AI Magazine Volume 7 Number 5 (1986) (© AAAI)

Frank Lynch, Charles Marshall, Dennis O'Connor, and Mike Kiskiel II

AI in Manufacturing at Digital

Abstract The rapid advances in information technology are causing a fundamental change in the way we do our business. Within our manufacturing businesses today, various parts of the organization are "reasoning" about "engineered products." The everyday problem-solving activity within the organization can be thought of as conducted by a network of experts knowledgeable about the products and the physical and paperwork processes that constitute the business, that is, the *knowledge network*. The focus of our attention has not been just at the factory level; we have been addressing the order-process cycle: marketing, sales, order administration, manufacturing, distribution, and field service. This cycle can be thought of as the outer loop of the knowledge network Also, we recently began addressing the inner loop. This loop is the *product life cycle* marketing and new product requirements, design and manufacturing startup, and volume or steady-state manufacturing This article describes DEC's internal strategy for applying artificial intelligence (AI) to manufacturing processes and problems above the work-cell level. In addition to an overview of this knowledge network, we feature DEC's newest system in order processing: the configuration-dependent sourcing (CDS) expert Project experience on this system, which deals with the assignment of fulfillment sites (factories) to line items in computer system orders, is also described

A Broadened Perspective of Manufacturing: The Knowledge Network

In order to form a vision and a strategy, we took a broad new look at our manufacturing business. The perspective ranged from the customer at the point of sale through point of manufacture and point of distribution and back to the customer. In 1981 DEC coined the term knowledge network to represent this notion (O'Connor 1984) (see figure 1).

In many of these "pockets of expertise," within DEC or any other manufacturing business, the expertise and the reasons for making decisions are generally undocumented or are unavailable to all the parties needing the information.

Two Views of the Business

Within the knowledge network two major cycles are apparent: the order-process cycle and the product life cycle

The order-process cycle (see figure 2) is oriented around taking, manufacturing, delivering, and servicing an order.

The product life cycle (see figure 3) is oriented around products as opposed to orders (for example, the evolution of a new product and its manufacturing process).

What Has Been Going on at DEC

DEC's internal involvement in AI applications started in 1978 with the collaborative research done with Carnegie-Mellon University (CMU) on the now well-known R1/ XCON configuration system.

A little over two years ago, DEC collocated their major AI operations in one physical location in Hudson, Massachusetts, naming this facility, a collection of over 175 scientists, engineers, and business persons, the Artificial Intelligence Technology Center (AITC). The three main activities at the center are manufacturing and engineering applications, AI product development (VAX Lisp, VAX OPS, and AI VAX-Station), and marketing of Digital AI products and services. One of the major visions of the applications effort is the development of the knowledge network.

To realize this vision, we have been examining the various parts of the network in order to understand what people reason about when they perform their tasks, what information they need, and what they pass on to other parts of the network. What problem-solving approaches do people use in carrying out their decision making, and what are the requirements for integrating AI solutions into the existing business systems and processes?

The following subsections contain examples of how DEC is realizing the knowledge network. Some of the systems prsented are in production, some in the developmental stage, and some still in the research stage. We focus especially on a key application in the order-process cycle-deciding which factories will deliver on which line item, or CDS.

Frank Lynch is the group engineering manager of the Advanced Systems and Tools Group, Artificial Intelligence Technology Center, Digital Equipment Corporation (DEC), 77 Reed Road HLO2-3, Hudson, Massachusetts 01749 Charles Marshall is principal engineer, and Dennis O'Connor is senior group engineering manager of the Intelligent Systems Technologies Group at DEC Mike Kiskiel is manager of AI Order-Process Applications at DEC

Trademark Information The following are trademarks of Digital Equipment Corporation: VAX, VAX-11/780, VMS, XCON, and XSEL

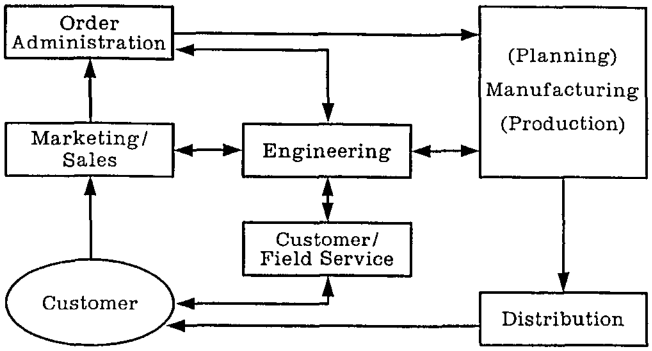

Figure 1. Knowledge Network.

