An Examination of the Impact of Expert Systems on the Firm: The Case of XCON

By: John Sviokla
Graduate School of Business Administration
Harvard University
Soldiers Field
Loeb 22
Boston, Massachusetts 02163

Reprinted by special permission of the MIS Quarterly, Volume 14, Number 2, June 1990. Copyright 1990 by the Society for Information Management and the Management Information Systems Research Center at the University of Minnesota.

Executive Overview

Most IS professionals would agree, I think, that expert systems are destined to play an important role in the design of effective information systems for organizations. Nevertheless, a number of important questions remain concerning the development of these systems—such as their costs and benefits, the extent to which they should be embedded within routine operational systems, and their organizational and behavioral impacts.

The XCON system developed by Digital Equipment Corporation is generally regarded as one of the most successful applications of the expert systems methodology within an operational system. XCON is used to "configure" customer orders for computers—i.e., to determine the precise composition, physical placement, and interconnection of parts needed to assemble a given computer having a specified set of optional features. Owing to the huge number of possible variations, configuring a customer order turns out to be a very complex problem.

This article provides a detailed examination of the effects of XCON. The study involved a number of interviews and field studies within Digital, as well as a review of related documents. Conditions prior to the introduction of XCON (1978) are compared to the post-XCON situation (1985). Prior to XCON, configuration was performed by "technical editors" within the responsibility of manufacturing management. XCON largely replaced this manual process. This article looks at such matters as the quality of decisions made by the system, the benefits achieved, and the organizational impact.

The effects of XCON were not altogether predictable. Although the system reduced the number of technical editors, this reduction was more than made up for by the large technical staff (around 30 people) needed to develop and maintain XCON. Despite this rather substantial increase in the staff, the benefits of XCON have more than justified the costs. The accuracy of configuration reports has improved greatly, and a great deal of additional information is provided. This has led to increased efficiency in the assembly process and field support. The added accuracy and detail played a major part in allowing Digital to assemble equipment at the customer site rather than at an internal staging facility where errors or missing parts could be corrected more easily than in the field. A substantial reduction was achieved in the average and variability of the leadtime required to deliver products—increasingly an important factor in the highly competitive computer market. Automation of the configuration process permitted greater centralization with its attendant tighter control and greater standardization. XCON has now become an indispensable part of Digital's logistic system and arguably, an important strategic asset of the firm.

This case study exemplifies the likely results of embedding greater "intelligence" within a firm's mainline information systems: the effects can be profound, widespread, and strategic. It serves as a valuable reminder that big wins often come from doing a better job in the trenches.

An Examination of the Impact of Expert Systems on the Firm: The Case of XCON

By: John J. Sviokla
Graduate School of Business Administration
Harvard University
Soldiers Field
Loeb 22
Boston, Massachusetts 02163

Abstract

This article presents a new and different perspective on a familiar story. XCON, Digital Equipment Corporation's (Digital) expert system for configuring PDP and VAX computers, is a well-known commercial expert system. However, much of what has been written about XCON has been from either a technical or development perspective. In a different vein, this research examines how the use of XCON changed the task (configuration) it was designed to support, how it altered the roles and responsibilities of the technical editors (the individuals who performed configuration), and how the system helped eliminate an entire step in the company's manufacturing process, thereby saving Digital an estimated $15 million.

Drawing upon the concepts of organizational information processing, as operationalized by Jay Galbraith, this study shows that use of XCON increased the information processing capacity of the organization. Moreover, the system altered the local execution of the configuration task and provided for a company-wide shift in the process of configuring. XCON allowed a more centralized, controlled, and detailed application of configuration knowledge. This augmented knowledge application and had the benefit of supporting Digital's product strategy. The main drawback was an increased need to commit resources to ongoing maintenance of the evolving XCON system.

Keywords: Expert systems, knowledge-based systems, artificial intelligence, organizational impacts, technology impacts, organizational information processing, XCON, manufacturing systems

ACM Categories: K.4.3, K.6.0, H.1.2, I.1.2.0, I.1.2.1, J.1

Introduction

Many consider XCON to be one of the most successful commercial expert systems in use. Its developers and users document $15 million plus in savings from 1980 to 1985. Yet few people understand how XCON created these benefits. There is even less understanding of the effects XCON use created for the underlying business task it supports. The purpose of this article is to provide an in-depth analysis of the effects of XCON use on the process of computer configuration at Digital and the resulting managerial implications of these effects.

More generally there has been little or no attempt to relate ES to other issues that affect the information processing capacity of a firm. Moreover, there are omissions, particularly with regard to the organizational issues of ES building and long-term use. This lapse probably exists because ESs are so new to industry that, until recently, business academics and practitioners have had little chance to explore their complexities. Mumford (1987) suggests the need to study how expert systems are actually used as opposed to how designers expect them to be used.

This study uses pre/post field investigations to discover how use of XCON altered the execution of the task it was designed to support. It also examines the changes XCON engendered in the information processing capacity of the firm.

Literature on Expert Systems

Use of ES in commercial organizations is a relatively new phenomenon, and examination of their implications is just beginning to be explored. Currently, there are three categories of literature concerning ES. The first describes how to build expert systems. *Building Expert Systems* (Hayes-Roth, et al., 1983), a collection of articles de-

signed to aid ES development and use, and *Expert Systems: A Practical Introduction* (Sell, 1985), which presents formulae for building ES based on the author's experience at Digital, are examples of a growing number of books written on the subject. Some publications focus on tools and applications (Harmon and King, 1985; Harmon and Rex Maus, 1988). Others focus on specific problem areas (e.g., Mockler, 1989).

The second category of literature, which chronicles the work of computer scientists, primarily describes theoretical issues, such as the methods of knowledge representation (Minsky, 1975), search strategies (Knuth and Moore, 1975), and so on. Often an ES has many developers, each of whom contributes to the lineage of writings that tracks the system over time. The series of articles from Shortliffe (1976) through van Melle, et al. (1981) for the MYCIN/EMYCIN project is one example of such a stream of research.

