.D. (Chair)	Information Systems
Efraim Turban, Ph.D.	Information Systems
Michael Chung, Ph.D.	Information Systems

ACCEPTED AND APPROVED ON BEHALF OF THE UNIVERSITY

C. J. Walter, Ph.D.
Dean, College of Business Administration

California State University, Long Beach

May 1996

ABSTRACT

AN INTELLIGENT EIS WITH A DATA WAREHOUSE FACILITY

by

Steven Joseph Gramme

May 1996

In this thesis, I propose a system that integrates executive information system (EIS), intelligent agent, mathematical modeling, and data warehouse technologies. This intelligent EIS system will use data stored in an organization's data warehouse to answer pertinent executive questions. Unlike current EIS systems, this system is not limited to individual departmental questions; this intelligent EIS will answer executive questions company wide. This EIS will interact with a data warehouse via a user interface and coordinator and use intelligent agents and mathematical models in the EIS to manipulate data in the data warehouse into meaningful executive solutions. The intelligent agents and mathematical models include a rule-based reasoning mechanism, a case-based reasoning mechanism, a simulation modeling mechanism, a statistical modeling mechanism, and an optimization modeling mechanism. This thesis first discusses the framework of the proposed Intelligent EIS with a Data Warehouse Facility (IEISDW), then a case example that determines the best internet service provider for an organization will be presented which utilizes the tools in the system.

ACKNOWLEDGMENTS

I would like to acknowledge and thank my thesis advisor, professor, and friend, Robert T. H. Chi for all his help in the past two years with my research and writing of this thesis. He has given me priceless help and guidance.

I would also like to acknowledge my family, including my mother, four sisters, three brothers, and in-laws for their continual support via encouragement and happy times.

I would also like to acknowledge the support I received from my deceased father in the first 18 years of my life in which he encouraged my education and success.

Finally, I would like to acknowledge my high school friends, Ara Palanjian, Robert Suruki, John Williams and all of the friends I have made in my college career for their support and encouragement in my education.

TABLE OF CONTENTS

LIST OF FIGURES

CHAPTER I

INTRODUCTION

Executive Information Systems (EIS) have been successfully implemented in many organizations. Of all the various EIS commercial products, only one (Executive Edge) presents limited artificial intelligence (AI) capabilities. Yet, the ability to include various problem solving agents for collaborative information processing, filtering and presentation, is highly desirable. It is possible that successful EIS of the future will be built around AI components (Chi and Turban 1995). Additionally, a facility is needed which will allow for EIS to examine data company-wide. Current EIS utilize departmental facilities and do not have access to company-wide data without painstakingly switching and/or downloading data from one department to another.

EIS have been initially designed to focus on, monitor, filter and organize executive's information, so that executives can make more effective use of computerized information. In general, the goals of EIS (Watson, Rainer and Koh 1991) are (1) to reduce the amount of data bombarding the executive, (2) to increase the relevance, timeliness, and usability of the information that reaches the executive, (3) to focus a management team on critical success factors, (4) to enhance executive follow-through and communication with others, and (5) to track early warning indicators: such as competitors' moves or customer demands. In using an ideal EIS, an executive should be able to easily query a problem such as a reduction in sales for the Western Region of the United States and drill down through the system until he/she arrived at the source of the reduction in sales.

Problem solving agents (intelligent agents) of EIS include agents for case-based reasoning and rule-based reasoning. Since much of executive processing involves complicated problem domains, a single intelligent agent's effort may be insufficient when the information is broad in scope and complicated in nature (Chi and Turban 1995). For such situations an integration of both agents and the use of mathematical models is necessary to solve complex problems. Some mathematical models that are of significant use when there is a complicated problem domain are statistical models, optimization models, and simulation models.

Another barrier in solving executive problems that have complicated problem domains is accessing the data needed to answer the questions. The development of analytical models and decision support systems in organizations is often hindered by the lack of data in a convenient and relevant form (Subramanian, et al. 1996). The EIS of today are connected to individual departments of an organization and thus are not able to analyze company-wide data. Executive decisions are decisions made based on long-term perspectives of an entire organization. A facility is needed which will allow for EIS to examine data company-wide. The database technology being used today that stores historical data of an entire organization is the data warehouse. A simple definition of a data warehouse is: a database that is optimized for decision support (Saylor and Bansal 1995). Whereas databases in current EIS present data by itself, using a data warehouse as a facility of EIS would allow the user to manipulate data into meaningful information. The environment of the data warehouse is ideal for EIS analysis. In a relatively short period of time, data warehouses have become the technology of choice for building the data management infrastructure for decision support systems (DSS) and executive information systems (McFadden 1996).

Research indicates that various commercial EIS products give a very stable viewpoint of a companies environment. The data that they present cannot be manipulated to answer pertinent questions of executives. Business is a changing environment and thus being able to manipulate data to answer various questions is important. EISs need to present data from a dynamic viewpoint. The utilization of an intelligent EIS with a data warehouse facility would allow data to be presented dynamically.

The EIS would be composed of intelligent agents and mathematical models. As mentioned above, these agents and models include mechanisms for case-based reasoning, rule-based reasoning, statistical modeling, optimization modeling and simulation modeling. An executive would utilize a user interface of an EIS to manipulate the data in the data warehouse. An executive would download data from a data warehouse into an EIS and then manipulate the data to answer vital organizational questions. The executive would first try to answer the question intuitively based on his/her own past experiences. If the executive never experienced a particular problem then he/she could use the reasoning mechanisms or mathematical models in which they think would solve the problem. If they still cannot arrive at a solution, then they can engage the EIS to solve the problem using the agents and models collaboratively to solve the problem. The EIS would first consult a case-base reasoning mechanism to answer a given question. If there was not a previous situation that could match the current situation then a rule-based reasoning mechanism would be engaged to determine a solution. The rule-based reasoning mechanism would generalize the current problem and then try to match this problem with any previous problems. If the rule-based reasoning mechanism could not determine a solution, then a statistical model would be engaged to determine possible solutions to a given problem. The statistical model uses regression analysis and other mathematical analysis to determine a pattern of similar previous cases that may match the current case scenario. In many statistical investigations, the main goal is to establish relationships which make it possible to predict one or more variables in terms of others (Freund, 1988). If there are not enough factors to determine a solution based on regression analysis, then an optimization model would be engaged. An optimization model would try to define the solution to a problem given a set of variables. If the problem domain is too large for numerical optimization techniques then a simulation mechanism would be engaged to simulate a possible solution. A simulation model and the optimization model are tools of the EIS which could be used to simulate "what if" scenarios asked by executives. If an executive does not receive a solution to his/her problem after the request has passed through all the intelligent agents and models, then the system will continue to submit the problem to the intelligent agents and models until there are possible solutions or a request has passed through the system ten times. The interative processing allows for the intelligent agents to work collaborative to solve a problem. If after ten times passing through the loop of intelligent agents and models the system still has no solution, then the system will output a "no solution" response.

In this thesis, an Intelligent EIS with a Data Warehouse Facility (IEISDW) is proposed to determine solutions to executive problems. In this system, solutions will be determined by the coordination of an EIS, intelligent agents and models, and a data warehouse. The rest of this thesis is organized as follows. In Chapter 2, a literature review of current research in EIS, intelligent agent and mathematical models, and data warehouse technologies is presented. A discussion of distributed artificial intelligence is presented in Chapter 3. Chapter 4 discusses the strategic management of executive decisions. The architecture and design of IEISDW is proposed in Chapter 5 while Chapter 6 demonstrates the problem solving process of IEISDW.

Chapter 7 gives an illustration of IEISDW capabilities, and this paper concludes with some observations concerning the direction and potential for future works.

CHAPTER II

LITERATURE REVIEW

EIS Technology

Executive Information systems (EIS) have been initially designed to focus on, monitor, filter, and organize executives' information, so that the executives can make more effective use of computerized information (Chi and Turban 1995). In general, the goals of EIS (Watson et al. 1991) are (1) to reduce the amount of data bombarding the executive, (2) to increase the relevance, timeliness, and usability of the information that reaches the executive, (3) to focus a management team on critical success factors, (4) to enhance executive follow-through and communication with others, and (5) to track early warning indicators: such as competitors' moves, or customer demands.

Current executive information systems (EIS) give a very stable viewpoint of a companies environment and they are subject to a bounded rationality. Since the business environment is constantly changing and is so vast, being able to manipulate and control data to answer various questions is important. An ideal EIS presents data from a dynamic viewpoint and thus can easily control and manipulate data into meaningful information.

Limitations of EIS

Current executive information systems are insufficient to deal with very complex problems.

> The capacity of the human mind for formulating and solving complex problems is very small compared with the size of the problems whose solution is required for objectively rational behavior in the real world or even for a reasonable approximation to such objective rationality (Simon 1957).

Executive information systems of today display summaries of data which must be interpreted into meaningful information by executives. EIS do not go through the process of reasoning the solution to a given problem. In many situations, the information displayed on EIS is too much information for a human mind to handle. At some point, an executive must have some help in their reasoning process. The limitation of a human minds processing capacity is called bounded rationality. "bounded rationality" implies that both the information a person can absorb and the detail of control is limited. As tasks grow larger and more complex, means must be found to effectively limit the increase of information a person sees and the complexity of control (Fox 1981).

Most executive information systems (EIS) have extended the "rationality" of a single human problem solver, but still have limited processing capabilities. The inference engine in executive information systems can process only a limited number of instructions within a given time period. Further, in order to maintain the efficiency of information retrieval, the problem domain of EIS is usually restricted to a narrowly defined area. Hence, EIS exhibit symptoms similar to the bounded rationality exhibited by human problem solvers when processing capacities are exceeded.

Intelligent Agents and Mathematical Models

Since executives (and their assistants) perform a variety of tasks, it is logical to propose a support system with multiple expert systems. The logic for such an architecture is that expert systems (ES) and or other intelligent agents are confined to narrow domains. Therefore a single ES cannot support all tasks (Chi and Turban 1995). This section will discuss the intelligent agents for case-based reasoning and rule-based reasoning. Additionally, it will discuss the mathematical models used to reason mathematical solutions. Each intelligent agent and mathematical models proposed in the IEISDW follows.

Case-Based Reasoning

Case-based reasoning (CBR) is a reasoning method which uses previous experiences to solve new problems. It allows a problem solver to focus on attributes of a problem which were experienced to be important to solve a similar type of problem. Therefore, new problems can be solved by matching the essential features on the old cases that were successfully solved before, such that possible incorrect lines of reasoning may be avoided. The main process of CBR methods involves remembering and adapting. Other important research issues of CBR methods include knowledge representation in CBR (Koton and Chase 1989), how to focus on important features, how to perform the matching process by applying different search strategies, and how to select the best old case if more than one match exists. Some of the developments in CBR methods include THE MEDIATOR (Kolodner and Simpson 1989), THE PERSUADER (Sycara 1988), PLEXUS (McCartney 1993), COACH (Ketler 1993), CASEY (Ketler 1993), JULIA (Hinrichs 1989; Kolodner and Simpson 1989), and MACE (Hunt 1995). These CBR methods address problems ranging from planning and design to legal and medical. Each of these uses different control strategies to manipulate the adapting process, but neither of them uses a mechanism to generalize previous cases in order to cover a wider problem space. In short, CBR methods can improve the problem solving process by allowing the executive to solve a problem by inferring from previous experiences; it allows the executive to focus on important features of the problem domain; and it avoids possible wrong directions which may lead to previous mistakes.

