



MIT LIBRARIES DUPL 2



3 9080 00579029 7

BASEMENT



HD28
.M414

No. 1791
86

DEWEY

DEC 21 1989
LIBRARIES

WORKING PAPER
ALFRED P. SLOAN SCHOOL OF MANAGEMENT

INTELLIGENT MANAGEMENT SYSTEMS:
DESIGN AND IMPLEMENTATION

BY

ZENON S. ZANNETOS

July 1986

1791-86

MASSACHUSETTS
INSTITUTE OF TECHNOLOGY
50 MEMORIAL DRIVE
CAMBRIDGE, MASSACHUSETTS 02139

INTELLIGENT MANAGEMENT SYSTEMS:
DESIGN AND IMPLEMENTATION

BY

ZENON S. ZANNETOS

July 1986

1791-86

MIT LIBRARIES
DEC 21 1989
RECEIVED



Massachusetts Institute of Technology
Sloan School of Management
50 Memorial Drive
Cambridge, Massachusetts, 02139

This paper was written as the Keynote Speech
for Economics and Artificial Intelligence,
the 1st International Conference
September 2-4, 1986
Aix-en-Provence, France

INTELLIGENT MANAGEMENT SYSTEMS: DESIGN AND IMPLEMENTATION

by

Zenon S. Zannetos

Professor of Management
Sloan School of Management
Building E52-476
50 Memorial Drive
Cambridge, Massachusetts 02139

(617) 253 - 7164

INTRODUCTION

During the last six years we have seen the evolution of a major thrust toward the development of "expert systems." These are systems whose knowledge base is developed and introduced into computers by "experts." As a result the answers provided by the system represent the opinion of the expert or experts who designed it.

Professor Simon made a prophesy about twenty years ago, that computers will simulate human thinking and that by 1985 machines will substitute humans "for any and all human functions in organizations." (Simon 1965:30) This was a very bold statement and while it has not become a reality as of now, and there is a doubt in my mind as to whether it will ever be so "for any and all human functions," the movement is in the right direction. Several firms today are staking their future on artificial intelligence and expert systems and many others are preparing to introduce these wonders of information technology into their operations as soon as these become available.

So great is the momentum generated by information technology, that many economists and businessmen view this area as the source of new growth in productivity in the United States and the rest of the world. Already certain states, such as Massachusetts, have avoided unemployment because of the birth of this industry. Many new "start ups" have sprung up to develop and manufacture machines and firmware that process lists and symbolic statements, preferably in parallel.

In the area of software, the excitement has been even greater, sending venture capitalists in search of new opportunities for investment to take the slack left by biotechnology. All in all, there is a lot of excitement.

In this presentation I will attempt to: (a) examine some of the fundamental premises upon which the value of expert systems rests; (b) briefly draw a distinction between "expert systems" and "intelligent systems" and propose that it is the latter, that is, systems which serve as an extension of human capability that the managers need; (c) describe some of the attributes of intelligent systems and (d) carry out some of the most important consequences of the introduction of intelligent systems for the functions of the executive, and, for the organization structure of the firm and of macroeconomic subentities.

EXPERT SYSTEMS

We have briefly defined expert systems as systems which emulate the decision making process and protocol of the expert. As such their value is in the availability of this expert knowledge to the user. They can serve, if properly designed, to advise and educate the user on:

- A. Problem identification methodology and process, assuming, of course, that the manager knows that he/she has a problem and needs to ask for help.