The third category of works focuses on commercial uses of ES. These "system biographies," written by practitioners, review the commercial ramifications of their handiwork. *The AI Business* (Winston and Prendergast, 1984), a collection of papers by managers and their academic allies, and the book *The Rise of the Expert Company* (Feigenbaum, et al., 1988) are two recent examples of this genre. The wealth of personal commentary and experience embodied in these works makes them valuable first-hand accounts of pioneering real-world efforts. Yet there are no unifying research themes in their observations because the authors are practitioners describing their experiences, not researchers seeking effects. Nor are any investigators looking at what effects these expert systems have on the user organization.

There are many possible ways to look at the impact of information systems on the organization, and there is a rich tradition of research in the more general area of information technology (IT) in organizations. In a classic review article, Kling (1980) highlights two schools of thought with fundamentally different assumptions about the organization: segmented institutionalist and systems rationalist. Kling points out that these assumptions drive very different analyses and interpretations of the phenomena observed. He notes that the technology is "malleable though not entirely plastic" (p. 100). Thus, to speak of the "social impact of computing" is as invalid as it is valid. The role of referent perspectives shapes the investigation.

In an analysis of citations in the MIS field, Culnan (1987) notes a distinct sub-field concerning the impacts of computers on the organization. Citing Culnan and building on Kling (1980), Markus and Robey (1988) analyze the underlying causal agencies, beliefs about logical structure (static versus dynamic theory), and levels of analysis (micro, macro, or mixed). By going a level deeper in the review and analysis of existing impact literature, they observe that some of the differences result from the implicit themes and analyses of the different researchers. These observations are consonant with Fry (1982), who notes the wide variety of constructs for IT and organizations. Therefore, it is important to characterize the motivation for this study.

Expert systems explicitly set out to capture and assist or automate complex, non-algorithmic decisions. In an organizational context, the manager creating such a system is consciously designing a decision-making process for the firm. Simon (1945) and Huber (1984) both note the central importance of designing the decision processes to the effectiveness of modern firms. In a theoretical review and analysis of the strategic systems literature on MIS, Bakos and Treacy (1986) suggest that systems that help "shift the limits of rationality" (p. 109) beyond the bounded rationality of humans (Simon, 1945; Cyert and March, 1963) are valuable opportunities for increasing the efficiency and effectiveness of the organization. Therefore, a large-scale, successful system, such as XCON, which explicitly attempts to re-design an existing decision process, should shift the information processing capabilities of the organization.

Design of Current Study

Given the complexities described above, the first research design question was to choose the best method to gather an understanding of the impacts of ES use on the organization. Benbasat, et al. (1986), in a review of case research in information systems, note that one important motivation for case research occurs when "the research is interested in complex problems which cannot be removed from the setting without destroying the phenomenon" (p. 7). Bonoma (1985)

advocates the use of case research in environments "where the existing body of knowledge is insufficient to permit the posing of causal questions, and when a phenomenon cannot be studied outside the context in which it naturally occurs (p. 207).

The impact of expert system use on a firm's information processing capacity is intertwined with an understanding of the business process it is designed to support. Hence, a case study seemed appropriate.

The next research design question was to choose measures—within the case approach—to provide a useful organization for the data. After a review of the variety of measures from the literature, two measures of impact were chosen: *performance programs* (March and Simon, 1958) and *information processing capacity* (Galbraith 1973; 1977). Within the performance programs, two levels of analysis were chosen: individual and organizational. Rousseau (1985) defends this type of mixed-level analysis when the effects are not simply macro or micro. (See Markus and Robey, 1988.) In this exploratory case work on XCON, it was important to be open to potential changes at the micro and macro levels of the task.

March and Simon (1958) provide the most complete discussion of performance programs. They emphasize the limited nature of humans to process information and the various sources of uncertainty that an organization faces. One of the critical mechanisms that an organization uses to make decisions and carry out work is performance program, which March and Simon define as follows:

> We have seen that under certain circumstances the search and choice processes are very much abridged. At the limit, an environmental stimulus may evoke immediately from the organization a highly complex and organized set of responses. Such a set of responses we call a *performance program,* or simply a *program* (p. 141).

March and Simon (1958) have some helpful ideas on how to look for performance programs in organizations. Specifically, they suggest observing behavior, interviewing members of the organization, and examining documents. The data gathered for this study that identify performance programs closely parallel March and Simon's suggestions.

Thirty-three interviews were conducted during more than a dozen field visits over a period of five months. In addition to interviews, observational and archival data were collected. These three sources were analyzed by the author to create a before-and-after picture of the configuration process and XCON's role within Digital.[1] For data collection and analysis purposes, pre-XCON was defined as 1978—the period just before XCON was introduced. Post-XCON is defined as 1985—five years after the system was in active daily use.

The second measure used within the case research was information processing capacity. The information processing capacity of the firm was chosen as a means to capture the changes in input and output of the target task. As affirmed by Bakos and Treacy (1986), it was expected that the use of an expert system would expand the ability to the human problem solvers and increase the information processing capacity of the organization for the target task. In an extensive review of the information processing literature, Zack and McKenney (1989) note that there have been very few operationalizations of the information processing concept in field research. There is no completely satisfactory model for the information processing capacity of a firm (Galbraith 1973). It is not the focus of this research to find and prove such a model. Rather, the information processing capacity of the firm and performance programs are used as organizing concepts for data collection, analysis, and comparison to understand the impact of expert systems use on the firm.