In general, the CBR method provides a forceful reasoning method for experience rich and knowledge poor problem domains. Frames are used to organize knowledge based on shared characteristics. An object is defined by its shared characteristics and can be related to any other object in the database that shares those characteristics. Most knowledge is represented in the frames and rules are needed to perform inferencing.

Each frame can be represented in terms of the language of first predicate calculus (Nilsson 1980) for the purpose of matching premises in production rules. In IEISDW, each predicate contains two elementary components: Function-symbols and Variable-symbols. A specific value associated with an attribute in a frame can be represented by a predicate with a Function-symbol with two Variable-symbols. Function-symbols are used to represent the attributes of frames, while the first Variable-symbol is the name of the frame and the second

Variable-symbol is the value of that attribute. Consider the following example:

Frame name: computer Brand: IBM

Model: 486

The above frame can be represented as Brand(computer, IBM) and Model(computer, 486), where "Brand" and "Model" are Function-symbols and "computer," "IBM" and "486" are Variable-symbols. Variable-symbols can be values of "variables." Model(computer, X) is a predicate with a variable of the computer styles. This means the computer models can match any symbols or numbers (Chi and Gramme 1995).

Rule-Based Reasoning

Many EIS use a rule-based reasoning mechanism to solve problems associated with deductive reasoning.

The rule-based system has been used most in current expert systems. The knowledge base of a rule-based system is specific to the domain of application and consists of two different types of knowledge: facts and rules. Facts are true to the domain, and rules describe relations or phenomena in the same problem domain. Rules represented as if-then statements often define logical relations between concepts of the problem domain (Bratko, 1986).

Psychological modeling is designed to create a program that embodies a theory of human performance of simple tasks. From the performance records of experimental human subjects, the modeler formulates the set of production rules minimally competent to reproduce the behavior (Buchanan and Shortcliffe, 1984). If no general rules or logical relations can be derived from the problem under analysis, it will be difficult to apply, the rule-based system to this problem domain.

Statistical Modeling (Regression Analysis)

A statistical model (regression analysis), is used for finding relationships among variables. In many statistical investigations, the main goal is to establish relationships which make it possible to predict one or more variables in terms of others. The problem of predicting the average value of one variable in terms of the known value of another variable (or the known values or other values of other variables) is called the problem of regression (Freund 1988).

Whenever possible, we try to express, or approximate relationships between known quantities that are to be predicted in terms of mathematical equations.

Whenever we use observed data to arrive at a mathematical equation which describes the relationship between two variables--a procedure known as curve fitting--we must face three kinds of problems:

1. We must decide what kind of curve, and hence what kind of predicting equation we want to use.

2. We must find the particular equation which is "best" in some since.

3. We must investigate certain questions regarding the merits of the particular equation, and of predictions made of it.

The first kind of problem is usually decided by direct inspection of the data, we plot the data on a graph, and we decide by visual inspection upon the kind of curve (a straight line, parabola,...) which best describes the overall pattern of the data.

The equation for a line passing through the data points is called a linear equation. It is often in the form: $y = a + bx$, where "a" is the y-intercept (the value of y for $x = 0$) and "b" is the slope of the line (the change in y which accompanies the increase in one unit of x). Linear equations are useful and important not only because many relationships are actually of this form, but also because they provide close approximations to relationships which would otherwise be difficult to describe in mathematical terms.

Once it is decided to fit a straight line to a given set of data, the second kind of problem must be addressed, namely that of finding the equation of the particular line which in some since provides the best possible fit. To illustrate what is involved, consider the following sample of data of a fictitious IS course and the relationship in the number of years they have studied IS to their scores on an IS proficiency exam.

Number of years (x)	Score on test (y)
3	57
4	78
4	72
2	58
5	89
3	63
4	73
5	84
3	75
2	48

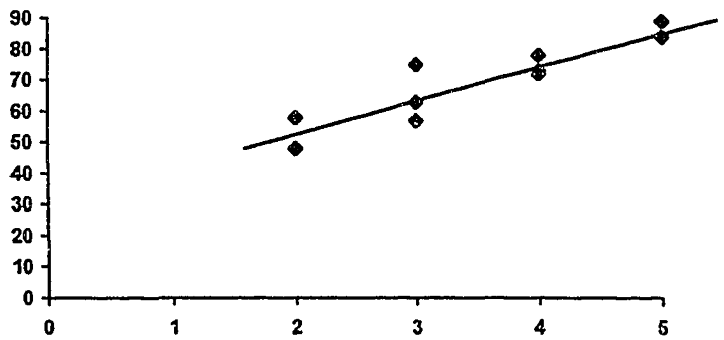

Figure 1. Regression Analysis

If the points are plotted which correspond to these ten pairs of values as in Figure 1, we observe that, even though the points do not fall on a straight line, the overall pattern of the relationship is reasonably well described with the line. There is no known departure from linearity in the scatter of the points, thus it would be justified to say that the straight line is a suitable description of the underlying relationship.

We know face the problem of finding the equation of the line. The criterion which is used almost exclusively for defining a "best" fit is known as the method of least squares. This method requires that the line which we fit to our data be such that the sum of the square of the vertical distances from the points to a line is a minimum. Thus we must solve for a and b in the equation: $y = a + bx$.

An example of a statistical modeling and simulation software currently being used is "Prophesy! 2.0" made by Abstraction Software (Stearns 1995).

Optimization Modeling

Ideally, an optimization model will generate in seconds, an optimal order level. For structured situations there is very user-friendly software to conduct such an analysis. Optimization is limited because it will work only if the problem is structured. Specifically, such a model will specify the required input data, the desired output, and the mathematical relationship in a precise manner. Obviously, if reality differs significantly from the model, optimization cannot be used. DSS deals with unstructured problems. But this does not preclude the use of optimization. Many times it is possible to break a problem into sub-problems, some of which are structured enough to justify optimization. Also, optimization can be combined with simulation for the solution of complex problems (Turban 1995).

Mathematical programming is the name for a family of tools designed to help solve managerial problems in which the decision maker must allocate scarce resources (such as labor, capital, machines, or water) among various activities to optimize a measurable goal.

The uses of mathematical programming, especially of linear programming, are so common that "canned" computer programs can be found today in just about any organization that has a computer.

DSS development tools, such as Lotus 1-2-3 or IFPS Plus, can be used to model and solve linear programming situations or have the capacity to interface with a "canned" LP program.

Every LP problem is composed of:

1. Decision Variables. The variables whose values are unknown and are searched for. Usually they are designated X1, X2, and so on.

Objective Function. This is a mathematical expression, given as a linear function, that shows the relationship between the decision variables and a single goal (or objective) under consideration. The objective function is the measure of goal attainment. Examples of such goals are total profit, total cost, share of the market, and the like.

2. Optimization. Linear programming attempts to either maximize or minimize the values of the objective function.

3. Coefficients of the Objective Function. The coefficients of the variables in the objective function are called the profit (or cost) coefficients. They express the rate at which the value of the object function increases or decreases by including in the solution one unit of each of the decision variables.

4. Constraints. The maximization (or minimization) is performed subject to a set of constraints. Therefore linear programming can be defined as a constrained optimization problem. These constraints are expressed in the form of linear inequalities (or, sometimes,

equalities). They reflect the fact that resources are limited or they specify some requirements.

5. Input-Output (Technology) Coefficients. The coefficients of the constraints' variables are called the input-output coefficients. They indicate the rate at which a given resource is depleted or utilized. They appear on the left hand side of the constraints.

6. Capacities. The capacities (or availability) of the various resources, usually expressed as some upper or lower limit, are given on the right hand side of the constraints. The right hand side also expresses minimum requirements (Turban 1995).

7. Example. These major components of a linear programming model are illustrated in the following:

Find X1, X2 (decision variables) that will minimize the value of the linear objective function:

cost coefficients

$$Z = 40X1 + 10X2$$

decision variables

subject to the linear constraints:

$$2X1 + 2X2 \geq 300$$
$$3X1 + 0X2 \geq 250$$

input-output coefficients — capacities or requirements

An example of an optimization type problem is an electronic circuit with several components, the properties of which dither around a specific nominal value. The dithering is the result of the production process of the components and cannot be kept below a certain bound, which means a percentage of electronic circuits do not meet the specifications. Optimization determines the nominal value of the critical components in a way that maximizes production yield. The optimum must be found while operating within specific constraints, which is called "constrained optimization." Optimizing these problems using a computer requires that they first be formulated mathematically; most of algorithms used for constrained-optimization problems require that the problem be formulated through an objective (or 'cost') function that has to be minimized (Storn 1995).

Simulation Modeling

The term "simulation" means to assume the appearance of characteristics of reality. In MIS it generally refers to a technique for conducting experiments (such as "what-if") with a digital computer on a model of a management system.

To begin, simulation is not strictly a type of model; models in general represent reality, while simulation usually imitates it. In practical terms, this means that there are fewer simplifications of reality in simulation models than in other models.

Second, simulation is a technique for conducting experiments. Therefore, simulation involves the testing of specific values of the decision or uncontrollable variables in the model and observing the impact on the output variables.

Simulation is a descriptive rather than a normative tool; that is there is no automatic search for an optimal solution. Instead, a simulation describes and/or predicts the characteristics of a given system under different circumstances. Once these characteristics are known, the best among several alternatives can be selected. The simulation process often consists of the repetition of an experiment many, many times to obtain an estimate of the overall affect of certain actions. It can be executed manually in some cases, but a computer is usually needed.

Finally, simulation is usually called for only when the problem under investigation is too complex to be treated by numerical optimization (e.g., because the assumptions do not hold) or the formulation is too complex.

Simulation involves setting up a model of a real system and conducting repetitive experiments on it. The methodology consists of a number of steps. The following is a brief discussion of the process:

1. Problem Definition. The real-world problem is examined and classified. Here we should specify why simulation is necessary. The system's boundaries and other such aspects of problem classification are attended to here.

2. Construction of the Simulation Model. This step involves gathering the necessary data. In many cases, a flowchart is used to describe the process. Then a computer program is written.

3. Testing and Validating the Model. The simulation model must properly imitate the system under study. This involves the process of validation.

4. Design of the Experiments. Once the model has been proven valid, the experiment is designed. Included in this step is determining how long to run the simulation. This step deals with two important and contradictory objectives: accuracy and cost.

5. Conducting the Experiments. Conducting the experiment may involve issues such as random number

generation, stopping rules, and presentation of the results.

6. Evaluating the Results. The next step is the evaluation of the results. Here, we deal with issues such as "What do the results mean?" in addition to statistical tools, we may use a sensitivity analysis (e.g., in the form of "what-if" questions).

7. Implementation. The implementation of the simulation results involve the same issues as any other implementation. However, the chances of implementation are better since the manager is usually more involved in the simulation process than with analytical models.

Four simulation modeling software packages currently being used in the business world are Scitor Corp's Process Charter 1.0.2, Modell Data AS' Powersim 2.01, High Performance Systems' ithink 3.0.61, and That!'s Superior Extend+BPR 3.1 (Tyo 1995).