- B. Definition of an appropriate decision-making context to facilitate a solution of the identified problem, which may include:
 1. A mental model of the complex system the manager wishes to understand and "control"

 2. A description/definition of:
 - a. The critical elements of the system,

 - b. The probable states of the critical elements,

 - c. The unique events that may result from the combination of these elements in their various probabilistic states,

- d. The unique events that may cause or alter the state of these elements in their various probabilistic states,
 - e. The specific or probable relationships (temporal; spatial; cause and effect) that may exist between the various elements in their probable states, and
 - f. The specific or probable relationships that may exist between the "system" and the problem-solving objectives of the decision maker.
- C. Identification of the information needed for the resolution of the complexity inherent in the problem
- D. Choice of alternative models and methodologies which experts use to generate the information necessary for obtaining alternative solutions to the identified problem
- E. Development of the criteria for evaluating the alternative solutions
- F. Definition of the data needed for driving the alternative models and methodologies as identified
- G. Collection of the data needed and offering of advice on how to apply them to test alternative solutions to the problem. The above activity, which in a sense is a hypothesis testing operation, may need to be carried out

at the subsystem level. As we all know the elements of a system as well as the subsystems, described under B (a) through (f), are systems in their own right. Therefore, information (data, models/methodologies) may need to be generated to identify and test the significance of elements, states, events, relationships, and the consequences of all these. Often the complexity of the system, within which the problem is embedded, is such that simplification and elimination of elements, states, events and relationships may be necessary.

- H. Iterations which may be necessary for refinement of definitions, systems, models, information, data and measurements, and choice of solution.

PREMISES UNDERLYING THE USEFULNESS
OF EXPERT SYSTEMS

If we look at all the aspects of the uses that executives can make of expert systems, several critical underlying assumptions emerge. These assumptions refer to the role of the executive, the role of the expert, and the role of information. I will now deal briefly with each one.

- A. The role of the executive: As Simon (1960;1965) pointed out, executives deal with two types of decisions; (1) structured and (2) unstructured. The boundaries between these decisions change as people learn from experience and incorporate into models and systems whatever intelligence they have gathered. In fact, not only the boundaries change given the general context of the decision-making occasion, but the context itself should change and also the nature and scope of the decisions as knowledge

is acquired. We can say, therefore, that managers are pattern recognizers who attempt to manage intelligently the complexity manifested inside their organization, and inherent in the outside world with which they interface. This complexity of course is greater; the greater (a) the change initiated within an organization and in the outside world into which it is embedded, and (b) the number of complementary activities the executive perceives and attempts to manage. That is why in the high technology field, with the dominant role of R&D, the internal task of the manager is more difficult than that of his/her counterpart in the commodity area.

So the manager given his/her perceptive capabilities defines the context of the managerial system, which affects the definition of the relevant system and the choice of the critical elements, their states, the events that relate them and result from their interaction, and the constraints that are assumed to govern their behavior, i.e. the type of relationships that are hypothesized to govern their interaction within the system and relate the behavior of the system to the objectives of the organization. The executive in one way or another, rightly or wrongly, consciously or unconsciously, bounds the context, always walking the thin line between the very meaningful but untractable generality and the limiting specificity which is necessary for analytical solutions.

As a pattern recognizer the manager is the most qualified person to describe and help define the context which generates meaning out of data, models and analytical techniques. Whether this definition of the decision-making context is the best or the most enabling for creativity under universally applicable circumstances, is another story. It must be stressed,

however, that if the problem solvers do not correctly identify and perceive the executive's context, they may succeed in elegantly solving the wrong problem, at least as far as the manager is concerned. And even if the context of the problem solvers is correct, as long as the executive's context is different, the executive will not derive meaning out of the data and the models and methodologies that generated the data. In other words, the subjective context of the executive represents the world as he/she sees it and this view affects both the choice of the "objective" models and analytical techniques as well as the meaning and use of the data generated. No matter how much information the expert derives out of data -- the output of his/her "rational" or "objective" models -- through the use of his/her subjective context, the executive will most likely be unimpressed. Unless the subjective context of the executive is changed to be consistent with that of the expert, the way the latter looks at the world, uses the data, and derives information out of them is more or less irrelevant. The "expert information" will probably be meaningless to the executive because the proper associative context is missing. If the mental model is different between two people there will be no communication and no transfer of knowledge.