For this study, the specific operationalization of this measure draws on the work of Galbraith (1973; 1977), who defines the information processing capacity of an organization to be the number of different input resources, the diversity of the outputs, and the level of task performances (Galbraith, 1977). These definitions are,

---

[1] Historical information—especially interviews about past events—presents methodological challenges: people forget the correct facts; they remember the facts incompletely; or, with time, they recollect the past in a manner inconsistent with the facts. Archival data and multiple interviews were used in this study to cross-check the recollections and verify important historical issues (Bonoma, 1985). Detailed process maps were drawn to create a flow of the decision process in the firm. Each iteration of these maps was reviewed for accuracy by three different experts within the firm. This article presents a verbal description of these maps and describes the significant differences from pre-XCON to post-XCON.

for this study, pragmatic and simple: *inputs* are defined as the number of data sources used and people consulted in the delivery of the product or service; *outputs* are defined as the number of different versions of the product/service; the *level of task performance* is the management's assessment of performance and any available objective measures. In addition to the two measures (organizational programs and information processing capacity), an in-depth study of the business situations was made in order to help the researchers interpret the data gathered.

The Company Setting

With 1986 revenues of $7.6 billion, net income of $617 million, and more than 89,000 employees worldwide, Digital Equipment Corporation was the second-largest computer manufacturer in the United States. In an industry that experienced phenomenal growth during the past three decades, Digital had built a reputation for engineering excellence and an ability to deliver reliable computers with superior performance for the price. While most computer manufacturers restricted their offerings to a few standard systems with limited options, Digital designed its machines and marketing to give customers a broad and flexible set of basic systems with many add-on options. Using this strategy, Digital grew rapidly. In the 13 years from 1972 to 1985, its equipment revenues increased 29 percent per year compounded.

Because of the product line's flexibility, many Digital employees—especially those in manufacturing and sales—needed constant retraining and updating of their product knowledge. This almost-continual learning process represented a significant price that the product strategy necessarily exacted.

The Issue—Configuration

Configuration was one of the key controls Digital used to ensure that the products its salespeople sold were, in fact, buildable. The term "configuration" means translating a customer's needs into a complete computer system. In 1974, when a customer bought a Digital computer, the customer worked with a salesperson to identify needs and tailor a system and a list of options to match the customer's precise requirements. This process was the first step in configuration.

A typical configured order contained one or more main components, each with one to 100 (or more) options.

There were three basic tasks in configuration:

1. Translating customer needs into Digital products
2. Checking the completeness and accuracy of the sales order
3. Designing the specific placement and connection of all parts in the order

The configuration issue was important to Digital because sales orders—the raw material of configuration—constituted the most important customer information that flowed through the company. Configuration was *the* major means used to verify, process, and ensure the accuracy of sales orders.

As early as 1974–75, the complexity of the configuration process was substantial. There were approximately 50 types of central process with 400 core options, each of which had approximately 10 versions—that is, 4,000 options overall. No one knew for certain how many possible working configurations existed, but the number was estimated to be in the millions. Some customers needed minor variations, such as three tape drives instead of two. However, in the overall order population, many different types of configurations passed through the firm—hence the complexity.

Misconfiguration posed many possible problems. For example, large customers might place two identical orders but the delivered machines might have different assembly. This difference in assembly might be cosmetic, such as a different color cabinet, or functional, such as a different placement of a disk drive. This lack of standardization did not communicate the quality control Digital wanted to present to all its customers.

In creating an organizational program for configuration at Digital, it was discovered that, pre-XCON, every configuration was reviewed in at least three places in the order flow: initial technical edit, manufacturing technical edit, and final assembly and test. After a salesperson drafted a contract with a price, an initial technical editor—usually located in a field sales office—scanned the order to ensure that all the necessary components were included. After this initial verification, the order was "accepted," which meant that Digital committed to deliver the stated

system on the specified dates and for the agreed-upon price.

Manufacturing management assigned primary technical editing responsibility to the second stopping point: the manufacturing technical edit group. Members of this group generated the configuration drawing and controlled the configuration through manufacture. The technical editors (TEs) within the manufacturing technical edit group usually began the configuration process by retrieving a sales order. The TE—equipped with a telephone, a bookshelf of technical manuals, and experience—was challenged with ensuring every order's technical completeness before creating a drawing showing the placement of all the system's modules. Because of the vast number of possible options and the exactness of assembly requirements, this task required intimate knowledge of an enormous amount of technical detail. The time spent on an order varied with its complexity, size, and frequency. On average, it took 15 minutes to configure an order by hand, but correctness was uneven. According to surveys conducted by the Intelligent Systems Technology Group (ISTG) within Digital, the human configurers were completely correct only 65 percent of the time. The ISTG criteria for correctness were relatively strict: a configuration could be scored "incorrect" if a cable was off by one foot.

The difficulty of configuring depended greatly on the newness, complexity, and frequency of the order. The actual time it took to edit an order varied from five to 10 minutes to an hour or more. The TEs and their management estimated that the cycle time to complete the technical editing process for an order—pre-XCON—was one to two days on a regular basis.

Task performance measures were largely subjective and locally determined by each technical editor. Given the immense number of possible configurations, each technical editor developed a repertoire that worked best for him or her. There were no optimal configurations; the evaluation of any particular configuration was simple: if it worked, it was right. Consequently, two people could configure the same systems differently, yet both could be correct according to the rating criteria.

A major method—perhaps *the* major step taken to ensure that systems were built correctly—was final assembly and test (FA&T), the third stopping point in the performance program for configuration. Before XCON, approximately 90 percent of Digital's system volume passed through FA&T, where all configuration changes and errors were settled by creating and testing the system. In a real sense, this was the final technical edit.

Because of a large backlog of orders for Digital products during this time as well as frequent changes to computer orders, it was not unheard of for the total order-processing time to extend four to six months—although such a situation was uncommon. Usually, the cycle time from when an order reached manufacturing technical edit until it was shipped to the customer was approximately 10 to 15 weeks. Of that time, once the parts were ordered and received, approximately four to six weeks were spent with the machine in FA&T.