Data Warehouse Technology

The term "data warehouse" was first introduced by William Inmon (Baer 1995), and proposed as a technology by IBM in 1991. Typically, there are two types of data warehouses: traditional (or physical) and virtual (or logical) data warehouses. Traditional data warehouses use metadata menus to create a common language for different data sources, enabling information sharing from a high performance, central data warehouse. A virtual data warehouse does not have its information built from a central data menu. Instead, it has an easy-to-read front end that can query any database source split over the entire enterprise. "Data marts" (Demarest 1994), as an example, are parts of a virtual data warehouse.

The primary objective of having a data warehouse is to support decision making by providing information analysis capability. However, presently corporations are maximizing availability and access to data, but they are not maximizing the use of the data. For example, corporations are using data warehouse technology to solve problems of information accessing or to improve operations (Janah 1995). Other corporate uses of data warehouses have included project monitoring (Warren 1994), or statistical analysis (Bull 1995). And in a few corporations, virtual data warehouses have been set up for financial analysts to query a target company's database for performing acquisition analyses (Eskow 1995).

Data warehouses pull information from networked applications and databases into a data store, which provides a single place for running queries and mining business trends.

A data warehouse is a subject-oriented, integrated, time variant, nonvolatile collection of data in support of management's decision making process.

Raw elements of data entering the data warehouse come from the unintegrated operational environment (the older "legacy" systems environment) in almost every case. The data warehouse is always a store of data that is physically seperated from the old operational, legacy systems environment even though the data in the warehouse is transformed from the application data found in the oprational environment.

The following is a definition of the features and components of a data warehouse.

Features of a Data Warehouse

1. Subject Oriented. The first notable characterisitc of the data warehouse is that it is organized around the major subjects of the enterprise. The orientation around the major subject areas of the corporation causes the data design to be data driven.

The operational world is designed around applications and functions such as loans, savings, bank cards, and trust processing. The data warehouse world is organized around major subjects such as customer, vendor, product, and activity. The alignment of the data warehouse around major subject areas of the corporation affects the design and implementation of the data found in the data warehouse.

2. Integration. The most important feature of the data warehouse environment and its second distinctive characterstic, is that data housed within it is integrated. With integration, the data warehouse takes on a very corporate flavor.

The integration of data warehouse data shows up in many different ways--in consistent naming conventions, consistent measurement of variables, consistent encoding structures, and consistent physical attributes of data.

When the DSS analyst approaches the data warehouse, the focus of the analyst should be on using the data that is in the warehouse, rather than on wondering about the credibility of consistency of the data.

3. Time Variancy. All the data in the data warehouse is accurate as of some moment in time. This basic characteristic of data in the warehouse is very different from that of data found in the operational environment. In the operational environment, data is accurate as the moment of access. In other words, in the operational environment, when you access a unit of data, you expect that it will reflect accurate values as of the moment of access.

The time variancy of data warehouse data shows up in several ways. The simplest way is that data warehouse data represents data over a long time horizon--from five to ten years. Secondly, every key structure in the data warehouse, contains implicitly, or explicitly, an element of time, such as day, week and month. The element of time is almost always found at the bottom of the concatenated key found in the data warehouse. On occassion, the element of time will exist implicitly, such as the case where an entire file is duplicated at the end of a month, or the quarter. Thirdly, data warehouse data, once correctly recorded, cannot be updated.

4. Nonvolatility. The fourth defining characteristic of the data warehouse is that it is non-volatile. There are only two kinds of operations that occur in the data warehouse: the intitial loading of the data, and the access of the data. There is no update of the data in the data warehouse as a normal part of the processing.

There are some very powerful consequences of this basic difference between operational processing and data warehouse processing. At the design level, the need to be cautious of the update anamoly is no factor in the data warehouse, since update of data is not done. This means at the physical level of design, liberties can be taken to optimize the access of data, particularly when dealing with the issues of normalization and physical denormalization. Another consequence of the simplicity of data warehouse operation is in the underlying technology used to run the data warehouse environment. Having to support record by record update in a on-line mode requires technology that has a very complex foundation underneath a facade of simplicity. The technology supporting backup and recovery, transaction and data integrity, and the detection and remedy of dead lock is quite complex and unnecessary for data warehouse processing.

The source of nearly all data warehouse data is the operational environment. It is a temptation to think that there is massive redundancy of data between the operational and the data warehouse environment. But in fact there is little or no redundancy between the two environments.

Consider the following:

1. Data is filtered as it passes from the operational environment to the data warehouse environment. Much data never passes out of the operational environment. Only data that is needed for

DSS processing finds its way into the data warehouse environment.

2. The time horizon of data is very different from one environment to the next. Data in the operational environment is very fresh. Data in the warehouse is much older. From the perspective of time horizon alone, there is very little overlap between the operational and the data warehouse environments.

3. The data warehouse contains summary data that is never found in the operational environment.

4. Data undergoes a fundamental transformation as it passes into the data warehouse. It is not the same data that resides in the operational environment from the standpoint of integration (Inmon and Hackathorn 1994).

In light of these factors, data redundancy between the two environments is a rare occurrence, resulting in less than 1 percent redundancy between the two environments.

The Components of a Data Warehouse

The components of a data warehouse are: (1) Current detail data, (2) Old detail data, (3) Lightly summarized data, (4) Highly summarized data, and (5) Metadata.

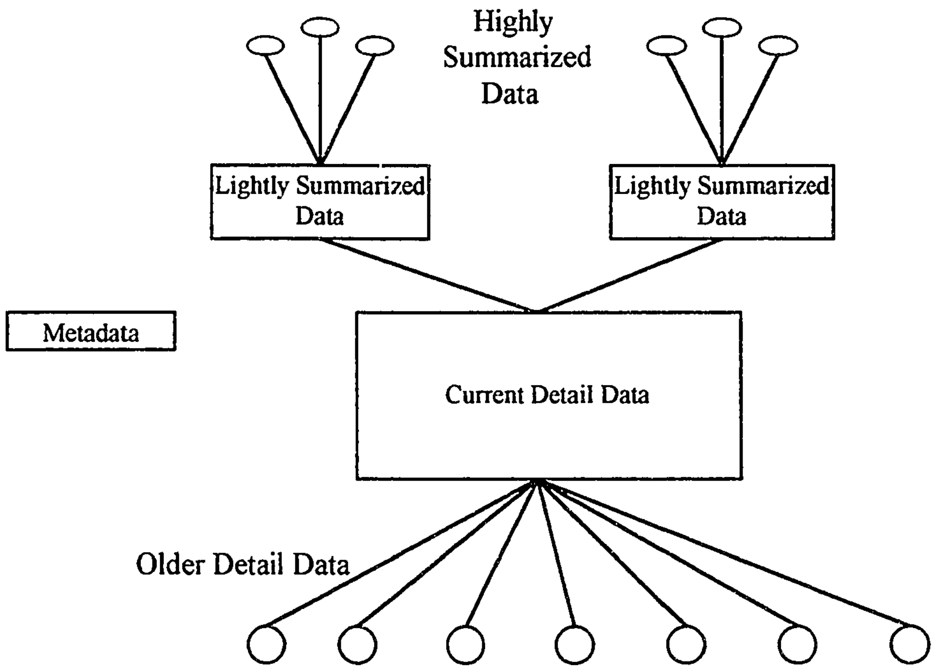

Figure 2. The structure of data in the data warehouse

1. Current Detail Data. The major area of concern in a data warehouse is the current detail data. Current detail data reflects the most recent happenings, which are always of great interest. Current detail data is voluminous because it is stored at the lowest level of granularity. Current detail data is almost always stored on disk storage, which is fast to access, but expensive and complex to manage.

2. Old Detail Data. Older detail data is data that is stored on some form of mass storage. It is infrequently accessed and is stored at a level of detail consistent with current detailed data. While it is not

mandatory that it be stored on an alternate medium, because of the anticipated large volume of data coupled with the infrequent access of the data, the storage medium for older detail data is usually not on disk.

3. Lightly Summarized Data. Lightly summarized data is data that is distilled from the low level of detail found at the current detailed level. This level of data warehouse is almost always stored on disk. The design issues facing the data architect in building this level of the data warehouse are the following:
"Over what unit of time is the summarization done?"
"What contents--attributes--will the lightly summarized data contain?"

4. Highly Summarized Data. The next level of data found in the data warehouse is that of the highly summarized data. Highly summarized data is compact and easily accessible. Sometimes this data is found outside the immediate walls of the technology that houses the data warehouse.

5. Metadata. The final component of the data warehouse is that of metadata. In many ways metadata sits in a different dimension then other data warehouse data, because metadata contains no data directly taken form the operational environment. Metadata plays a

special and very important role in the data warehouse, and is used in various ways.

It is a directory to help the DSS analyst locate the contents of the data warehouse.

It is a guide to the mapping of data as the data is transformed from the operational environment to the data warehouse environment.

It is a guide to the algorithms used for summarization between the current detail data and the lightly summarized data and the lightly summarized data and the highly summarized data, etc.(Inmon and Hackathorn 1994).

Today, building a data warehouse is costly and time consuming because data must be moved from existing data stores--sometimes manually--and metadata needs to be translated into a common format. Metadata is information about data from operational databases and other applications, and it typically is based on proprietary format. So, for example, it is not uncommon for one database to define sales numbers one way and another to do it a different way (Cole 1995).

Data warehousing technology focuses on the management of five (5) primary information flows. The first four flows (inflow, upflow, downflow, and outflow) are used to get data in from legacy systems up to a more compact form, down to archival storage, and out to consumers, and a fifth flow--metaflow--(is) used to manage the data warehouse itself (Hackathorn 1995). However, the technology has just become a popular subject recently because of improvements of database technologies and software and hardware accessing tools available.

Limitations of the Data Warehouse

Data warehouses are hot, but the technology is not without problems (Richman 1995). For example, although Relational Databases have been used to build data warehouses for a while, there are several limitations when using relational technology for data warehouse applications. The limitations include processing overhead for complex queries, demanding size of answer tables for complex queries, and lack of time-sensitive data manipulation (Thompson 1994). Other database technologies include multidimensional databases (Maurice 1994), and specialized, proprietary databases such as the Red Brick Warehouse from Red Brick Systems, Los Gatos, California.

Operational Data Storage (ODS)

An ODS is a (1) subject-oriented, (2) integrated, (3) volatile, (4) current or near current collection of data in support of detailed operational decisions.

The operational data storage is used almost exclusively for the operational information processing of the corporation. Once the operational data is outdated it is transferred into the data warehouse (see figure 3).

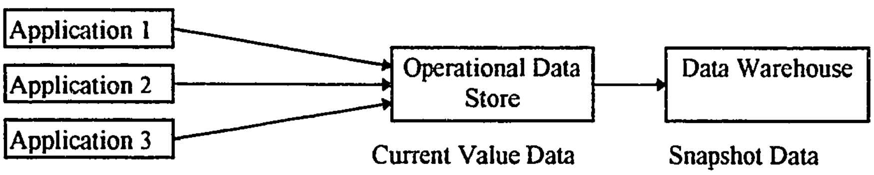

Figure 3. The Flow of Data to the Data Warehouse

The operational data storage is different from the data warehouse in that the data warehouse is non-volatile, and the operational data storage is volatile. Also, the data warehouse contains a long history of data while the operational data storage contains current and near current data.

