

Jeffrey Stone

## The AAAI-86 Conference Exhibits: New Directions for Commercial AI

The annual conference of the American Association for Artificial Intelligence (AAAI) is the premier U.S. gathering for artificial intelligence (AI) theoreticians and practitioners. On the commercial side, AAAI is the only event with a comprehensive exhibition that includes most significant U.S. vendors of AI products and services. In 1986 some 5100 people attended AAAI—a very good showing considering that the 1985 International Joint Conference on Artificial Intelligence (IJCAI) drew about the same number of people even with its substantial international support.

The commercial exhibits at AAAI-86 (110 exhibitors; 80,000 square feet) gave us the opportunity to take a snapshot of an industry in transition. What I saw was a dramatic increase in the commercialization of AI technology and a decrease in the mystique, smoke, and hype. A preliminary tour of the AAAI-86 exhibits indicated that participants could expect substantial changes from the situation at IJCAI-85.

At the 1985 conference, many vendors showed off their AI software and hardware products through demonstrations of simple application prototypes and games. At AAAI-86, however, product demonstrations showed real applications. Some vendors, including Intellicorp and Teknowledge, held private application forums, giving visitors a look at a variety of successful applications developed with their tools.

To match the new realism of the vendor exhibits, a new kind of user visited these exhibits. According to several exhibitors, the 1985 conference attendee was a corporate staffer, a person looking for tools to build prototypes that would demonstrate the value of AI technology. The 1986, user was interested in building operational systems that bring a return on an investment.

The bottom line is clear: Corporate America has been sold on AI and now expects to use and benefit from the technology. Various AI market watchers have estimated the percentage of Fortune 500 firms making significant investments

in AI to be from 40 to 80%. Regardless of the exact figure, I believe that the majority of large U.S. organizations are currently preparing to put AI technology into operation or have already done so.

### IBM Unveils Its AI Plans

As if to testify that AI is now legitimate technology in corporate America, IBM chose AAAI-86 to unfurl its dedication to AI. In the keynote address, Herb Schorr, IBM's AI czar, presented IBM's recently organized program for AI activities.

Schorr leads IBM's AI Project Office, a new type of IBM organization that will permeate all IBM activities and organizations related to AI. Schorr's organization currently consists of 12 people with full responsibility within IBM for (1) developing AI products for internal and external use, (2) marketing AI products outside IBM, and (3) applying AI tools and technology within IBM. The company currently has 6 expert system applications in operation and over 60 under development, including computer system work-flow management, a mass storage system test, customer system relocation cost estimation, S/38 capacity planning, a business procedure adviser, and software fix coordination.

The AI Project Office reports to a steering committee of senior IBM management that meets monthly. This kind of access to top IBM management is rare within the corporation and is indicative of IBM's desire to make headway in AI technology, application, and markets.

Underscoring the company's interest in AI was the announcement at the conference of a joint study agreement with Carnegie-Mellon University (CMU) under which IBM will provide CMU with 225 IBM RT PCs, valued at \$5.5 million, to "foster the advance of AI." Aside from the cooperation with CMU, IBM is participating in 52 other AI studies with 33 universities for a commitment of \$8.7 million over several years.

To help customers use IBM AI technologies, IBM has established an expert system marketing center in Princeton, New Jersey, and a support center in Cambridge, Massachusetts. The company maintains an internal AI support center

---

Jeffrey Stone is publisher and editor of *Professional Computing Highlights*, a monthly newsletter. His business address is: 331 Barton Place, Menlo Park, California 94025. The opinions expressed here are his own.

in Palo Alto, California, that trains IBM staff, offers AI support services, and provides consultation to internal groups. When you consider the scope of his charter and the size and makeup of the IBM organization, Herb Schorr will have his hands full in bringing AI to IBM.

### Toward Standard Computing Environments

Returning to the vendor exhibition floor, I found the most obvious trend among the expert system tool vendors was their continued movement toward standard computing environments. Users might be willing to pay the price for a Lisp machine as a super expert system development environment, but they are clear on requiring integration of applications with standard information systems. Users consider:

**Mandatory** Delivery of AI applications on standard corporate computing systems—primarily personal computers (PCs) and mainframe computers,

**Very Desirable** Delivery of AI applications in standard corporate computing languages—such as C (the current leader),

**For the Future** Ability to embed AI technology in existing and new corporate applications.

The advantages of delivering applications on standard computing systems are obvious: AI can be linked to corporate data, and corporate investments in equipment and experience can be reused. The advantages of delivering in a language such as C are efficiency in execution speed and memory utilization. These efficiencies, enable expert system applications to run on corporate mainframe computers and minicomputers with considerably reduced effect on other applications. The same expert system applications running in a Lisp environment on the same hosts are often completely unfeasible because of the system resources they consume.

Of the four leading expert system tool vendors, three are supporting or will support C delivery.

- Teknowledge has switched to C and no longer ships Lisp-based products.
- Inference demonstrated a C version of their Automated Reasoning Tool (ART) at AAI-86.
- The Carnegie Group has announced C delivery versions of KnowledgeCraft will be available in late 1987.