Solving the configuration problem—XCON

By 1975, management recognized configuration as a serious issue with an overwhelming magnitude of potential problems. To support the massive FA&T effort necessitated by the growing volume of equipment sales, the company had created an FA&T facility in Westminster, Massachusetts. This facility covered more than 13 acres under one roof and cost between $15 and $20 million to build. Inventory to stock the plant cost approximately $20 million more. By 1975, Digital was grossing $433.2 million in equipment sales, and by the early 1980s that figure was projected to quadruple.[2] Such growth would necessitate four or five more FA&T plants. If the FA&T "solution" to the problem continued, tens of millions of dollars would be tied up in plant, equipment, and inventory simply to support the assembly operation.

Organizational dilemmas also confronted Digital managers. Estimates of the number of configuration experts at Digital ranged from 50 to 60 experienced individuals. The number of people projected to be required by the early 1980s ranged between 50 and 200. Furthermore, the underlying problem continued to become more and more difficult. It was clear that no human being could possibly remember and manage all the relevant technical data required to perform con-

[2] In fact, Digital exceeded $2.4 billion in equipment revenue in 1981.

figurations while simultaneously tracking and learning about the onslaught of new product offerings and options.

The impetus to seek a solution began with the realization of the enormity of the FA&T cost. This search led to the development of XCON, which began with a collaboration between Digital and John McDermott, a professor of computer science at Carnegie Mellon University. During the period from December 1978 to April 1979, McDermott—with the help of Digital personnel—created a prototype with 210 rules and 100 components. The system demonstrations were impressive enough to garner support to fund the XCON project, which was to provide the system with the basic knowledge of configuration. McDermott remained the principal designer/developer. Toward the end of 1979, the system was validated with 50 "representative orders" which were reviewed by six experts. The system made two "significant" and 10 "low-level" mistakes.

Conceptually, the system is very straightforward. There are three basic inputs to XCON:

- All VAX orders and selected PDP-11 orders
- Digital product information
- Engineering/marketing configuration guidelines

In processing an order, XCON performs the following activities, with the data provided from the three inputs above:

- Adds or deletes components necessary to make the order complete and correct (includes comments on the reason for each change)
- Assigns each component to its correct location
- Configures cable layout, lengths, and connections
- Calculates cable layout, lengths, and connections
- Calculates the address vectors (the logical addresses of the options in the computer)

The outputs from XCON are:

- Order line item summary
- Detailed configuration drawings at the cabinet, box, and backpanel levels
- Explanatory text regarding:
  - — Parts not needed in configuration
  - — Spare parts
  - — Additional parts needed for completeness
  - — Parts that, for technical reasons, were not configured
  - — Address and vector settings
  - — Unused controller capacities
  - — Unused box power

The extent and precision of the output was much more complete than previous output by human configurers. By the end of 1979, Digital was using XCON in a pilot project at its Salem, New Hampshire, manufacturing plant. In 1980, the system was in active use at Salem. By early 1981, the system was in regular use throughout the company.

Performance program: post-XCON

There were two major changes to the performance program post-XCON. First, fewer areas in Digital continued doing technical editing. The coordination and control of configuration became standardized and centralized. Second, the roles and responsibilities of the technical editors changed dramatically. With the use of XCON, 90 percent of all orders were put through the system. There was no need to have the sales office check a configuration. Nor was there a need to make substantial configuration adjustments at the FA&T operation. Manufacturing technical edit was the one and only major configuration generation point.

This centralization led to more control of the configuration process. It also standardized the execution of configuration policy on each and every XCONed order. Moreover, the best configuration expertise began to migrate to the ISTG organization—the software development team responsible for keeping XCON up to date.

The second major shift in performance programs came at the TE's desk. The homogenous task became differentiated, and three processes of configuration emerged: (1) XCON; (2) modified; and (3) manual.

In 1986, the vast majority of Digital orders went directly through the XCON system, as shown in the Salem plant's data (Figure 1).

"XCON" meant that the TE checked the XCON output. When the order—a list of line items—flashed on the screen, the editor would scan for part numbers, order completeness, and overall correctness. Often salespeople would include extra parts or parts for another system on the same

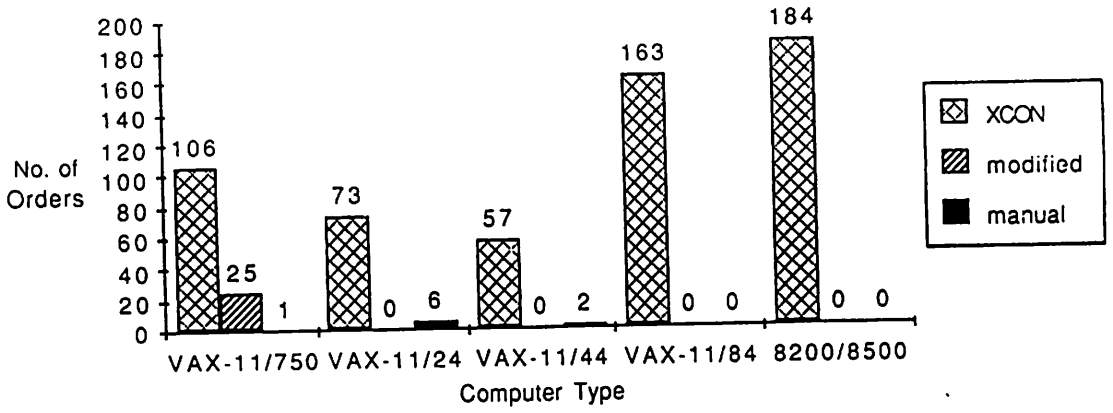


Figure 1. Configuration Data—Digital Salem (4 months in 1986)

order. When this occurred, these extra parts needed to be marked as "spares" so that XCON would not attempt to configure them and find them as mistakes. Consequently, a task for the technical editor was to ensure that all extra parts were correctly marked "spare." In cases where the spare components were similar to those requiring configurations, the TE might have to ask the salesperson for clarification.