The operational data storage forms a foundation for operational information processing. There is a difference between the informational processing found in the data warehouse. The informational processing found in the operational data storage is for the clerical community, making detailed, up to the second decisions. There is a secondary audience for informational processing in the operational data storage and that is the manager who is making corporate-wide, up to the second decisions, such as in purchasing, reordering, restocking, and product manufacturing.

The type of informational decisions being made out of the data warehouse is quite different. The data warehouse is used for long-term, analytical processing. The data warehouse serves primarily the management community looking at the larger perspective.

Nowhere does the difference between the operational data storage and the data warehouse show up more clearly than in the use of EIS. When EIS is used in the operational data storage, EIS is aimed at managers who have to make up to the second decisions. When EIS is used in the data warehouse, the analysis created out of the EIS is used to make long-term, directional decisions.

Information and the Organization

The end user can be at any level of the organization. The forms of informational processing vary from one level of management to the next. The type of informational processing of top management is typically EIS and key indicator analysis. Information for top management is usually very summary and strategic in nature. Top management's emphasis on informational processing is on long-term trends (Inmon and Hackathorn 1994).

Utlizing a data warehouse allows information to be retrieved organization-wide. If top management needs to analyze data across departments, the data warehouse will allow this. Data can be analyzed from the marketing, finance and accounting departments for example. Data from each of the departments is conglomerated in the data warehouse.

Analyzing organization-wide data allows an executive to make decisions for the organization as a whole. Being able to make decisions for an organization as a whole rather than piece mealing information from each department allows for better decision making for the organization.

CHAPTER III

DISTRIBUTED ARTIFICIAL INTELLIGENCE

Distributed artificial intelligence involves the cooperation of multiple intelligent agents. The multiple intelligent agents enable the processing of more complicated information by the cooperative efforts of various agents. In fact, many executive tasks contain various aspects of problem domains which cannot be supported by a single data retrieving and processing mechanism (Chi and Turban 1995).

According to distributed artificial intelligence studies, distributed problem solving consists of two forms (1) task-sharing and (2) result-sharing (Smith and Davis 1981). In task-sharing systems, nodes assist each other by sharing the computational load for the execution of sub-tasks of the overall problem, while in result-sharing systems, nodes assist each other by sharing partial results which are based on somewhat different perspectives on the overall problem. The result-sharing system is used in the IEISDW and is discussed in detail in the following section (Chi 1991).

Result-sharing Systems

Result-sharing is a form of cooperation in which individual nodes assist each other by sharing partial results, based on somewhat different perspectives of the overall problem (Smith and Davis 1981). In this type of system, control is typically "data-oriented." At any instant, the computation done by a certain agent is used to satisfy an unknown variable which is needed for the computation of another sub-task (see figure 4).

Therefore, an explicit hierarchy of task-sub-task relationships does not exceed between individual nodes. Typically, one of the agents acts as the group planner (or the coordinator), and each of the other agents sends all pertinent information to this agents in order to form a global plan for problem solving. The main issue of these systems is how to guide and coordinate the interactions among the participating agents, so that the problem can be solved jointly by the group.

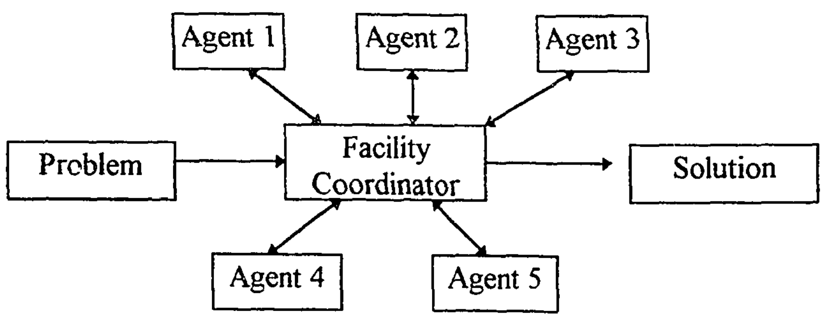

Figure 4 Framework of a Result-sharing System

In general, result sharing distributed artificial intelligence systems are most useful in problem domains in which (1) results achieved by one node influence or constrain those that can be achieved by another node (i.e., the results are relevant to each other), (2) sharing of results drives the system to converge to a solution of the problem (i.e., results received from remote nodes do not cause oscillation), and (3) sharing of results drives the system to a correct solution of the problem (Smith and Davis 1981).

Organizational Structure of Agents

The organizational structure of the agents determines the amount of information processed and the coordination necessary for the agents to operate efficiently (Shaw and Fox 1991). There are two types of structures:

Centralized Structure

This type of structure consists of coordinators and problem solving agents. The coordinators are responsible for performing the management function in the system which includes job assignment, setting job priority, current status posting and evaluations. A coordinator can give commands directly to other agents or to other sub-coordinators in a hierarchical structure. Therefore, a multiple hierarchical structure may be used if different groups of agents are involved. Centralized structure has advantages such as having one (or a few) coordinator spending its resources on reasoning about coordination and having coordinators enforce consistent views. However, the coordinator may become a potential bottleneck and the entire network could be prone to collapse if it fails.

Decentralized Structure

Another type of organizational structure of agents is called the decentralized structure. Each agent in this type of structure coordinates itself. However, having each agent coordinate itself, although more reliable, could incur large amounts of overhead since more information must be exchanged and agents may duplicate another agent's reasoning about coordination.

Message Communication Between Agents

Intelligent agents communicate messages to one another in different ways and protocols are set up to standardize the communication channels.

Communication Channels

The communication channels facilitate message passage between agents. To improve the communication efficiency, blackboard systems and mailbox systems can be employed.

1. The blackboard system. A blackboard system provides a public bulletin board (or data base) so that all agents can share information. An advantage of this type of channel is more efficiency for message passing. However, it may cause a traffic jam if more agents intend to update or retrieve information from the blackboard.

2. The mailbox system. A mailbox system provides a temporary storage device for incoming messages and face-to-face communication channels between different agents. This type of communication channel provides a more efficient way for message sending if only two agents are involved (i.e., the receiver does not have to interact with the blackboard for knowledge retrieving). However, it takes more time if several receivers are involved (i.e., the sender has to send the same message several times).

Communication Protocol

In general, the communication protocol involves using signaling to indicate the status and the request for problem solving of agents. For example, in a concurrent system, an agent may send a request signal for problem solving while the target agent is busy working on another request. The request signal may be returned by the target agent with a busy signal. Two communication protocols have been developed for distributed problem solving systems.

1. Sequential problem solving. Each agent processes it problems sequentially. In this type of system, the whole system has to wait for the process of each agent. Therefore, no signaling agent is necessary.

2. Concurrent problem solving. Agents can process their problems simultaneously. Since a particular agent may be busy performing certain tasks when a message is sent to this agent, a signaling control system should be implemented. Shaw and Fox (1991) employed a status data base to indicate the busy status of agents.

CHAPTER IV

STRATEGIC MANAGEMENT

As described in Chapter 2, the goals of EIS are (1) to reduce the amount of data bombarding the executive, (2) to increase the relevance, timeliness, and usability of the information that reaches the executive, (3) to focus a management team on critical success factors, (4) to enhance executive follow-through and communication with others, and (5) to track early warning indicators: such as competitors' moves, or customer demands.

In order for EIS to reach these goals they must be encoded with information regarding the critical success factors and early warning signs of an organization. To extract these criteria, interviews are conducted with top managers to determine the factors they believe are critical to the success of the organization and what the possible things to look out for. Once these critical success factors and warning signs are determined, they will be encoded into the IEISDW and set the constraints and rules for the intelligent agents. The determination of the critical success factors and warning signs of a firm starts with the understanding of the overall strategy of the firm. This chapter will discuss the dimensions of strategy, the strategic planning process, the definition of business strategy and planning scenarios.

The Dimensions of Strategy

Strategy can be seen as a multidimensional concept that embraces all the critical activities of the firm, providing it with a since of unity, direction, and purpose, as well as facilitating the necessary changes induced by its environment (Hax and Majluf, 1991). The following include critical dimensions that contribute to a unified definition of the concept of strategy.

1. Strategy as a coherent, unifying, and integrative pattern of decisions.

2. Strategy as a means of establishing the organizational purpose in terms of its long term objectives, action plans, and resource allocation priorities.

3. Strategy as a definition of the competitive domain of the firm.

4. Strategy as a response to external opportunities and threats, and internal strengths and weaknesses, in order to achieve competitive advantage.

5. Strategy as a channel to differentiate managerial tasks at the corporate, business and functional levels.

6. Strategy as a definition of the economic and non-economic contribution the firm intends to make to its stakeholders (Hax and Majluf, 1991).

These critical dimensions of strategy give an overall understanding of what strategy is. Next I will discuss the process of strategic planning.

The Strategic Planning Process

The strategic planning process is a disciplined and well-defined organizational effort aimed at the complete specification of a firm's strategy and the assignment of responsibilities for its execution. The planning process appropriate for a single business firm with a purely functional organizational structure is quite different from one suitable for addressing the strategic tasks of a highly diversified global corporation. There are, however, basic commonalties found in the formal planning process of most business firms. These are the hierarchical levels participating in the process, the planning tasks at each one of these levels, and the sequence in which the task should be executed (Hax and Majluf, 1991).

A formal planning process should recognize the different roles to be played by the various managers within the business organization in the formulation and execution of their firms strategies. There are three basic conceptual hierarchical levels which have always been identified as the essential layers of any formal planning process: corporate, business and functional levels (see figure 5).

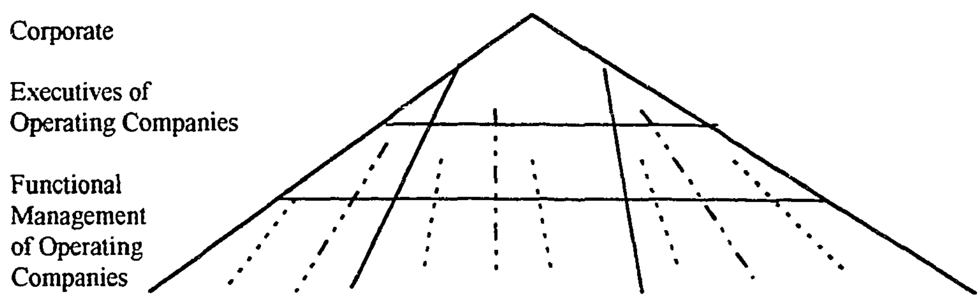

Figure 5. Management Levels

At the corporate level resides the decisions which, by their nature, should be addressed with full corporate scope. These are decisions that cannot be decentralized without running the risk of committing severe sub-optimization errors. Those who operate at lower levels of the firm do not have the proper vantage point to make the difficult tradeoffs required to maximize the benefits for the corporation as a whole, mainly when confronted with situations that affect adversely their own unit in the organization. It should be noticed that the decision maker at the corporate level is not necessarily the isolated CEO. Depending, among other things, on the management style of the CEO, corporate strategies might be shaped and implemented by the core team of top executives.

At the business level reside the main efforts aimed at securing the long-term competitive advantage in all the current businesses of the firm. Business managers are supposed to formulate and implement strategic actions congruent with the general corporate directions, constrained by the overall resources assigned to the particular business unit.

Finally, functional strategies not only consolidate the functional requirements demanded by the composite of businesses of the firm but also constitute the depositories of the ultimate competitive weapons to develop the unique competencies of the firm.