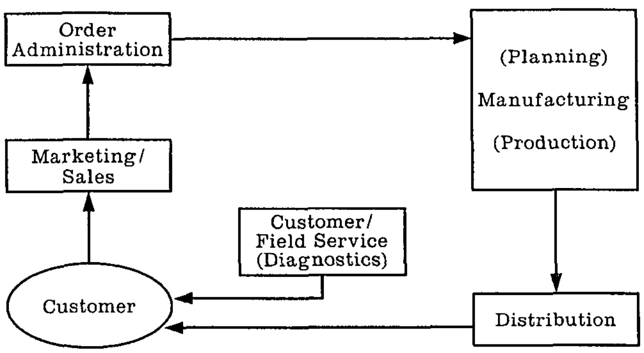

Figure 2 The Order-Process Cycle.

The Order Cycle

Until recently, DEC has concentrated its development activities on the outer loop of the knowledge network—the order cycle.

All the systems described in the order cycle were developed and built using the VAX-11/780 system running under the VMS operating system. They were all written in VAX OPS5. The prototypes of these systems were developed jointly with CMU under the direction of Dr. John McDermott.

Requirements, Specification, and Design Requirements, specification, and design are classic parts of a product design process (in our case, configuring preengineered components into a custom system design). Getting the requirements straight (XSIZE), taking an order that is complete and manufacturable (XSEL), and placing an order into production planning which can be built (XCON) all involve much uncodified expertise and reasoning about the product and its intended use.

XSIZE involves determining a customer's requirements. This problem poses the question of how one maps an understanding of a customer's planned usage onto, perhaps, *plan fragments* that describe and satisfy the requirements and situations. This research, sponsored by DEC, is being conducted at CMU.

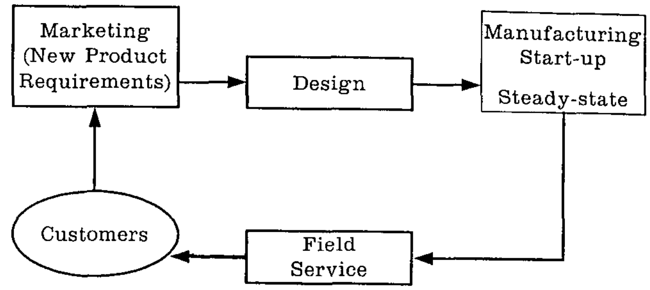

Figure 3. The Product Life Cycle.

XSEL is a computer sales assistant. It interactively refines the specification of the order to deduce function, to completely define a configurable system, and to map a generic DEC system into a specific order. Over 25% of the sales force has been trained in the use of XSEL. Within the past year, XSEL has been used to generate quotes for more than 15% of all VAX system orders. It is planned for use in the customer dialogue for every order.

XCON is an expert configuration engineering tool. It is the world's largest expert system in use for production. It contains over 5000 rules (previously undocumented design rules) plus a data base of over 20,000 components. XCON is used daily in production in U.S. and European plants, to assemble virtually all major DEC systems.

A typical order deals with 50 to 250 components and 25 to 125 pieces of information about each component, and typically, 1000 steps are necessary to configure a system. XCON accounts for such things as cable-length constraints, sizes and layout of cabinets, and power requirements (Bachant 1984).

Intelligent Manufacturing IMACS, an intelligent management system for computer system manufacturing, is focused within the four-wall manufacturing environment. It is a set of cooperating expert systems dealing with such things as paperwork management, capacity planning, floor-loading throughput management, and inventory management. IMACS was a very successful experiment, and some of its modules are being used on the factory floor today. One module, DISPATCHER, dispatches parts to the manufacturing floor; the other module, MATCHER, matches a final build order to a new order when an older order is canceled (thus avoiding the cost of tearing down or stocking the canceled order.)