At the other extreme, there were some configuration or planning tasks that could not be performed well on the machine. For XCON, this occurred relatively infrequently—only about four to six times per month per TE in the Salem sample. When it did happen, the Salem technical editors had to configure manually, as was done pre-XCON. There were also some special customer orders for which the technical editors decided not to use XCON at all, again configuring the entire order manually. For instance, military orders often specified the configuration in the purchase contract. In these cases, the TEs simply translated the contract specifications onto the configuration sheets because the military specifications might or might not agree with XCON's results.

The challenges arose in the second category of tasks, "modified" (for which XCON created less flexibility). Modifications to XCON content often occurred on brand new VAX models. Modifying an XCON configuration meant going back through the entire output—often 20 pages or more—to trace every reference to a particular line item or series of line items and their cable needs, address vectors, and other specifications. In the XCON case, this hybrid system of part ES, part manual seemed to be the most time-consuming and frustrating activity of all. As one TE noted:

> It is *so* boring! Often I have to take the output, lay it on a table, and get the white-out. I look through the entire thing, making all the necessary changes. People often expect to be able to see the XCON output on the computer, so I have to adjust it on the screen. To do this I call up the output file, put it in a word processor, and edit it line by line. It takes forever. We have been working with the people in Hudson [the XCON developers] to get an editor [software], which could do all that for us—especially tracking down all the dependent details (Source: author interview).

In general, TEs were guided by rules of thumb gained from experience to separate orders requiring careful review from those requiring less. For example, early versions of new products such as the 8600 series of VAX computers often contained mistakes. Therefore, technical editors, alert to the 8600 series, would carefully inspect orders containing one. Overall, the technical editor's job had changed significantly.

Information processing capacity: Post-XCON

With the XCON system in place, the second measure—the information processing capacity of the firm—increased significantly. Table 1 shows the inputs, outputs, and task performance com-

Table 1. Digital's Pre- and Post-XCON Task Performance Components for Configuration

Pre-XCON	Post-XCON
Inputs—Increased	
Sales order	= Sales order
Memoranda/guides	= Memoranda/guides
Training	= Training
Experts consulted	+ Experts consulted
Technical Editor[a] (23 in 1981)	+ + Technical Editors (17 in 1986) XCON Development Group[b] (30 in 1986)
Outputs—Significantly Increased	
Phone follow-up	- Phone follow-up
8″ x 10″ hand drawing of configuration	+ + Configuration including: parts to be added, exact layout, assembly information, vectors, cabling, etc., explanations
	+ + Information for XSEL and other systems internal to Digital
Task Performance—Significantly Increased	
In 1981, average review of 100–300 orders/editor/year	+ + In 1986, average review of over 1,000/editor/year
Correctness range 65% to 90%[c]	+ + Correctness 95–98%
Total volume ~ 4,000–5,000 VAX orders, ~ 15,000 other orders	+ + Order volume ~ 60,000+ orders
Even distribution of orders through TE	+ + Short time windows; distribution uneven with high peaks
	+ + Shorter cycle time (10–12 weeks in assembly to 2-3 weeks)

The indicators of effect and direction are shown for each item: + + is a large increase, + an increase, = stayed the same, - a decrease.

[a] This number excludes managers of technical editors.

[b] There were approximately 40 people in the XCON/XSEL development group as of July 1986. Most of them worked on XCON.

[c] This estimate and the estimate of 95–98% correctness are based on tests conducted by the Intelligent Systems Technology Group at Digital.

ponents for configuration—pre- and post-XCON—at Digital.

Inputs

Post-XCON, the inputs to the technical editing process were more numerous. It is important to remember that in the information processing assessment, an increase in inputs through specialization of tasks increases the information processing need of the task because each new input must be coordinated. With the XCON system, the information processing needs increased in this area because the XCON development group was another input to the process. Of the 40 or more

developers in the Intelligent Systems Technology Group (ISTG), approximately 30 were involved with the creation and maintenance of XCON—a net increase of 24 people with greater skill specialization. Simultaneously, the number of technical editors decreased from 23 to 17 in the period from 1981 to 1986. This decrease was more than offset by an increase in the number of technical staff needed to create XCON. Within the working definition of information processing, this translates into an increase in inputs.

It should be noted that the focus of the information processing view of the firm is different from traditional views of input. In a traditional input/output analysis, an increase in inputs means more material, labor, or capital. In the information processing analysis, inputs and outputs are not productivity measures but rather are approximations of the complexity of the information processing task. Consequently, the framework is useful for describing information processing effects but is significantly different from traditional productivity approaches.

With XCON, manufacturing, engineering, and marketing were also involved in creating the rules for systems' configurations. As of early 1985, all products had to have established and complete configuration rules in the XCON system before release. Overall, there was an increase to the inputs for XCON.

Outputs

The outputs from the XCON-aided configuration process increased dramatically. Outputs are defined as the number of different versions of the product/service. XCON clearly increased outputs. Configuration done by the manufacturing technical editors was always intended to help assemblers and technicians build machines. However, XCON output was so superior to previous hand configurations in terms of its detail and scope that manufacturing managers attributed savings of at least two hours of assembly time, per system, to the XCON output. There were additional salary savings in terms of being able to use less-skilled individuals to assemble systems.

The pre-XCON configuration paperwork was a one-page hand-drawn document of the basic system layout. Post-XCON documents were 10+ pages of detailed specifications and drawings indicating where every module and component was to reside. A more consistent and complex set of documents reflects an increase in information processing capacity, as described by Gailbraith (1973) and operationalized here.