In figure 6 a model for the strategic planning process is presented which recognizes the three essential layers of managerial decision making. It also serves to illustrate the different nature of planning tasks undertaken by each level, and a possible sequence for the execution of those tasks. Individual responsibilities have to be assigned at all levels in the organization, for developing, implementing, and controlling the proper strategic tasks.

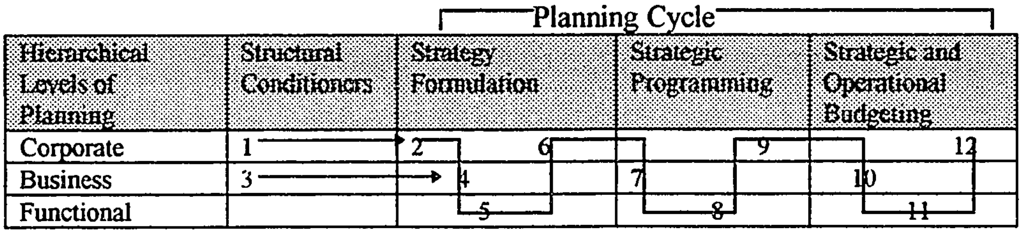

Figure 6. The Strategic Planning Cycle

1.(a) Vision of the firm, business segmentation, horizontal and vertical integration, corporate philosophy, special strategic issues. (b) Managerial infrastructure, corporate culture, and management of key personnel.

2. Strategic posture and planning guidelines: corporate strategic thrusts, planning challenges at corporate, business and functional levels, and corporate performance objectives.

3. The mission of the business: business scope, ways to compete, and identification of product market segments.

4. Formulation of business strategy and broad action plans.

5. Formulation of functional strategy: participation of business planning, concurrence or non-concurrence to business strategy proposals, broad action programs.

6. Consolidation of business and functional strategies, portfolio management, and assignment of resource allocation priorities.

7. Definition and evaluation of specific action programs at the business level.

8. Definition and evaluation of specific action programs at the functional level.

9. Resource allocation and definition of performance measurements for management control.

10. Budgeting at the business level.

11. Budgeting at the functional level.

12. Budgeting consolidations, and approval of strategic and operational funds.

The emphasis of this paper will be on the decisions being made by executives since the IEISDW is a system developed for executive decision making. Therefore, according to this model, executive decision are decisions dealing with the vision of the firm, managerial infrastructure, consolidation of business and functional strategies, resource allocation and definition of performance measures for management control, budgeting consolidations, and approval of strategic and operational funds.

The Definition of a Business Strategy

The first task consists in defining the mission of the business, in terms of its product, market, and geographical scope and a way to develop the necessary unique competencies that will assure a sustainable competitive advantage. Then we face the two central analytical tasks: The environmental scan which leads to the identification of opportunities and threats; and the internal scrutiny, which defines the basic strengths and weaknesses at the SBU level. When these core activities are done, they naturally lead to the formulation of the business strategy, supported by a set of multi-year broad and specific action programs. The corporation will evaluate these programs, allocate resources, make a formal commitment through the agreed upon budget figures, and define the performance measurements needed to carry out an intelligent strategic management control.

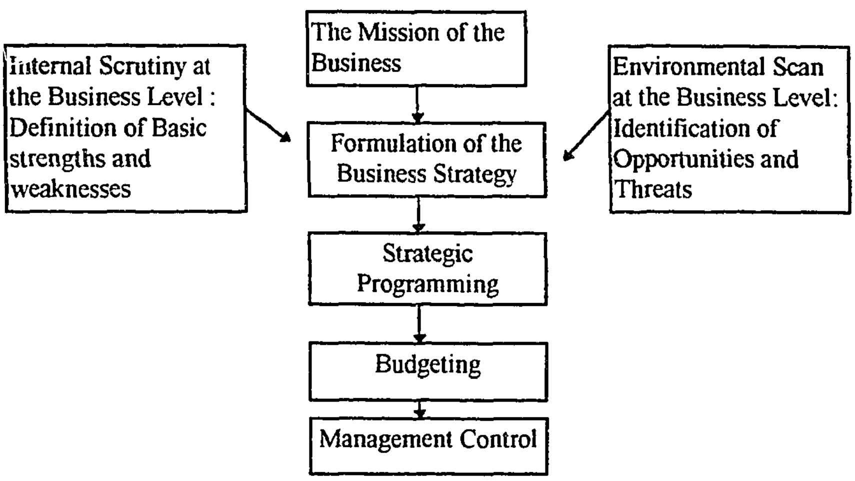

Figure 7. Business Strategy Flow

The Mission of the Business

The primary information that should be contained in the statement of mission is a clear definition of current and future expected business scope. The other important piece of information that should be contained in the mission statement of a business is the selection of a way to pursue a position of either leadership or sustainable competitive advantage.

Environmental Scan at the Business Level

Definition of the industry and competitive analysis. The four basic methodologies to perform this analysis include: (1) Michael Porter's framework for the structural analysis of industry, (2) an Environmental scan at the business level based on external factors analysis, (3) strategic group analysis, (4) The financial state analysis framework.

1. Porters Framework. Five forces that typically shape the industry include intensity of rivalry among competitors, threat of new entrants, threat of substitutes, bargaining power of buyers, and bargaining power of suppliers.

2. External Factors. The environmental scan at the business level attempts to identify the degree of attractiveness of the industry in which the business belongs.

3. Strategic Group Analysis. Identifying strategic groups in an industry. Porter suggests the following dimensions to identify differences in firm strategies within an industry: specialization, brand identification, push vs. pull marketing approach, channel selection, product quality, technological leadership, vertical integration, cost position, service, price policy, financial and operating leverage, relationship with parent company, relationship to home and host government.

4. The Financial State Analysis Framework. A method to gather quantitative intelligence at the level of the firm, based on financial statements analysis (i.e. balance sheets, income statements, and statement of changes in financial position).

Internal Scrutiny at the Business Level

The Value Chain. The focus of analysis of the value chain is SBU. The underlying principle is that all the tasks performed by a business organization can be classified into nine different broad categories. Five of them are so-called primary activities and the other four are labeled support activities.

Primary activities include innbound logistics--Receiving, storing, materials handling, warehousing, inventory control, vehicle scheduling, and returns to suppliers. Operations--Transforming inputs into final product form (e.g. machining, packaging, assembly, equipment maintenance, testing, printing and facilities operations). Outbound Logistics--Distributing the finished product (e.g. finished good warehousing, material handling, delivery vehicle operation, order processing and scheduling). Marketing and Sales--Induce and facilitate buyers to purchase the product (e.g. advertising, sales force, quoting channel selection, channel relations and pricing). Service--Maintain or enhance value of product after sale (e.g. installation, repair, training, parts supply and product adjustment).

Support activities include: Procurement--Purchasing of raw materials, supplies and other consumable items as well as assets. Technology Development--Know-how, procedures, and technological inputs needed in every value chain activity. Human Resource Management--Selection, promotion, and placement; appraisal, rewards; management development; and labor-employee relations. Firm Infrastructure--General management, planning finance, accounting, legal, government affairs, and quality management.

Internal Scrutiny is supported by the following tasks: Identification of most relevant SBU competitors; determination of critical success factors (CSF); development of a competitive profile for the SBU; by measuring the business strengths and weaknesses against each of the most relevant competitors; preparation of the summary assessment and identification of overall strengths and weaknesses associated with the SBU.

Critical Success Factors: Direction for Development

A tool that can be used for analyzing an organization's broad needs is the critical success factors method. In this method, interviews are conducted with top managers to determine the factors they believe are critical to the success of the organization. These are the limited number of areas in which results, if they are satisfactory, will ensure successful competitive performance for the organization. They are the few key areas where things must go right for the business to flourish.

An organization may have CSFs arising from various areas, including the structure of the industry, the competitive strategy of the firm, the company's current competitive positioning in the industry, actions of competitors, geographical location, the level of the general economy, government actions, the flow of raw materials, energy and other key resources, and time-related factors (Vargo and Hunt, 1996).

Planning Scenarios

The determinants in the development of a planning scenario are much like the determinants that must be considered when developing database for an EIS. The steps in developing a planning scenario include:

1. The first step in the planning cycle is to develop assumptions about the external environment.

2. Some key environmental factors used are macroeconomics--GNP; inflation rates: interest rates; political/regulatory--taxes; environmental: health and safety regulations: land access, liability exposure; industry--supply and demand: prices: import/export balances; plant operating rates.

3. Multiple Scenarios should be constructed. These should be plausible, internally consistent, pertinent to the business, and sufficiently different from one another to cover a wide range of possible outcomes.

Formulating Operating Company Strategies

Strategic plans may be formulated in relation to the base scenario only. While it is technically feasible to work out plans for alternative scenarios as well, the resulting workload is generally too great to justify the effort. In developing its plan an operating company must:

1. Objectively identify and analyze its shortcomings and successes in the past year, new threats and opportunities which may be developing, and the availability of resources.

2. Make a similar analysis of its competitors.

3. Devise specific action programs to improve its competitive position.

Another important planning tool is the use of key factors for success.
The key factors depend in particular on:

1. The type of industry (whether the cost structure is mainly fixed or variable, whether there are many or few competitors, and their strengths.

2. The degree of maturity of the industry (i.e. embryonic vs. growth vs. mature vs. declining.

3. The actions taken by the current and potential competitors themselves. (e.g. whether competitors are changing their own strategies, modifying the services they supply to the market).

The determination of the critical success factors and warning signs of a firm starts with the understanding of the overall strategy of the firm. Once critical success factors and warning signs are extracted from executives, they will be encoded into the IEISDW and set the constraints and rules for the intelligent agents.

CHAPTER V

A CONCEPTUAL MODEL FOR AN IEISDW

The Framework of an IEISDW

The framework of the IEISDW is described in figure 8. The framework of the IEISDW consists of a user interface, a facility coordinator, a data warehouse management system, the data warehouse facility, the operational data storage, and the various intelligent agents and mathematical models.

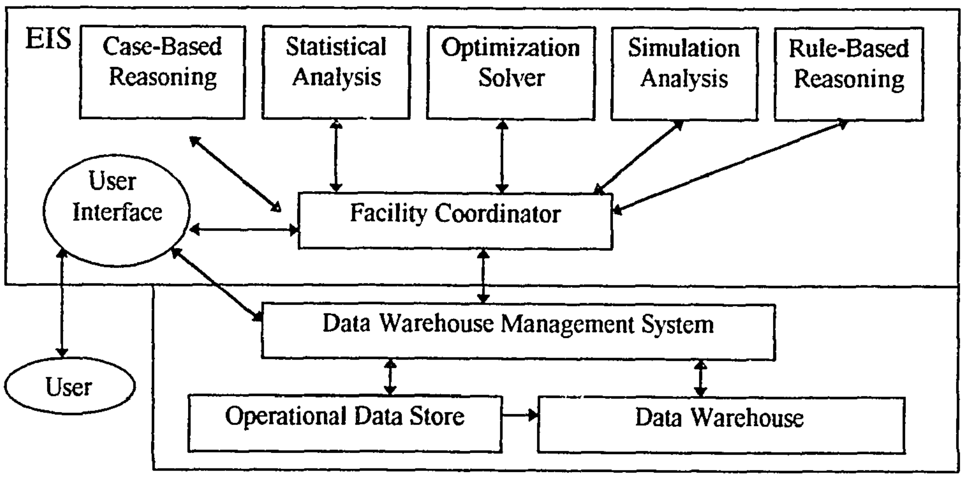

Figure 8. The Framework of the IEISDW

The functions of each component of the above framework are described as follows:

User Interface

A user interface is both the hardware and software that provide communication between the executive and the DWEIS.