Order Logistics ILOG is an intelligent logistics assistant. It is designed to capture the general knowledge that is used to manage the current order-flow logistics. It is meant to assist management by handling exceptions to the process, from receipt of order in manufacturing to arrival of goods at the customer's site. It deals with problems before and after the four-wall factory setting. The system develops and evaluates alternative plans, such as partitioning of orders to various plants (using limited XCON information), transportation management, and distribution plan management (Orciuch 1984). ILOG successfully passed the experimental stage and is now being incrementally developed and introduced into manufacturing in the form of CDS, a system for managing the logistics of orders with highly configuration-dependent manufacturing requirements

CDS Within scheduling, there is a task called *sourcing,* which is the selection of a factory or warehouse, a *source,* to fulfill a line item in an order. This task is done by production control (PC) personnel, who are the experts and users in this domain.

The CDS project is the development of a program to replace the sourcing module in the traditional order-scheduling system. The key difference between the CDS modules and the traditional modules is that CDS derives which line item or portions of a line item must be built together in the same factory from a connection to the configuration knowledge rules in XCON. In addition, CDS uses heuristic information that was acquired from PC personnel about how they do their work in sourcing (for example, parts with a request date different from the system request date are sourced as spare parts). The chief advantage of CDS is correct source locations assigned to all the line items of an order.

Before CDS Sourcing prior to the CDS project was a semiautomatic process. A program used lookup tables for preferred source locations for parts and assigned those locations to the line items of an order. It is important to note that this approach viewed each order line item independently, as though a particular line item had no relationship to any other line item. This element is only true for some DEC products; many line items in other products or system orders are by convenience or necessity grouped together at the factory or warehouse. This interrelation of line items is complex because only a portion of some line items need to be grouped with other line items. For example, a line item might contain three of DHU11-AP, a communications-related part. It is feasible that one DHU11-AP should be installed in the CPU, one DHU11-AP should be installed in the UNIBUS expander cab, and one DHU11-AP should be considered a spare part. If this were true, however, the CPU factory, the expander cabinet factory, and the spares warehouse would each be required to fulfill part of the line item. This fulfillment results in three sources for one line item. To complicate matters, the split of CPU and expander cabinet is determined by much how space is available in each cabinet, that is, what else has been ordered on the system.

The old method sourced the line items in the order with 75 to 80% accuracy. If an arbitrary line item had an 0.8 likelihood of being correct, then a whole order consisting of 30 line items had an 0.8**30 = 0.00123, or one in a thousand, chance of being correct. In practice, when following the old method, virtually all orders needed to be fixed.

How CDS Works Because some line items are related to other line items, CDS looks at the whole order instead of a line item at a time. Some of the relations, or rules, are based on technical configuration knowledge. Thus, in practice, parts appear integrated or built into other parts. For instance, memory modules are installed and tested inside the CPU cabinet. Because XCON contains technical configuration knowledge, CDS has a link to XCON for that information. For this category of relations, CDS is almost product independent because product dependencies are maintained in XCON.

Other relations concern special packaging "configurations." One example is a terminal that is composed of four to eight parts, with each part packed in a separate box. The exact number of parts is specified in the line items on the order. One combination of six parts is so popular that if the combination occurs, even as parts of other line items, then the six boxes are packed into another box at a location that specializes in shipping these particular six parts as one part. Another example is MICROVAX systems, which are completely configured by XCON, including which modules go where in the cabinet. For shipment purposes, however, some of the modules are packed separately and shipped with a standard MICROVAX package for final assembly of the "custom" system by the customer. Thus, some XCON-related information needs to be ignored. All of these packaging relations are represented and encoded in CDS.

Before the actual sourcing of the order begins, all of the relations of line items or parts of line items with other line items are completely developed for both technical and packaging configurations. This is the bulk of the expertise used by the production controllers. This process is called *order partitioning.*

After partitioning, preferred sources are assigned first for whole partitions, then for parts without partitions, then for parts that are excepted from partitions. Any parts unsourced at this point are sourced using the rules of thumb outlined earlier (for example, parts without explicit sources whose part numbers begin with Q are sourced as software, and so on).

Performance CDS and XCON run on a VAX 8650 computer as a node in a VMS cluster in the order-scheduling system. A system order with 30 to 50 line items sources in about 20 to 25 seconds. XCON configures orders at an average of 45 seconds per order. The old sourcing program sourced a system order in about 6 seconds. Obviously, it requires more computer time to source with CDS.

Process performance is the important consideration in implementing any system. Seventy percent of the orders can proceed through the sourcing step automatically without manual intervention. Before CDS, the figure was 5%. A configuration-related sourcing error once took 10 to 15 minutes to correct. CDS has reduced that time to near zero for 70% of the orders.