Post-XCON output also had more audiences. Pre-XCON output was used primarily by manufacturing assembly people. Post-XCON output was used not only by manufacturing for assembly but also in the field engineering's assembly process. In addition, field service representatives used the customer's copy of the configuration as a trouble-shooting aid. It was assumed that there were many other users that ISTG simply did not know about. During the interviewing at the company, for example, it was discovered that the shipping department used XCON output as part of its training and packaging routine to assure that all the correct components were in the shipment going to the customer site.

Besides being more complete and precise, XCON output was also broader in scope. Before XCON, the manufacturing engineers or manufacturing technicians calculated address vectors and cable lengths. However, issues such as excess or unused capacities were not addressed. The XCON configuration document generated these numbers for every computer. This information could be communicated back to the salesperson to show the expandability of the system.

Another output of the technical editing process was a continual feedback of information to other areas in Digital. For example, XCON was used to support XSEL, the most important other use of XCON in Digital. XSEL is an expert system that helps a salesperson translate a customer's business needs into specific parts and computers. When a salesperson uses XSEL, the XSEL system accesses the configuration information and rules embedded in XCON. In this way, the XCON system helps support Digital's sales force in configuring the Digital product line.

There was one category of outputs—telephone follow-up—that decreased. The TEs at the Salem manufacturing plant stated that they were not called as frequently by marketing, sales, and manufacturing for advice on systems. However, this decrease did not change the assessment that outputs substantially increased overall.

Task Performance

The most striking effect of XCON was an increase in task performance. In 1981, 23 technical editors

created approximately 4,000 VAX and an estimated 15,000 other system configurations by hand. With XCON, in 1986, 17 technical editors created over 60,000 system configurations with increased quality and detail. Quality estimates varied widely, but all pointed to improved performance. Technical editors estimated that pre-XCON, they were correct about 90 percent of the time, whereas the personnel at ISTG figured correctness at about 65 percent. The difference in estimates was largely a reflection of the increased expectations of the configuration output that XCON generated versus an unaided TE. For example, pre-XCON configurations did not contain address vectors for the components; post-XCON they did. Consequently, post-XCON the basic configuration was more thorough. In any event, the estimates for XCON's percentage of correct configurations were at approximately 95 percent to 98 percent in 1986 on all orders.

In addition, there had been a change in Digital's business environment that increased the information processing need of the task even more. During most of the 1970s, Digital had a large backlog of computer orders that allowed at least some pre-production planning time. In the computer market of the early 1980s, the firm had to be more nimble because the backlog had evaporated and customers were expecting—and getting—faster delivery. Not only were the lead times shorter, but customers were waiting until later in each quarter to book orders. In 1986, the majority of Digital orders were booked and filled in the last month of each fiscal quarter. The order capture rate, as it was called, was more compact with sharp peaks. This volatility meant that more orders were driven through the system within a shorter time frame. XCON was an essential component in allowing the organization to cope with these peaks, which had been generated by the changes in the business environment.

Another major improvement in task performance was a reduction in an order's cycle time. The estimated cycle time between configuration and actual order shipment decreased dramatically—from an average of three to four months during the late 1970s to a minimum of one to two days in 1986. Cycle time for a VAX computer varied, but the minimum time elapsed through the POM (point of manufacture) part of the manufacturing process—from technical edit to shipment—was estimated to be two to three days. The average time was approximately three to four weeks, reduced from 10 to 12 weeks using FA&T.

Many factors besides XCON have contributed to this dramatic decrease in cycle time of the order. Competitive pressures forced mini-computer makers to strive to be more responsive in their manufacturing. During the early 1980s, the market for computers began to soften; this rekindled Digital's efforts to improve the speed of the order-fulfillment cycle. Nevertheless, XCON was a significant enabling mechanism in the firm's improved cycle time.

In summary, looking across the three components of information processing, inputs increased while both outputs and task performance significantly increased.

Organizational Implications of XCON Use

The organizational adaptations fostered by the use of XCON had four revealing features.

Technical editing system replaced

First, the organization replaced a system of technical editing, which relied heavily on experts and apprentices to manage a complicated problem on a decentralized basis, with a centrally supported system of intelligent software plus human operators. The latter system was able to cope with a higher level of complexity than had ever been managed by the company's previous configuration support.

By creating the XCON software, Digital created an asset that had a different capacity from the previous method of configuration. Before XCON, configuration capacity was determined by the number of staff available to process configurations. Total configuration capacity was limited by the number of trained individuals. Furthermore, the total number of configurers would have to be sufficient to match peak demands; this might have necessitated inventorying talent to meet high volume. In contrast, XCON's capacity was limited only by the hardware and software constraints and could, in practice, cope with high peaks in configuration throughout the organization. Therefore, Digital did not need to inventory configuration talent to serve the new market de-

mands. This allowed the firm to pursue a customization strategy that might not have been feasible if it had to build the organizational structure to support this strategy with human configurers.

Roles and responsibilities shifted

Second, roles and responsibilities shifted when XCON technology was introduced. In general, those associated with XCON's creation seemed to gain responsibility for development and maintenance. As the use of the system became an integral part of operating the business, responsibility for supporting the software became very real. As a software manager at Digital expressed it: "XCON is the only ES I know of today where the financial lights would dim if it were unplugged. That is a tremendous responsibility" (Source: author interview).

Transferring responsibility to users and developers, however, was not universal. In the case of XCON, the technical editors' job seems to have lost clout and become more clerical. As one of the TEs explained.

> It was more fun before XCON, when you had to figure out each system. You got to keep in touch with many parts of the company—engineering, sales, and marketing—to know what was happening. We still do that now, but not so much. Also, we have to do all that correcting (Source: author interview).

The most explicit recognition of the decreased responsibility was that, pre-XCON, technical editors could prevent an order from being built until all configuration issues were solved—post-XCON they could not. Also, the TEs used to be consulted more often. They did continue to get calls and electronic mail from different Digital people wanting information on how to configure a particular order, but this happened less frequently.