Facility Coordinator

A Facility Coordinator is the management system that provides scheduling of tasks among intelligent facility units. This facility units include mechanisms for Case-Based Reasoning, Rule-based Reasoning, Statistical Modeling, Optimization Modeling, and Simulation Modeling.

Data Warehouse Management System

A Data Warehouse Management System is an intelligent interface between the Coordinator and the Data Warehouse. The Data Warehouse would forward the coordinator's requests to the right warehouse and deliver back the information that the coordinator requests. While the Coordinator keeps track of the order of intelligent analyses by intelligent units and mathematical models, the Data Warehouse Management System keeps track of the order of requesting and delivering of information to intelligent units and mathematical models.

Data Warehouse

A data warehouse can be a physical or logical data warehouse. The warehouse could include remote data marts which is electronically linked to the Data Warehouse Management System.

EIS systems are designed for solving unstructured problems. When a problem is input, the Facility Coordinator will handle the processing of the problem. Depending on the problem type, the Facility Coordinator can schedule intelligent units to process the problem in any order, and in any combination. For instance, simulation is usually combined with optimization to juggle between optimizing structured problems and simulating unstructured problems to achieve a possible best alternative.

Operational Data Storage (ODS)

The operational data storage is used almost exclusively for the operational information processing. Once the operational data is outdated it is transferred into the data warehouse. The operational data storage forms a foundation for operational information processing. The informational processing found in the operational data storage is for the clerical community and for managers making detailed, up to the second decisions.

Intelligent Agents & Mathematical Models

The intelligent agents and the mathematical models in the IEISDW system include mechanisms for Case-based reasoning, Rule-based reasoning, Statistical Modeling, Optimization Modeling, and Simulation Modeling. The following includes an explanation each of the intelligent agents and mathematical models in relation to the IEISDW.

1. Case-Based Reasoning Mechanism. Case-Based Reasoning has been verified as an approach to intelligent information retrieval (Daniels and Rissland 1995). For instance, inputs are accepted in a standard frame-based representation of a problem case, and texts of relevant cases are output. Then, the query is automatically formed by submitting , in text form, a set of highly relevant cases, based on CBR analysis. And a solution is proposed. After proposing the solution, the current situation can be stored for the next case-based reasoning analyses (Riesbeck and Schank 1989).

2. Rule-Based Reasoning Mechanism. A rule-based reasoning system uses rules to explain and generalize old cases in a case base. Rule-based reasoning is deductive knowledge that can be retrieved from the knowledge of a domain expert.

3. Statistical Modeling Mechanism. Statistical Analysis is a facility that includes statistical

functions to provide best fit models, regression analyses, statistical hypothesis testings of the current situation. The Case-based Reasoning unit would determine what statistical models to use and for what purposes.

4. Optimization Modeling Mechanism. Optimization Modeling is a facility that proposes the best or optimized solution by testing all other normative (verifiable) alternatives. Since strategic decisions most often deal with uncertainty and risk, optimization allows executives to evaluate situations which involve a finite and not too many alternatives. Optimization techniques include decision trees, linear and other mathematical programmings.

5. Simulation Modeling Mechanism. Simulation is a facility that imitates reality of the proposed solution by conducting experiments. Simulation is usually applied to non-routine or descriptive-type problems to obtain an estimate of the overall effect of certain actions. Once these characteristics are known, the best among several alternatives can be selected. However, this best alternative is not the best alternative among all alternatives. It is only the best alternative deducted from several experiments. Major simulation techniques are probabilistic simulation, time dependent versus time independent simulation and visual simulation.

CHAPTER VI

THE PROBLEM SOLVING PROCESS OF IEISDW

Figure 9 gives an illustration of the problem solving process of the IEISDW. Each of the steps of the process are explained below.

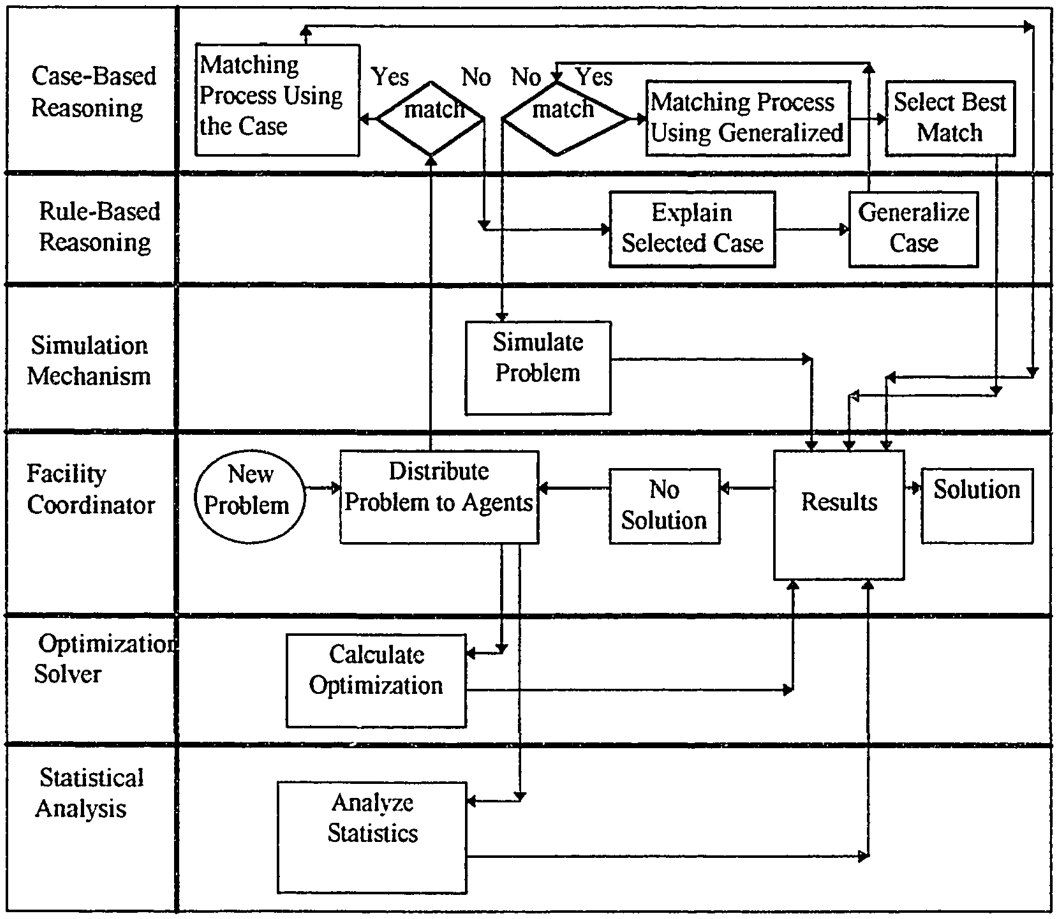

Figure 9. Procedures of Problem Solving in a IEISDW

When facing an executive question, an executive will first try to solve a problem intuitively based on their own knowledge and past experiences. An executive may consult the IEISDW when the problem domain becomes too complex and/or is unknown by an executive.

An executive may query the IEISDW for important decision making questions through the user-interface. The user interface includes a computer terminal which is connected via a LAN network to the data warehouse. If the executive knows which problem solving mechanism to use to solve a given problem, then he/she may access that mechanism and solve the problem. If the executive has no idea of how to solve the problem, then the executive enters the question into the system. The facility coordinator will then download the data necessary to answer the question form the data warehouse to the computer terminal. Then the Facility Coordinator will first consult the Case-based reasoning mechanism for a match to the given case. If there was not a previous situation that could match the current situation then a rule-based reasoning mechanism would be engaged to determine a solution. The rule-based reasoning mechanism would generalize the current problem and then try to match this problem with any previous problems. If the rule-based reasoning mechanism could not determine a solution, then a statistical modeling mechanism would be engaged to determine possible solutions to a given problem. The statistical modeling mechanism uses regression analysis to determine a pattern of similar previous cases that may match the current case scenario. In many statistical investigations, the main goal is to establish relationships which make it possible to predict one or more variables in terms of others. If there are not enough factors to determine a solution based on regression analysis, then an optimization modeling mechanism would be engaged. Optimization analysis could be used in a problem domain where many factors are unknown. If the problem domain lacks sufficient determinants, an optimization modeling mechanism would be engaged to solve a given problem. An optimization model would try to define the optimal solution to a problem given a set of variables. If the problem domain is too large for numerical optimization techniques then a simulation mechanism would be engaged to simulate a possible solution. A simulation modeling and the optimization modeling mechanisms are tools of the EIS which could be used to simulate "what if" scenarios asked by executives. If an executive does not receive a solution to his/her problem after the request has passed through all the intelligent agents and mathematical models then the IEISDW will continue to submit the problem to the intelligent agents and mathematical models until there are possible solutions or a request has passed through the system ten times. The interative processing allows for the intelligent agents and mathematical models to work collaborative to solve a problem. If after ten times passing through the loop of intelligent agents and mathematical models the system still has no solution, then the system will output a "no solution" response.

CHAPTER 7

AN ILLUSTRATION OF AN IEISDW CAPABILITIES

A Internet Service Provider Selector is implemented using IEISDW to decide which internet service provider would fit an organization's needs in respect to their information needs, budget constraints and ease of use. In general, Internet Service Providers have various backgrounds in terms of the information provided, the cost of the service and the simplicity of use. The decision as to what Internet Service Provider to use is a very important one because in order to keep up with current technology and the information explosion a company must choose an Internet Service Provider that is competitive and fits their specific needs.

In recent years, Internet Service Providers have been successfully providing information for executive managers in a growing number of organizations. Internet Service Providers have become very popular among organizations and the features and amount of Internet Service Providers has increased dramatically. Internet Service Providers have many different features and it has become very difficult for organizations to determine which Internet Service Provider will best fit the needs and constraints of the individual organization. There is not a current system that organizations can access to determine which is the best Internet Service Provider for them.

I propose the use of the IEISDW to determine the best Internet Service Provider for individual corporations.

Internet Service Providers have many of the same features and some features that are unique to each provider. The following illustrates the types of Internet Service Providers and well as the types of services provided and some pricing structures.

Types of Service Providers

Users interested in accessing the Internet have several options including commercial On-Line services, national service providers and local providers. The type of provider selected will depend on how the individual plans to use the Internet and their location. On-Line services provide the easiest access at competitive prices, but performance is slow. National providers fall into two categories consisting of young companies concerned primarily with the Internet market and computer and communications colossuses attempting to move into the new market. Local providers cover a spectrum from small local providers to major regional operations. Local providers provide the greatest flexibility, local information and competitive pricing. The major commercial On-Line services and national service providers are described with details about pricing, special features, software and connections (Norr 1996).

Commercial On-Line Services

Commercial On-Line services make accessing the Internet easiest, they offer special features such as parental controls and their prices are competitive. But they deliver the slowest performance.