B. The Role of the Expert

As indicated above, an expert system attempts to provide the user with the latest theory and methodology applicable to the respective decision-making situation. The decision-maker, who uses an expert system has, in effect the expert as a personal advisor at his/her will, if only communications can be

conducted in a meaningful way. Let us remember that the knowledge base of the expert system is that of the expert, and so are the contextual associations. The expert may also instruct the user on how to look at a problem.

To illustrate further the role of the expert, it may be instructive to examine the general components of an expert system. Invariably, one finds that the expert system software provides:

1. Methods for facilitating the interface between the user and the system
2. A knowledge base which contains the models, data that are necessary for the models, descriptions or attributes of models and data, and constraints or relationships.
3. Methods for deriving logical inferences.

Underlying, therefore, the construction and use of effective expert systems are the diagnostic and associative routines which enable the system to: (a) diagnose the right problem; (b) understand the decision-making context of the user; (c) engage the user in a Socratic dialogue to assist him/her in developing the appropriate subjective context if there is a discrepancy between the identified problem and the decision-making context, (d) associate alternative models, methodologies and heuristics with the solution of the identified problem; and (e) carry out the logical implications of the results of the application of such models, etc. within the identified context.

One of the practical criteria that I would like to propose for testing expert systems is to treat the system as if it were a student taking an examination. We, as professors, pose the problem to students, see the results and grade their capabilities in the chosen field. If the description of the problem is not complete enough to invoke the necessary decision-making context, the perceptive student either asks questions or, if he/she cannot, makes appropriate assumptions before he/she can associate alternative models and methodologies with the problem to be solved. These assumptions are part of the student's enabling associative context. They are stated, are transparent, and are part of the solution. Without them the student cannot effectively communicate with the professor and the latter cannot judge how expert the student is.

And so as with the expert systems, students must have the capability of performing diagnostics, be associative, have knowledge of objective models and methodologies and derive logical inferences.

C. The Role of Information

Information results from the association of relevant data with an appropriate decision-making context. It is used to reduce the complexity and uncertainty surrounding unstructured situations by constraining the problem. One way of looking at the relationship between information and data is as follows:

Information (Level 1) = Data Associative Context (Level 1) where: Data are the output of "objective" or "rational" models as defined or decreed by the "experts".

: Associative Context, is the result of subjective processes or the output of subjective models. It is the "way" particular decision makers "look at things" and use data.

The more unique the associative context is, the more personalized and unique the information that results from the use of common data. All of us may look at the same "objective data" but will not necessarily derive the same information. That is what distinguishes good from bad decision makers; the use of relevant rational models and objective data driven by a powerful subjective associative context. If, however, the context of the unique and successful decision maker were to be imitated and accepted by the experts, it would then become part of an "objective" or "rational" model. As a result what was before "Information (Level 1)" will become "Data (Level 2)", and new subjective contexts at a higher level must be created for uniqueness. That is how innovation, organizational knowledge, change and the consolidation of learning are generated continually.

Another characteristic of information is that it is probabilistic. There is nothing deterministic about "the subjective context", the assumptions that are part of it and the postulated relationships. The greater the complexity of the decision-making situation the greater the number of probable contexts. Even the so-called "rational" or "objective" models are often probabilistic.

It is clear from what we have said, therefore, that real expert systems should be diagnostic, elicit information so that they emulate (model) the decision-making context and protocol of the user, be associative-perceptive on context, protocol, problems, causes, and possible solutions, be an educator, on the appropriateness of context as well as objective models given a context, be probabilistic, learn from the experience of the user, and finally restructure (self-organize) themselves to incorporate the acquired knowledge. These are the desirable attributes of expert systems as indicated by the premises underlying their use. Unfortunately most, if not all, of the commercially available expert systems fall short of these specifications.

EXPERT VERSUS INTELLIGENT SYSTEMS

As mentioned earlier, expert systems attempt to make available to the user the advice of the expert. These systems could distribute scarce expert opinion regarding alternative decisions and models the experts use to make decisions in the prescribed situations. This assumes that the decision maker has properly identified the problem or through proper diagnostic routines at the user interface level, the problem has been properly identified. If not, the expert system will at best only serve the role of providing expertise in solving well defined and well structured problems. It can bring forth models used by experts, in the context of the expert. In effect the decision-maker will be the expert.