More subtly, XCON use enriched and reinforced the experts' assessment and role. This specific instance lends credence to Huber's (1984) general prediction that information technology would expand the reach of an expert's purview. It helps knowledgeable individuals communicate more effectively because it fosters a common language about computer configuration. Performance of a complex task improves when a common language is developed to state succinctly the relevant criteria and facts in the task domain (Macy, et al., 1953). Similarly, ESs themselves build a common language for a complex task. They codify opinions and integrate actions associated with the task, thus building a common set of recommendations, designs, and judgments in addition to the shared descriptive vocabulary.

Formal maintenance system developed

Third, the organization needed to formally cope with the maintenance of the configuration knowledge base. For both pre- and post-XCON, there was no specific training for technical editors, nor was there a prescribed "inventory" of TE talent. An individual with experience in manufacturing or order processing often served a brief apprenticeship to become a TE. Post-XCON, a staff of approximately 40 programmers maintained the XCON knowledge base. Because of the changing nature of product knowlege, over one-half of the XCON knowledge base was rewritten annually (Barker and O'Connor, 1989).

This centralization of the creation and maintenance of configuration knowledge raised some potential risks. For example, there was a risk that the reviewers of the system output would not understand the task well enough to modify XCON's results when modifications were needed. In 1985, the organization still had a mix of pre-XCON and post-XCON employees who could, together, bridge the operating procedures from one era to the next and make relevant adjustments to XCON output. The long-term effects of XCON use on the available pool of talent for configuration was an issue that had yet to be resolved in 1985.

Also, there was a risk that the expert system was a short-term solution for a long-term problem. Perhaps by 1985 it was time to revamp the entire XCON approach to order processing and product strategy; but because XCON could reasonably cope with the problem, the issue was never raised. This issue is always a problem when an expert system is used to augment a complex process. Designers and managers will always be confronted with the trade-off between managing the complex process with the aid of an expert system and simplifying the process itself.

"Progressive structuring" initiated

The fourth effect on the organization seen after XCON use is what I term *progressive structuring.* In the process of capturing and codifying the data, opinions, and suggested actions necessary to make XCON a useful tool, Digital personnel "progressively structured" the problem of configuration. The progressive structure of the problem articulates specific inputs, outputs, and increases in task performance. In short, progressive structuring increases the information processing capacity of the firm as defined by Galbraith (1973, 1977). The impact of progressive structuring on configuration was felt across the organization. The configuration acted as the focal point from which more and more structuring of the task emerged. Initial success with the system led to further development of new information processing tasks, thereby progressively structuring the total task of configuration.

With the task reasonably structured, new organizational options became available. ISTG managers felt that XCON was an integral part of Digital's move from an FA&T manufacturing strategy to a point-of-manufacture (POM) strategy. The FA&T strategy required that most systems (80 percent to 90 percent) be assembled before shipment, while the POM strategy shipped the same high percentage of the systems as separate components to be fully assembled for the first time at the customer site. XCON's robust, logical description of the configuration was an integral ingredient in coordinating and controlling the assembly necessary to execute a POM strategy successfully. ISTG attributed to XCON 20 percent of the savings ($15 million-plus) made by changing from FA&T to POM. Information processing capacity, in a sense, substituted for physical capacity. Formalization of this complex decision task allowed the organization to pursue a strategy with high information processing needs.

Conclusion

The use of XCON at Digital changed the management and execution of the knowledge-intensive task of computer configuration. The pre-XCON management challenge was to keep a sufficient supply of technical editors available to configure and reconfigure Digital computers. Because the relevant base of configuration knowledge was constantly changing as the company introduced new products, the quality of the configuration decisons was spotty. Since many derivative decisions, such as sourcing and assembly, are built upon the configuration, not surprisingly, other departments responsible for derivative decisions undertook their own configuration efforts to double-check the pre-XCON configuration.

Post-XCON, the control and application of the most up-to-date configuration knowledge was more consistent and correct, and the need for derivative examination of the configuration document dwindled. In fact, the robustness of the knowledge applied by XCON played a central role in the organization's removal of the last stage of the manufacturing process (e.g., final assembly and test). Meanwhile, the management challenge shifted from keeping staff up-to-date to keeping the software up-to-date.

From an organizational perspective, the information processing capacity of the firm increased. In particular, the combination of the knowledge-based system and the individual decision makers allowed for quicker dispersion of expertise and faster, higher-quality decisions. The first leverage point created by the knowledge-based system seemed to be the leverage of scarce expertise, and the second leverage point seemed to be the elimination of the second-guessing that occurred when a critical decision was not robustly executed. Thus, the organizational implications showed up locally in changed roles and responsibilities of the individuals performing the task and globally in the tasks that depended on the decisions made by the local experts.

Many organizations are dependent on individuals (such as technical editors) who must make decisions in knowledge-intensive domains. Product configuration, product engineering, and product service are three areas in which organizations train and deploy scores of people to meet the decision needs of the organization. Often, the initial decisions of these individuals drive derivative activities; therefore, it is critical to get the decision correct the first time. In domains where knowledge is rapidly changing and where no algorithmic solution is available, organizations invest in both multiple checks and rechecks of the decision and in considerable training of key employees. This type of solution is not only costly, it creates ambiguity and uncertainty around a critical business process.

Knowledge-based systems may be a powerful tool to aid organizations in addressing this problem. The creation of the system can help articulate a common solution to a complex problem. Use of the system can help standardize execution of corporate policy and directives. Overall, a combination of system knowledge and individual knowledge may lead to a more effective and efficient system for executing critical business decisions.