National Service Providers

Two types of companies dominate the national service provider arena. One is a group of relatively young companies little known outside the Internet market. They got their start providing Internet access for business customers, gradually built up nationwide networks, and have in recent years capitalized on the growing popular interest in cyberspace. Companies matching this profile include UUNET Technologies, Performance Systems International (PSI), and Net-Com On-Line. The other group consists of computer and communication giants including IBM and MCI that are attempting to grab the share of this new market on the strength of their existing infrastructures, expertise and customer loyalty. AT & T is about to launch an access program too.

National providers are notable for their high performance systems and low flat rate plans. They target more experienced users and business customers, although, they have been recently reaching out to wider audiences and offering software that is easy to use. And while On-Line services such as CompuServe and America On-Line add low-cost, Internet-only services, look for national provider prices to continue to drop.

Local Providers

It doesn't take much to become an Internet service provider, and entrepreneurs are setting up shop everywhere. There are now more than 1,000 Internet service providers across the United States. Some have matured into major regional operations; others are strictly local, even neighborhood based businesses. And despite mounting competition from larger and more established players, new service providers are appearing all the time.

Local and regional providers are the most diverse group, offering competitive pricing, flexibility, local information that may not be available anywhere else in cyberspace, and a personal touch that larger services cannot match.

Commercial On-Line services, national providers, and local providers all give you quick and easy access to the Net. As the competition among them heats up, you can expect even lower rates, easier connections, and more service options tailored to specific needs.

Major Service Categories and Pricing

Although there are many categories of services provided by Internet Service Providers, some of the most common include the following: Finance, Shopping, Entertainment, Home and Leisure, Travel, Professional forums, Sports, Reference and Education, News and Weather and, Media.

The pricing for Internet Service Providers ranges significantly based upon the time a user spends on the On-Line system. CompuServe, for example, charges $25 per month with three hours of free access and any additional hours being $1.95 an hour. Dow Jones charges a $19.95 annual fee, a start up fee of $29.95 a month for three hours of free usage and $1.50 per 1,000 characters. The demand for Dow Jones is so great that they even offer discount rates if you spend $1,000 to $15,000+ per month on its services.

Utilizing the IEISDW

New problems (Internet Service Providers) are collected with various background information in the format of a frame. In the case base, cases of current Internet Service Providers are recorded. Each case contains the information about the Internet Service Provider's pricing, type of provider, type of services, and ease of use.

This information is stored in three types of slots: Critical Slots, Common Slots, and Decision Slots. Each Critical Slot stores the value of an attribute that is necessary for a successful match. A Common Slot contains the value of an attribute thought not necessary, which contributes to a better quality match. Each Common Slot has a weight assigned to that attribute indicating its importance. The decision made in the past associated with each Internet Service Provider is solved in the Decision Slot which contains the solution plan.

In this example, we suppose that services provided and pricing are essential features to the final decision, while, service provider type and ease of use are less essential. The services provided and pricing information about an Internet Service Provider are stored in the Critical Slots while service provider type and ease of use is stored in the Common Slots. We further suppose that the attribute of ease of use has a weight(wt.) of 2, and the service provider type has a weight of 1, which suggests that ease of use be more important than service provider type when selecting a match. The match succeeds if the values in all of its Critical Slots are identical to their counterparts associated with this new case. The values stored in the Common Slots are not necessary for the match but are useful to decide which match is "better'' than others. Consider the following examples in which we have found two matching old cases for the new case under consideration:

Old case 1:

Pricing	(Critical Slot):1000-5000
Services	(Critical Slot):Financial
Ease of Use	(Common Slot,wt.2):Easy
Service Provider Type	(Common Slot,wt.1):National
Service Provider	(Decision Slot):Good

Old case 2:

Pricing	(Critical Slot):1000-5000
Services	(Critical Slot):Financial
Ease of Use	(Common Slot,wt.2):Difficult
Service Provider Type	(Common Slot,wt.1):Local
Service Provider	(Decision Slot):Fair

New Case:

Old case 1:

Pricing	(Critical Slot):1000-5000
Services	(Critical Slot):Financial
Ease of Use	(Common Slot,wt.2):Easy
Service Provider Type	(Common Slot,wt.1):Local
Service Provider	(Decision Slot):Good

In this example, both old cases have acceptable values of Critical Slots as the ones of the new slots. Considering the values of the Common Slots in both old cases, we can make further comparisons between them. That is, the first old case matches the new case in the values of Ease of Use which has a weight of 2, while the second old case matches the new one in the value of Service Provider Type, with a less weight of 1. Therefore, the first old case is a better match with respect to this new case. Besides Common Slots and Critical Slots, there are Decision Slots. Decision Slots contain the decisions or plans that the new case adapts. The decision made in the old case that is determined to be the best match will be adopted for the new case. In our example, a "good" decision that was made in the first old case will be adopted as this old case is a better match than the second one.

Domain Theory

Domain Theory is used to explain or interpret the values associated with each slot in the old case frames so as to generalize the old cases. With this explanation, old cases can be generalized and extended to cover a larger problem domain. Therefore, the generalized old cases can cover larger problem spaces depending on the explanation of the domain theory and become more powerful in terms of problem solving capabilities. The Domain Theory is represented by a set of production rules which reflect the policy of explanation. In this example, the policies include:

Policy 1: Any pricing that is from less than $1,000 a month will be treated as a "good" selection, a "fair" selection if less than $5,000 per month but greater than $1,000, and a "poor" selection if above $5,000 per month. (this policy is reflected by rule 1, 2 and 3).

Rule 1

Pricing (X, Y) & $Y \leq 1{,}000 \rightarrow$ pricing(X, good)

Rule 2

Pricing (X, Y) & $1{,}000 < Y \leq 5{,}000 \rightarrow$ pricing(X, fair)

Rule 3

Pricing (X, Y) & $Y > 5{,}000 \rightarrow$ pricing(X, poor)

Policy 2: If the services provided has a rank of 4 or higher, it is a "good" selection, greater than or equal to 2 and less than or equal to 3, it is a "fair" selection, less than or equal to 1, a "poor" selection (this policy is reflected by rules 4, 5 and 6)..

Rule 4

Services(X, Y) & Rank (Y, Z) & $Z \geq 4 \rightarrow$ Services (X, good)

Rule 5

Services(X, Y) & Rank (Y, Z) & $2 \leq Z \leq 3 \rightarrow$ Services (X, fair)

Rule 6

Services(X, Y) & Rank (Y, Z) & $Z \leq 1 \rightarrow$ Services (X, poor)

where Services(X, Y) means the Services of company X is Y, and Rank(Y, Z) means Z is the Rank of Y. The rank of the Services are as follows:

Frame Name	**Rank**
Financial	5
Travel	4
Professional Forums	3
Media	2
Reference & Education	1

Policy 3: The "Ease of Use" of the Internet Service Provider. The Ease of use is divided into three categories: "Very Easy," "Easy," "Difficult," If any two of the same types of Internet Service Providers fall into

the same category, they can be treated equally; therefore both have the same Ease of Use (Rule 7, 8 and 9 are used to describe policy 3).

Rule 7

Ease of Use(X, Y) Very Easy → Ease of Use(X, good)

Rule 8

Ease of Use(X, Y) Easy → Ease of Use(X, fair)

Rule 9

Ease of Use(X, Y) Difficult → Ease of Use(X, poor)

Policy 4: The "Service Provider Type" of the Internet Service Provider. The Service Provider Type is divided into three categories: "Commercial On-Line," "National," "Local." If any two of the same types of Internet Service Providers fall into the same category, they can be treated equally. (Rule 10, 11 and 12 are used to describe policy 4).

Rule 10

Service Provider Type(X, Y) Commercial On-Line → Service Provider Type(X, good)

Rule 11

Service Provider Type(X, Y) National → Service Provider Type(X, fair)

Rule 12

Service Provider Type(X, Y) Local → Service Provider Type(X, poor)

Explanation of Old Cases

The explanation of old cases involves iterations. Each iteration contains three steps: namely, (1) Matching, (2) Exploration, and 3. Instanciation.

Matching Process: This process involves matching the Variable-symbols in premises of the explanation rules against the values of attributes in a frame if the Function-symbol in the premises of explanation rules is the same as the attribute in the frame. A successful match will result in execution of explanation rules.

Exploration Process: The Exploration process finds the extended frames to explain values associated with the slots of a certain frame. Explored frames are frames which are explored by the Exploration process.

Instanciation process: This process simply associates instance with the explained frames which include the original and the extended ones. In this process, all Variable-symbols are updated based on the actual values. Consider the following example:

Explanation of the first policy: Consider old case 1, which is described as follows:

Iteration 1

1. Matching process: The first premise of rule 1, "Pricing (X, Y)" is matched with the case with the slot "Pricing" and X is matched with ABC company, while the value of "Pricing" slot "1000-5000" is matched with Y.

Frame name: Old case 1: ABC company → X

Pricing	(Critical Slot):1000-5000 → Y
Services	(Critical Slot):Financial
Ease of Use	(Common Slot,wt.2):Easy
Service Provider Type	(Common Slot,wt.1):National
Service Provider	(Decision Slot):Good

2. Exploration process: Because no extended frame is associated with the value "1000-5000," no exploration process is necessary.

3. Instantiation process.

ABC company = X

$1000 < Y \leq 5,000$

Pricing	(Critical Slot):Good
Services	(Critical Slot):Financial
Ease of Use	(Common Slot,wt.2):Easy
Service Provider Type	(Common Slot,wt.1):National
Service Provider	(Decision Slot):Good

<u>Explanation of Second Policy:</u> This case can be further generalized by the second policy:

Pricing	(Critical Slot):Good
Services	(Critical Slot):Financial
Ease of Use	(Common Slot,wt.2):Easy
Service Provider Type	(Common Slot,wt.1):National
Service Provider	(Decision Slot):Good

1. <u>Matching process:</u> The first premise of rule 4 "Services(X, Y)," is matched with the case with the slot "Services" and X is matched with ABC company while the value of rental amount slot "Financial" is matched with Y.

Frame name: ABC company → X

Pricing	(Critical Slot):Good
Services	(Critical Slot):Financial → Y
Ease of Use	(Common Slot,wt.2):Easy
Service Provider Type	(Common Slot,wt.1):National
Service Provider	(Decision Slot):Good

2. <u>Exploration process</u>: The value of the slot "Services" is "Financial" which can be explored by retrieving an extended frame to explain "Services."

Frame name: Financial → Y

rank: --

3. <u>Instantiation process:</u>

ABC company = X

Y = Financial

<u>Iteration 2</u>

1. <u>Matching process:</u> The second premise of rule 4 "Rank(X, Y)," is matched with the extended frame "Financial" and the value "5" associated with the slot "Rank" is instantiated with Z.

Frame name: ABC company → X

Pricing	(Critical Slot):Good
Services	(Critical Slot):Financial → Y
Ease of Use	(Common Slot,wt.2):Easy
Service Provider Type	(Common Slot,wt.1):National
Service Provider	(Decision Slot):Good

Frame name: Financial → Y

Rank → Z

2. <u>Exploration process</u>: Because the value "Rank" is instantiated with Z, it will further be explored by looking for the frame with the name "Rank."