Only market leader IntelliCorp is holding firm with Lisp. At AAI IntelliCorp president Kehler commented on the Lisp-C question: "It is very important to go after high-value applications. Speed is secondary to giving high value." He also admitted that IntelliCorp was working with the Kyoto Common Lisp, a Common Lisp system that cross-compiles Lisp to C.

### Embedability

If the hot buzzword in 1985 was "C," in 1986 it must have been *embedability*. Embedability refers to the ability of expert system tools to be built into user applications to provide

these applications with the advantages of knowledge-based operation. This technique means bringing AI technology to the standard corporate applications: payroll, decision support, management information systems, project management, factory automation, and the rest.

Eventually, the use of knowledge bases will be similar to the use of databases. Vendors will provide knowledge base maintenance and administration systems as stand-alone tools (just as the database management system [DBMS] vendors do today) and will offer user-application developers the facilities to embed knowledge base access in their own applications (just as the DBMS vendors provide user-application access to databases today). The facilities for embedding knowledge base access in applications will include subroutine libraries and host language support. Several expert system tool vendors with whom I discussed embedability indicated that they believed the market for "embedable" tools would eventually be as large as the market for stand-alone expert system tools, if not more so.

On the conference exhibit floor, almost every vendor was ready to talk about embedability. It turned out, however, that talk was much easier to find than expert system tools which support true subroutine-type interfaces. Many tools can be invoked from user applications, but most do not provide closed subroutine linkage that enables the user application to have complete control over what happens, particularly over the interface between the end user and the application.

Teknowledge, consistently a leader in integrating expert system technology into mainstream computing, once again assumed a leadership role with the announcement of its Copernicus Open Architecture. Copernicus will be a set of 100% embedable knowledge engineering subroutine libraries that user organizations can link into their applications. Teknowledge called the Copernicus announcement a "statement of direction" and provided little technical detail. Nonetheless, company officials said that Copernicus is not just smoke; Teknowledge tools S.1 and M.1 have themselves been largely constructed from modular subroutine libraries that will be the foundation of Copernicus products.

Another company, Software A&E (Arlington, Virginia), has an interesting expert system shell that is said to be fully embedable. Their Knowledge Engineering System (KES) received considerable attention at AAI-86 primarily because of its support for three inference techniques: production rules, hypothesis and test, and Bayesian (statistical). KES hypothesis and test capabilities enable expert system developers to describe objects in terms of their characteristics (for example, a medical disorder in terms of its symptoms) and are particularly useful for classification problems that make up a large part of all expert system applications.

An important element of KES is its simple developer interface, which Software A&E designed to enable application domain experts, rather than computer or knowledge engineers, to develop applications. KES is coded in C and is

available on a variety of micro, mini, and mainframe computers, including MS-DOS PCs (\$4,000), UNIX, and DEC VAX VMS systems. Run-time version and volume discounts are offered. KES was the only AAAI-86 product that was said to be capable of being fully embeddable in the sense of having an interface callable from a closed subroutine.

### **One Cannot Live by Rules Alone: Object-Oriented Programming**

Software A&E is ahead of the crowd in several respects. First, it already supports true embeddability; second, the company has realized that pure rule-based systems might not be the best solution for many problems. This latter point was echoed by a number of other vendors as well. It seems that rule-based development works well when a problem can be divided into a number of areas or entities, each of which has a well-defined operational focus. If the relationships between the entities are not simple, however, then an object-oriented programming metaphor is often found to greatly simplify application development. Although high-end tools such as Intellicorp's knowledge engineering environment (KEE) have supported object programming for some time, it now seems that many more expert system tools will soon support this style of development.

At AAAI-86, we noted that Xerox's CommonLoops object-programming proposal was showing signs of real acceptance as a set of object-oriented programming facilities for Common Lisp. (CommonLoops was first proposed as an extension to Common Lisp by Xerox at IJCAI-85. An initial version was released to the Common Lisp community at no cost in February 1986.) Xerox has marshaled more Lisp vendors to support its CommonLoops, a set of object-oriented programming extensions to Common Lisp. Lisp vendors that announced their support for CommonLoops included Lucid; Franz; and Artificial Intelligence, Ltd. (Hertfordshire, United Kingdom). To date, CommonLoops has been ported to a number of Lisp implementations, including Digital Equipment Corporation's VAX Lisp, Franz Common Lisp, Hewlett-Packard Common Lisp, Kyoto Common Lisp, Lucid's Common Lisp, Spice Lisp, Symbolics Lisp, Texas Instruments Explorer Lisp, and Xerox Interlisp. Although PC Common Lisp vendor Gold Hill has yet to offer support for CommonLoops, the company did say that it was fully behind the Xerox proposal.