Quality of the sourcing process is another measure of performance. CDS sourcing is 99% correct with respect to configuration-related sourcing decisions In the old process, some errors in sourcing were not discovered, and some errors were introduced through the manual work done on the orders. Parallel tests demonstrated that sourcing problems which would otherwise be caught later in the process and corrected at greater expense are now solved by CDS. CDS is a cleaner process.

CDS Implementation At the beginning, PC groups pointed out problems in the sourcing area due to product configuration relations. The business trend was toward more configuration-related problems. It was determined that something like CDS was technically feasible and that a forward-chaining, rule-based system was right for the application. VAX OPS5 was chosen as the development language. CDS was prototyped in about two months and demonstrated to PC personnel. This demonstration clearly showed the feasibility of CDS to key people in production control and resulted in a legitimate project with a sponsor, program manager, funding, and key user representatives.

Working with the PC representatives, the project members specified the new sourcing process, developed a project plan, and continued to enhance the CDS prototype. Also, new functions were added to XCON to assist the CDS prototype. At the project midpoint, it was possible to do some computer performance measurements using the real order database to confirm expectations about hardware requirements. Because it has capacity for growth, VAX 8650 was selected to run XCON and CDS together.

The first implementation was cautious. This point in the order flow is very important to the business, and *fail soft,* the ability to have failure that will not disrupt the daily operation, is a requirement. Fail soft was implemented in three ways: First, CDS runs after the old sourcing program, which is still in the order flow; second, CDS can run with or without configuration information (even without configuration information CDS would still be able to source a little better than the old program); third, PC personnel have a process and tools to use if a fail-soft situation should occur.

The CDS system, which represents both the problem-solving process and the decision heuristics for sourcing, is encoded in a set of modules written in both traditional and AI languages and integrated into the traditional order-processing software system and the XCON configuration system. Approximately 125 VAXOPS rules are used to capture the sourcing heuristics. In addition, the more than 5000 XCON rules pertinent to the particular order being processed are activated to obtain technical configuration knowledge. Approximately 25 more rules define the sourcing process. The CDS rules and the XCON system are integrated with code written in BASIC to interface with the previously existing traditional order-process system software. The integration was greatly simplified by having languages that were all members of the VMS software family with full call-in and call-out capabilities.

The second implementation will be bolder in that CDS will be the only sourcing function. In addition, the order database will have a change in representation to cleanly support CDS function.

Current Status The first implementation of CDS was recently installed and run in the order-scheduling system, but there is some follow-up work to be done. Aside from the second implementation plan mentioned earlier, follow-up consists mostly of some additional heuristic sourcing decisions (such as new MICROVAX exceptions and new complex terminal products) that the PC groups will decide after extended use of CDS.

Learning Experiences In prior attempts at order processing, AI projects have not gone as smoothly as CDS has. The CDS project has followed what has become a straightforward project recipe for DEC, a recipe that forms the basis for DEC's "Guide to Expert Systems Program Management" (ISIG 1985). CDS has a sponsor (the PC manager with the business goals that CDS addresses). CDS has a program manager who is charged with total project responsibility and who reports to the sponsor. Users provide direction, functional definition, experts, and the development ownership. DEC has found that if one of these project roles is missing or misplaced, project productivity is low or nonexistent.

During the prototype stage, a development team emphasized functionality. Speed of execution was a secondary consideration. The most essential effort, however, was defining what the business process should be and what CDS should do. The people-software interaction during the prototyping and demonstrations helped the process design as well as the software design.

Productivity could have been improved in one area. The manufacturing information systems (MIS) group for PC are the current owners of CDS. However, programmers from the owner group were not allocated until CDS was well into the integration stage, and final implementation was substantially delayed. It would have been productive to include the owner programmers in the development of the prototype and demonstrations.