If it were true that we, the experts, not only know and teach about the relevant theory, methodology and tools of our respective fields of expertise but we can also make the best decisions, then we would not need managers. We,

the so-called experts, would conquer the managerial world, both at the firm and the national level, and be the latest in the line of succession of others who made similar claims in the past, such as the accountants, work simplification advocates, industrial engineers, operations researchers, organization behaviorists, management information system specialists, and even lawyers.

I must admit that I have some philosophical problems with terms such as expert systems, expert support systems, expert decision support systems. These terms, in my view, tend to portray the manager either as superfluous and replaceable by systems, or as a weakling needing to be propped up. They do not give enough credit to the beleaguered executive who is trying to recognize patterns, structure the world around him, develop enabling assumptions, relationships and contextual associations, in order to manage uncertainty and complexity especially in a technology-based world. Furthermore, these terms give the impression that the system is the driving force or is in control, and that the only thing an executive needs is to obtain one of these systems and his/her place in the Pantheon of successful executives will be guaranteed.

The world of the executives unfortunately is not simple, cannot be compartmentalized into mutually exclusive or independent classes of problems, and it is ever changing because of competition, technology and knowledge gained from experience. The more important the problem is, the more global, the more unstructured, the more integrated and the more complex it is, requiring hypotheses testing for learning. Therefore, the possibility is very remote that important problems can be solved by ad hoc expert systems, without the help of the executive to constrain them and avoid excessive and damaging distortions.

Expert systems, as known today, are for the most part specific purpose systems. The probability, therefore, that the purpose for which these systems have been designed to address and the needs of the decision maker will match is quite small. Unless, of course, the expert systems are generic, flexible, adaptive, or designed to assist the experts within organizations to design their own expert systems. These are attributes of what I call intelligent systems (Zannetos 1965).

INTELLIGENT SYSTEMS

What the executives in my view need are systems which in addition to having the above attributes can serve as an extension of their (human) capabilities and are part and parcel of the integrated management information system of the organization. This is in order to be (a) able to draw from global data bases and specialized models, (b) activated automatically by deviations from patterns of operational data and critical environmental factors, (c) brought to the attention of appropriate managers probable changes in the patterns of assumptions, operations and functional relationships with their probable implications and (d) self-organizing, i.e. classifying, testing hypotheses, learning and restructuring data bases and relationships.

A few years back I described in detail the attributes of intelligent systems (Zannetos 1965; 1968; 1978) so I will not repeat myself. The interested reader may wish to resort to the original sources for it is all there. I must stress, however, that the systems that I have advocated over the years have many unique features some of the most important of which,

although feasible, have not as yet drawn the attention of the designers of so-called expert systems. A possible and partial exception to this statement is the Operations Advisor of the Palladian Software Company, the debut of which I am awaiting with anticipation.

Four of these features which have a direct relation to the closing part of this paper, are:

1. The unification between accounting and statistical (probabilistic) data and analyses; thus enabling an integration between the on-going data gathering system of an organization, the operational models, and a probabilistic causal-diagnostic intelligent system.
2. The generation of intelligent information by the system when it is queried and even when it is not, enabling the information to be brought to the attention of the appropriate people within the organization.
3. The generation of signals out of operations which indicate the probable necessity for, and the direction of, changes in the organization structure (indicating the probable value of relative centralization or decentralization for particular types of decisions and specific locations within the organizations).
4. General applicability of the proposed iterative, hierarchical, recursive and associative models and protocols, thus allowing application at any level within the organization and expansion to other units or levels cocentrically to encompass the total organization.