However, these systems create management challenges. To ensure ongoing viability of an organization's processes, the manager will need to consciously redesign the jobs of decision makers affected by the knowledge-based system. Moreover, use of these tools has the potential to severely de-skill talented individuals and create a knowledge-worker version of mindless work. At the extreme, the organization could become over-dependent on the knowledge embedded in the system and find it difficult to update the knowledge. However, the manager can act to consciously "grow" individuals who understand the task in enough depth so that they can monitor and update the system.

More generally, organizations are creating and implementing software tools that embed greater knowledge and skill. The ongoing management challenge is to understand the organizational implications of the use of these tools and redesign business processes and individual responsibilities in light of evolving technological capabilities.

References

Bakos, J.Y. and Treacy, M.E. "Information Technology and Corporate Strategy: A Research Perspective," *MIS Quarterly* (10:2), June 1986, pp. 107–119.

Barker, V. and O'Connor, D. "Expert Systems for Configuration at Digital, XCON and Beyond," *Communications of the ACM* (32:3), March 1989, pp. 298–320.

Benbasat, I., Goldstein, D. and Mead, M. "The Case Research Strategy in Studies of Information Systems," unpublished paper, Harvard Business School, Boston, MA, May 1986.

Bonoma, T.V. "Case Research in Marketing Opportunities: Problems and a Process," *Journal of Marketing Research* (22), May 1985, pp. 199–208.

Culnan, M.J. "Mapping the Intellectual Structure of MIS, 1980–1985: A Co-Citation Analysis," *MIS Quarterly* (11:3), September 1987, pp. 341–353.

Cyert, R. and March, J. *Behavioral Theory of the Firm*, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1963.

Feigenbaum, E., McCorduck, P. and Nii, H.P. *The Rise of the Expert Company*, Times Books, New York, NY, 1988.

Fry, L.W. "Technology-Structure Research: Three Critical Issues," *Academy of Management Journal* (25:3), 1982, pp. 532–552.

Galbraith, J. *Designing Complex Organizations*, Addison-Wesley, Reading, MA, 1973.

Galbraith, J. *Organization Design*, Addison-Wesley, Reading, MA, 1977.

Harmon, P. and King, D. *Expert Systems: Intelligence in Business*, John Wiley and Sons, Inc., New York, NY, 1985.

Harmon, P. and Rex Maus, W.M. *Expert Systems Tools and Applications*, John Wiley and Sons, Inc., New York, NY, 1988.

Hayes-Roth, F., Waterman, D.A. and Lenat, D.B. *Building Expert Systems*, Addison-Wesley, Reading, MA, 1983.

Huber, G.P. "The Nature and Design of Post-Industrial Organizations," *Management Science* (30:8), August 1984, pp. 428–451.

Kling, R. "Social Analyses of Computing: Theoretical Perspectives in Recent Research," *Computing Surveys* (12:1), March 1980, pp. 61–110.

Knuth, D.E. and Moore, R.W. "An Analysis of Alpha-Beta Pruning," *Artificial Intelligence* (6:4), 1975, pp. 293–326.

Macy, J., Jr., Christie, L.S. and Luce, R.D. "Coding Noise in a Task-Oriented Group," *Journal of Abnormal and Social Psychology* (48), 1953, pp. 401–409.

March, J.G. and Simon, J.A. *Organizations*, Wiley, New York, NY, 1958.

Markus, M.L. and Robey, D. "Information Technology and Organization Change: Causal Structure in Theory and Research, *Management Science* (34:5), May 1988, pp. 583–598.

Minsky, M.L. "A Framework for Representing Knowledge," in *The Psychology of Computer Vision*, P. Winston (ed.), McGraw-Hill, New York, NY, 1975.

Mumford, E. "Managerial Expert Systems and Organizational Change: Some Critical Research Issues," in *Critical Issues in Informa-*

tion Systems Research, R.J. Boland, Jr. and R.A. Hirschheim (eds.), John Wiley & Sons, New York, NY, 1987, pp. 135–155.

Rousseau, D.M. "Issues of Level in Organizational Research: Multi-level and Cross-level Perspectives," in *Research in Organizational Behavior* (7), L.L. Cummings and Barry M. Staw (eds.), JAI Press, Inc., Greenwich, CT, 1985, pp. 1–37.

Sell, P.S. *Expert Systems—A Practical Introduction,* Wiley, New York, NY, 1985.

Shortliffe, E.H. *Computer-based Medical Consultations: MYCIN,* American Elsevier, New York, NY, 1976.

Simon, H.A. *Administrative Behavior,* The Free Press, New York, NY, 1945.

van Melle, W., Shortliffe, E.H. and Buchanan, B.G. "Emycin: A Domain-independent System that Aids in Constructing Knowledge-based Consultation Programs," in *Machine Intelligence, Infotech State of the Art Report 9,* No. 3, 1981.

Winston, P.H. and Prendergast, K.A. (eds). *The AI Business,* The MIT Press., Cambridge, MA, 1984.

Zack, M.H. and McKenney, J. "Organizational Information Processing and Work Group Effectiveness," working paper, #89-054, Harvard Business School, Boston, MA, 1989.

Acknowledgments

The author would like to thank Jim McKenney, Dorothy Leonard-Barton, Randy Davis, Izak Benbasat, and Donna Rosenbauer for their help in this work. He would also like to thank the Division of Research at Harvard Business School for support of this research.

About the Author

John J. Sviokla is assistant professor at the Harvard Business School. He received his B.A. from Harvard College and his M.B.A. and D.B.A. in management information systems from Harvard University. He has investigated issues in the strategic use of IT, IT planning, technology implementation and evaluation of IT projects. His current research focuses on the impacts of information technology on the management of expertise in organizations. Dr. Sviokla, whose articles have appeared in *DATA BASE* and the *Harvard Business Review,* has served as an invited speaker and presented papers at international conferences on information systems and is a member of the National Academy of Management and the Institute for Management Sciences. He is an active consultant and teaches regularly in corporate executive programs.