Realizing the Knowledge Network in the Order-Process Cycle

The earlier examples illustrate that our vision of collaborative expert systems is becoming a reality. Knowledge cap-

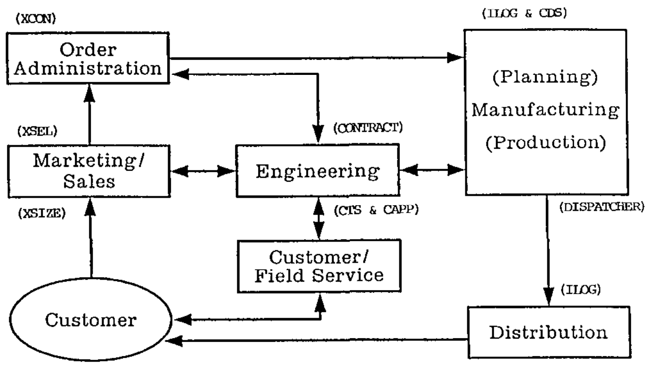

Figure 4 VAX-Based Knowledge Network.

tured in the order-administration technical edit stage (XCON configuration rules) has been shared with additional expertise both upstream (XSEL) and downstream (CDS) from this part of the cycle.

The Next Step: The Product Life Cycle

Recently, DEC turned its attention to the product life cycle. Although we have been sponsoring research in this area for several years, we have just begun to capitalize on our investment here.

The effort in this area is in the conceptual, research, and early development stages. The projects are all being developed on the VAX 11/780 system running under VMS. The languages being used are VAX Lisp and a commercial derivative of CMU's schema representation language (SRL). Most of the SRL-related work started as sponsored research and joint development at CMU under the direction of Dr. Mark Fox (Sathi 1985).

In the product life cycle, developing product requirements, specifying and designing new products, managing the transition into manufacturing, planning the manufacturing processes, and managing projects all involve expertise and representation of complex entities. We are actively pursuing advanced development and sponsoring research in the following areas: product description handoff to manufacturing with the configuration tracking system (CTS), computer-aided process planning (CAPP), and project management with the commitment negotiating and tracking tool (CONTRACT) (Marca 1985).

Status of the Knowledge Network

Many systems in varied stages from production use to research concept were described. A sketch of the knowledge network (see figure 4) indicates these projects and systems superimposed on the network.

Conclusion

Our vision at DEC is the evolutionary integration of expert systems that make up the knowledge network. We have made significant progress in the outer cycle, and our plan is to have the knowledge network evolve to meet business needs in manufacturing. The information technology is here, but to make it happen the evolution requires an understanding of, and an integration with, existing traditional business systems. These new AI systems are designed to co-reside with the excellent traditional business systems and to integrate across the DEC network architecture.

Our technology strategy has been and continues to be the support of university research for fundamental technology, in-house advanced development to transfer the technology to DEC, training of domain and software engineers in knowledge engineering, and the internal development by these knowledge engineers of real-world applications integrated into existing traditional applications.

Our business strategy is to increase indirect labor manufacturing productivity, to help manage inventory assets, to decrease cycle time within the operations, to increase customer satisfaction with products matched to customer needs and delivered in a timely fashion, and to decrease the time from new product conception to product shipment to the customer.

Our human resource strategy is to train and leverage DEC people assets in the new AI technology and to include the three dimensions of technology, business focus, and organizational dynamics in our training and development program (Clanon 1986). Without an overall strategy that considers these three perspectives, AI can not deliver solutions to our manufacturing-related problems and can not help in creating new opportunities which allow DEC manufacturing to attain a competitive advantage. We have seen the positive impacts AI can make in the order-process cycle of manufacturing and expect dramatic results as we attack problems in the product life cycle, all of which is gradually leading us to a computer-integrated business via the knowledge network.

References

Bachant, J , and McDermott, J 1984 R1 Revisited: Four Years in the Trenches *AI Magazine* 5(3): 21-32

Clanon, J M 1986 How Digital Trains Employees to Design AI Systems *Electronic Engineering Times*, 3 March 1986

Intelligent Systems Technologies Group. 1985 Guide to Expert Systems Program Management Artificial Intelligence Guide Series, Digital Equipment Corporation

Marca, D , and Cashman, P 1985 Towards Specifying Procedural Aspects of Cooperative Work Paper presented at Third International Workshop on Software Specification

O'Connor, D 1984 Using Expert Systems to Manage Change and Complexity in Manufacturing In *Artificial Intelligence Applications in Business*, 149-157 Norwood, N J : ABLEX

Orciuch, E , and Frost, J. 1984 ISA: Intelligent Scheduling Assistant Paper presented at IEEE AI Applications Conference

Sathi, A ; Fox, M ; and Greenberg, M 1985 A Representation Language for Activity Knowledge in Project Management In Pattern Analysis and Machine Intelligence IEEE Transactions