INTELLIGENT SYSTEMS, THE FUNCTIONS OF THE EXECUTIVE
AND ORGANIZATION STRUCTURES

Looking ahead to the day when intelligent systems will become extensively used within organizations, some very fundamental realizations emerge which will have a profound impact on the role and functions of the executive and on the organization structure. Whether the top executives become the proponents of the introduction of these systems within their organizations or not, will not in any permanent sense shield them from their impact. And even if only one competitor effectively uses these systems, because the visibility and the strategic competitive advantage gained will be so great, the rest of them will be unable to idly sit by. The value added (increase in productivity) because of effective decision making is potentially so enormous that the laggards will be overwhelmed. As in the case of many other managerial uses of computers, the major benefit is not in the reduction of cost, but in the value generated from effective decisions, some of which could not have been feasible or even possible otherwise.

To the extent that the greatest value of intelligent systems is in the replication of the organizational intelligence and of the appropriate associative decision-making context, these systems bear the "signature" of both the organization and the decision maker. They are evolutionary and learn from the experience of the user regarding both protocol and causal diagnosis. Therefore, the sooner an organization adopts them, the earlier the start toward building the organizational and individual subjective knowledge base. This is to be contrasted with the objective knowledge base which the experts

build into the system. Some organizations may be tempted to wait "until technology is perfected." Given the dominant role played by the subjective context, the small cost relative to the potential value of these systems, and the great opportunity cost of time delays, an organization need only look at the evolution of basic artificial intelligence hardware and software for timing. To many of us, it appears, that the time is now.

Over the years I have heard top executives lament that they have no time for strategy and planning because of the many crises emanating from operations. The more serious the crisis, the greater usually is the intensity of the lamentations. Delegation of decision making is normally not considered a viable solution, because they feel that their subordinate managers and staff are not capable of making the right decisions; especially if the organization is in trouble. In fact the top executives may feel that there is already too much delegation and decentralization, and that is why the organization is in trouble in the first place.

If we look carefully behind the symptoms of the problems of top executives, we will invariably find that they really have difficulty in managing the complexity that is manifested in their operations. More often than not, concentrating on symptoms will lead to the wrong solutions, excessive centralization, punitive controls, recriminations and paralysis. Even in organizations where there is a lot of basic strength, in terms of technology and divisional management, desperate, uneducated and fumbling measures by top management addressed at the symptomatic level, will lead to suspicions and alienation, unnecessarily delaying the process of recovery. Unfortunately, because our present information systems do not record "how

things would have been" but only "how things were", there is little or no "objective evidence" of the cost of missed opportunities and fumbling.

The role of uncertainty and complexity and their impact on organization structures has been pointed out many years ago (See among others: March and Simon, 1958; Zannetos 1965 (a) and (b); and Galbraith 1973). To the extent that intelligent systems will immeasurably cut complexity out of management systems, their impact on the organization structure and on the management process will be all pervasive. This is so because intelligent systems will:

1. Enable the decentralized use of the planning stock of knowledge of the organization.
2. Serve as a coordinative mechanism of (a) plans to be consistent with global objectives and (b) operations.
3. Signal interdependencies to all those who will be affected by critical decisions, determine the impact of decisions and projected actions and eventually propose appropriate actions with second order effects in mind.
4. Encourage learning because of the explicitness of the interorganizational and intraorganizational interdependencies.
5. Learn (by themselves) from the experience of the users and update the stock of knowledge of the organization.

6. Provide signals for more effective control of operations and on the necessity for reorganization, in order to reduce further the complexity and uncertainty in the organization.
7. Reduce the necessity of meetings.
8. Effect a matching between problems and the participants at problem-solving sessions, enable the participants to become effectively educated before such meetings, cut the time spent in meetings and increase their effectiveness.

The executive in an era of intelligent systems will be spending more time on strategy and seeing that the latter and the operations are integrated and less time on control of operations. This does not necessarily mean that he/she will be less able to exercise control. In fact the opposite. What it means is that with intelligent systems which serve as an extension of executive capabilities, managers will be able to spend less time on, and yet have more effective control over operations mainly because the system will bring to their attention only significant deviations from patterns. More significantly, executives will be able to better "control" decentralized decision making through integrated models, and therefore should feel more comfortable delegating. Thus "control" will be applied effectively at the point of allocation or commitment of resources and not on a post mortem basis.