Frame name: Rank → Z   Rank: 5

3. <u>Instantiation process:</u>

ABC company = X

Financial = Y

Rank = Z

Frame name: ABC company

Services Provided	(Critical Slot):Good
Pricing	(Critical Slot):Good
Ease of Use	(Common Slot,wt.2):Easy
Service Provider Type	(Common Slot,wt.1):National
Service Provider	(Decision Slot):Good

Explanation of the Third Policy: This case can be further generalized by the third policy:

Frame name: ABC company → X

Services Provided	(Critical Slot):Good
Pricing	(Critical Slot):Good
Ease of Use	(Common Slot,wt.2):Easy
Service Provider Type	(Common Slot,wt.1):National
Service Provider	(Decision Slot):Good

1. Matching Process. The first premise of rule 7, Ease of Use(X, Y), is matched with the case with the slot "Ease of Use" and X is matched with "ABC company," while the value of Ease of Use slot "Easy" is matched with Y.

Frame name: ABC company → X

Pricing	(Critical Slot):Good
Services	(Critical Slot):Good
Ease of Use	(Common Slot,wt.2):Easy → Y
Service Provider Type	(Common Slot,wt.1):National
Service Provider	(Decision Slot):Good

2. Exploration Process. Because no extended frame is associated with "Easy," no exploration process is necessary.

<u>Explanation of the Fourth Policy</u>: This case can be further generalized by the fourth policy:

Frame name: ABC company → X

Services Provided	(Critical Slot):Good
Pricing	(Critical Slot):Good
Ease of Use	(Common Slot,wt.2):Good
Service Provider Type	(Common Slot,wt.1):National → Y
Service Provider	(Decision Slot):Good

1. Matching Process. The first premise of rule 10, Service Provider Type(X, Y), is matched with the case with the slot "Service Provider Type" and X is matched with "ABC company," while the value of Service Provider Type slot "Easy" is matched with Y.

Frame name: ABC company → X

Pricing	(Critical Slot):Good
Services	(Critical Slot):Good
Ease of Use	(Common Slot,wt.2):Good
Service Provider Type	(Common Slot,wt.1):National → Y
Service Provider	(Decision Slot):Good

2. Exploration Process. Because no extended frame is associated with "Easy," no exploration process is necessary.

In the integrated system, cases and policies are stored separately in the knowledge base. This allows the system to be easily maintained or modified whenever it is needed.

So far, to this point in the example, the question of which Internet Service Provider for an organization to choose has passed through the case-based reasoning mechanism and the rule-based reasoning mechanism. In the case-based reasoning mechanism, old cases were matched up with the current case.

When no matches occurred, the cases were generalized and again an attempt at matching the generalized case to the old cases.

If there is still not a match using the rule-based reasoning mechanism, regression analysis would be used to plot the characteristics of previous old cases and try to match the linear regression to the current case.

If the previous cases do not match any cases in the regression analysis, then the optimization modeling mechanism would attempt to optimize a solution based on all the characteristics given in the new case. Therefore, characteristics of Pricing, Services Provided, Ease of Use, and Service Provider Type would all be given a weighted value and the optimization modeling mechanism would calculate an amount that gives an optimal solution.

If the problem domain of the case becomes too complicated because their are too many characteristics to deal with, then the simulation modeling mechanism would simulate a possible solution based on the characteristics.

CHAPTER VIII

CONCLUSION

In this paper, an Intelligent EIS with a Data Warehouse Facility (IEISDW) was proposed to determine solutions to executive problems. This system answers executive's questions through the coordination of an EIS, intelligent agents, and a data warehouse.

In Chapter 2 of this paper, a literature review of current research in EIS, intelligent agent and data warehouse technologies is presented. A discussion of distributed artificial intelligence is presented in Chapter 3. Chapter 4 discusses the strategic management of executive decisions. The architecture and design of IEISDW is proposed in Chapter 5 while Chapter 6 demonstrates the problem solving process of IEISDW. Finally, Chapter 7 gives an illustration of IEISDW capabilities.

This paper illustrates how it is possible for successful EIS of the future to be built around AI components. Additionally, it illustrates how a data warehouse facility can be utilized by an EIS to examine data company-wide.

Data warehouse technology can be integrated with EIS to provide strategic decision making capabilities. This concept could be enhanced in the future by several methods such as automating intelligent reasoning to search the data warehouse for potential opportunities among similar markets and suggest them to executives, or to suggest business problems which need to be corrected or improved.

In this thesis, I proposed an EIS framework with intelligent distributed information processing agents. The case-based reasoning agent and the rule-based reasoning agent of EIS enable the processing of more complicated information by the cooperative efforts of each agents. In fact, many executive tasks contain various aspects of problem domains which cannot be supported by a single data retrieving and processing mechanism. In other words, current EIS retrieve data directly from the database and present it with no automated data processing. By employing the proposed framework, a more intelligent EIS can be constructed by having a set of agents working cooperatively.

This research is expected to follow many directions in the near future, including:

Detailed experiments will be performed to evaluate the IEISDW framework in terms of enhanced problem solving performance and improved learning capabilities. A variety of case evaluation methodologies will be tested such as neural networks and multi-dimensional distance methods.

An implementation will be conducted by using object oriented languages so that each case can be used as an individual object in order to reduce implementation costs.

Real world applications will be used to test the system such as the problem domains of production facility reallocation, strategic investment evaluation, portfolio selection management, and strategic marketing segmentation evaluation.

REFERENCES

Baer, Tony. 1995. IBM discloses warehouse plans: in the backroom, 'old iron' still has a home. Software Magazine, 1 July, 24.

Bratko, I. 1986. PROLOG Programming for Artificial Intelligence. Reading, MA: Addison-Wesley.

Buchanan, B.G., and E. H. Shortcliffe. 1984. Rule-based Expert Systems. Reading, M.A.: Addison-Wesley.

Bull, Katherine. 1995. Warehousing isn't enough. Information Week, 26 June, 47.

Chi, Robert T. and Efraim Turban. 1995. Distributed Intelligent Executive Information Systems. Decision Support Systems 14 (Winter): 117-130.

Chi, Tsing-Hwa. 1991. MEDSS: A Distributed Decision Support System (DDSS) For Model Building. Ph.D. diss., University of Texas at Austin.

Chi, T. H. Robert, Minder Chen and Melody Y. Kiang. 1993. A Generalized Case Based Reasoning System for Portfolio Management. Expert Systems with Applications 6: 67-76.

Chi, Robert T. H. and Steven J. Gramme. 1995. Case-based Executive Information Systems with Deductive Facilities. AIS Conference. Pittsburgh.

Cole, Barb. 1995. Data Warehouse Shortcuts due from Sybase. Network World, 12 October, 40.

Daniels, J. J. and E.L. Rissland. 1995. Using CBR to drive IR. Proceedings of the National Conference on AI, 1995, 34-45. Menlo Park, California, USA.

Demarest, Marc. 1994. Building the Data Mart. DBMS 7(July): 44.

Eskow, Dennis. 1995. Winning regimen; GE Capital's warehouse helps it gobble up smaller competitors. PC Week, 19 June, 1.

Fox, Mark. 1981. An Organizational View of Distributed Systems. IEEE Transactions on Systems, Man, and Cybernetics 12 (January): 24-27.

Freund, John E. 1988. Regression Analysis. Modern Elementary Statistics. Englewood Cliffs, N.J.: Prentice-Hall.

Hackathorn, Richard. 1995. Data Warehousing Energizes Your Enterprise. Datamation 43 (February): 95.

Hax, Arnoldo, C. and Nicolas S. Majluf. 1991. The Strategy Concept and Process. Englewood Cliffs, N.J.: Prentice-Hall.

Hinrichs, T. 1989. Strategies for adaptation and recovery in a design problem solver. In Proceedings of the DARPA Case-Based Reasoning Workshop. San Mateo, CA: Morgan-Kaufmann: 115-118.

Hunt, John. 1995. MACE: A System for Construction of Functional Models Using Case-based Reasoning. Expert Systems with Applications. 9: 347-360.

Inmon, W. H. and Richard D. Hackathorn. 1994. Using the Data Warehouse. New York, N.Y.: John Wiley & Sons, Inc.: 1-25.

Janah, Monua. 1995. Data lessons learned: Stanford data warehouse improves operations. InfoWorld, 31 July,1.

Ketler, Karen. 1993. Case-based Reasoning: An Introduction. Expert Systems Applications 6: 3-8.

Kolodner, Janet L. and Robert L. Simpson. 1989. The Mediator: Analysis of an Early Case-Based Problem Solver. Cognition Science 13: 507-549.

Koton, P., and M.P. Chase 1989. Knowledge representation in a case-based reasoning system: Defaults and exceptions. In Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning, 57-63. Hawaii.

Maurice, Frank. 1994. A drill down analysis of Multidimensional databases. DBMS (July): 60.

McCartney, R. 1993. Episodic Cases and Real-time Performance in a Case-based Plan System. Expert Systems With Applications. 6: 9-22.

McFadden, Fred R. 1996. Data Warehouse for EIS: Some Issues and Impacts. Proceedings of the 29th Annual Hawaii international Conference on System Sciences 1996: 120.

Nilsson, Nils. 1980. Principles of Artificial Intelligence. Palo Alto, CA: Tioga Publishing Co..

Norr, Henry. 1996. The right way to the I-way. MacWeek, 12 January, 123.

Richman, Dan. 1995. Data Warehousing Challenge, Startup aims to tackle setup and usage problems. Information Week, 21 August, 21.

Riesbeck, Christopher K. and Roger C. Schank. 1989. Inside Case-Based Reasoning. Hillsdale, New Jersey: Lawrence Relbaum Associates.

Saylor, M.J., and S.K. Bansal. 1995. Open Systems Decision Support: An Overview. Data Management Review 5 (January): 20-24.

Shaw, M. and M. Fox. 1991. Distributed Artificial Intelligence for Group Decision Support: Integration of Problem Solving, Coordination, and Learning. Working Paper.

Simon, H. A. 1957. Models of Man. New York: Wiley.

Smith, R.G. and R. Davis. 1981. Frameworks for Cooperation In Distributed Problem Solving. IEEE Transactions on Systems Man and Cybernetics. 11: 1.

Stearns, Tom. 1995. Prophesy! Tools help build virtual networks. LAN Times, 6 November, 118.

Storn, Rainer. 1995. Constrained Optimization. Dr. Dobb's Journal 20 (May): 119.

Subramanian, Ashok, L. Douglas Smith, Anthony C. Nelson, James F. Campbell, and David A. Bird. Proceedings of the 29th Hawaii International Conference On System Sciences 1996: 54.

Sycara, K. 1988. Using Case-based Reasoning for Plan Adaptation and Repair. Proceeding DARPA Workshop on Case-based Reasoning: 425-434.

Thompson, George A. 1994. The house that data built: data warehousing gives corporations a competitive advantage with better data analysis. HP Professional 8 (December): 36.

Turban, Efraim. 1995. Modeling and Model Management. Decision Support and Expert Systems: Management Support Systems: 161-74.

Tyo, Jay. 1995. Simulation Modeling Tools. Information Week, 7 July, 60.

Vargo, John and Ray Hunt. 1996. Telecommunications in Business. Chicago, Il.: Irwin Book Team: 238-41.

Warren, Liz. 1994. Super stores (data warehousing). Computer Weekly, 8 September, 26.

Watson, Hugh J., R. K. Rainer and C. Koh. 1991. Executive Information Systems: A Framework for Development and a Survey of Current Practices. Management Science 33 (March): 99.