Lower level managers will also feel more comfortable making decisions, in that systems will permit the application of the latest decision-making

methodologies in a more global context. Decisions, as a result, will be more timely, more efficient and more effective. But above all, these decisions will be of a different scope.

Today, managers are often paralyzed in making decisions because (a) they do not have at hand the appropriate methodology and tools and (b) they cannot assess the consequences of their decisions on the total organization. Intelligent systems should alleviate to a very large extent both of these burdens and expedite effective decision-making.

Organization structures in the future will be flatter than what we find today. A lot of staff analysts and their functions will be eliminated by intelligent systems, but there will be a rise in the role of the Manager of Support Technology. The latter will be concerned with the introduction into the organization of the latest technology in computers, artificial intelligence and systems, and will serve as the catalyst, educator and transfer agent.

The most pronounced influence of the new breed of systems will occur in the relationships between managers. These relationships will be at the core of major organizational restructurings, only some of which will be reflected in formal organizational charts. For the most part, managers will be receiving signals from systems on the interrelationships between their decisions and operations and those of other managers. As a result, temporary reorganizations depending on the issues involved will be taking place, which will affect the relative centralization and decentralization of organizational subentities. There will be an ongoing and dynamic relationship between

strategy and structure, and it will be reflected in the intelligence of the integrated decision-making models of the organization and the context used by the respective executives. As knowledge is gained and incorporated into partial models, the systems must be flexible enough to integrate it into more global models, and the organization structures must be able to adapt to it without creating crises.

Several years ago we predicted some of the characteristics of these organization structures and we concluded that the dominant impact of computers and systems will be relatively decentralizing. (Zannetos and Sertel, 1970) However, what we timidly implied then, needs to be repeated and stressed. The managerial world in the future will be so different from what we know today that a comparative-statics analysis will be meaningless. On the other hand, progressive managers who are caught in this torrent of change will have made so many adaptations and so fast, that they will not notice the difference. Unless, of course, they pause and are left behind.

BIBLIOGRAPHY

Galbraith, J., (1973), Designing Complex Organizations. Addison-Wesley, New York, NY.

March, J. G., and H. A. Simon (1958), Organizations. Wiley, New York, NY.

Simon, H. A., (1960), The New Science of Management Decision. Harper & Row, New York, NY.

Simon, H. A., (1965), The Shape of Automation for Men and Management. Harper & Row, New York, NY.


Zannetos, Z. S., (1965a), "On the Theory of Divisional Structures: Some Aspects of Centralization and Decentralization of Control and Decision Making", Management Science Vol. 12, No. 4, December 1965, pp. 49-68.

Zannetos, Z. S., (1965b), "Measuring the Efficiency of Organization Structures: Some Implications for the Control System of the Firm". Sloan School of Management Working Paper 117-65, M.I.T., Cambridge, MA.

Zannetos, Z. S., (1968), "Toward Intelligent Management Information Systems", Industrial Management Review, Vol. 9, No. 3, Spring 1968, pp. 21-38. (A revision of Sloan School of Management Working Paper 155-65).

Zannetos, Z. S., "Intelligent Information Systems: A Decade Later". Sloan School of Management Working Paper 1028-78, M.I.T., Cambridge, MA.

Zannetos, Z. S. and M. R. Sertel, "Computerized Management Information Systems and Organizational Structures". Sloan School of Management Working Paper 486-70, M.I.T., Cambridge, MA. Also appeared as chapter in Management-Informationssysteme-Eine Herausforderung an Forschung and Entwicklung, Erwin Grochla and Norbert Szyperski (eds.), Wiesbaden: Betriebswirtschaftlicher Verlag Dr. Th. Gabler, 1971, pp. 695-720.



BASEMENT

Date Due

2-12-90

MIT LIBRARIES DUPL 1



3 9080 00579029 7

12