At the expert system shell level, Neuron Data (Palo Alto, California) and Gold Hill Computers introduced hybrid products that integrate object-oriented and rule-based programming. Neuron Data introduced an IBM PC-AT version (\$5,000) of its successful Nexpert Macintosh-based tool system (\$3,000). With its new PC version, called Nexpert Object, Neuron adds structured objects and a rich set of inheritance mechanisms to its original rule-based system. Nexpert Object runs on PC-AT-class machines under Micro-

soft Windows. In February 1987, Neuron released a version of Nexpert Object for DEC's MicroVax computers. The company also has a Macintosh version that provides all of the object-oriented features of Nexpert Object. Nexpert is coded in C.

Gold Hill Computers stepped out of its traditional role as a vendor of Lisp environments and announced an expert system building tool called Acorn. Acorn is a hybrid tool that combines frames, forward and backward chaining rules, and object programming into a single integrated system. Like Software A&E's KES, Acorn has been developed so that nonprogramming application-domain experts are able to create useful applications. Such nonprogramming users will generally utilize Acorn's highest level of facilities, which Gold Hill calls the SHELL. SHELL users can enter frames and rules as well as control and debug applications. A lower TOOLKIT level is intended for experienced developers and makes the knowledge representation language available to the developer. At the lowest level, GCLISP can be used. Acorn runs under Gold Hill's GCLISP 286 developer on PC-AT-class systems. Interfaces to dBASE, 1-2-3, and C will be provided. The product will sell for \$5,000 and is expected to ship in the first quarter of 1987.

Finally, even pure rule-based Teknowledge acknowledged that rules alone are not enough. As Copernicus-based products are introduced, future versions of Teknowledge products will include facilities to support object-oriented programming.

### **PC Expert System Tool Pricing**

Companies such as Software A&E, Neuron, and Gold Hill that have powerful expert system tools are pricing the PC versions of their products in the \$3000 to \$5000 range. Can this sort of pricing hold up for PC expert systems, or should we expect to see swiftly falling prices for these high-end PC products in 1987?

One clue about future pricing comes from Paperback Software (Berkeley, California). Paperback, the brainchild of Adam Osborne, is dedicated to publishing high-quality, low-cost software for the PC. Its products include the best-seller VP-Planner, a 1-2-3 work-alike, and VP-Info, a dBASE-compatible database manager. At AAAI-86, Paperback introduced its \$99 VP-Expert, an expert system tool that challenges the high-end PC products with its breadth, quality, and features. VP-Expert offers a rule-based system with an inductive front end for fast prototyping; rigorous confidence factor implementation; subscripted and indirect-reference variable addressing; database rules that apply across all database records; backward chaining and limited forward chaining; interfaces with popular PC spreadsheet and database files; floating-point arithmetic with trigonometric functions; a superior user interface featuring menus, color, and windowing; graphic representation of the knowl-

---

*. . . high priced vendors will soon have to decide whether to retain their premium pricing and settle for being niche vendors or to drop their prices considerably . . .*

---

edge base; and an internal text editor for knowledge base creation and maintenance (external text editors can also be used).

The advent of full-featured quality tools such as VP-Expert at reasonable PC software prices means that the high-priced vendors will soon have to decide whether to retain their premium pricing and settle for being niche vendors or to drop their prices considerably and address much broader markets. As some of the smoke and mystique surrounding AI technology clear, this decision will also be clear to the vendors. I believe that most will decide to drop their prices; in a year and a half, there will be few PC expert system tools selling for over \$1000.

### **The 80386**

In the hardware arena, the most significant AAI-86 development was Gold Hill's 80386-based Hummingboard, a PC XT or AT add-in board that runs GCLISP with up to 24MB of RAM. The board was developed jointly by Gold Hill and AI Architects (Cambridge, Massachusetts) and retails for \$7,000, including GCLISP software. For this \$7,000, you get a 16MHz Intel 80386; a 1MB on-board RAM that is expandable to 24MB with 1M-bit memory chips (currently to 6MB with 256K-bit chips); a 32-bit data path with 2K cache; memory sharing with the host PC processor using the Intel above-board standard; optional support for Intel 80286 or 80387 floating-point processors; a GCLISP interpreter and compiler; a GMACS editor; a GCLISP debugger; an online help system and the San Marco Lisp Explorer; and Gold Hill's 386 operating system environment.

When configured in an XT or AT, Gold Hill's Hummingboard runs at 3-4 million instructions per second (MIPS)—almost as fast as the new Sun III workstations with 25MHz 68020s and probably in the same ballpark as the current workstations from Symbolics. In addition to making the Hummingboard's super performance available to GCLISP, Gold Hill's operating system environment lets users run standard PC applications on the 386.

The Hummingboard is just the start of the 386 revolution. Compaq has already announced its 386-based AT-compatible system that sells for \$6,499, including 1MB RAM, a 40MB fast hard disk drive, and a 1.2MB diskette unit (display and display adapter extra). There is now a new crop of 386 boxes, most of which are able to run the full library of PC applications as well as very large AI applications at performance levels approximating that of the Symbolics workstations. The 386 is the last nail in the coffin for

Lisp machines as application-delivery vehicles. All that is left now for the Lisp machines is their powerful development environments, which are still unequaled by the portable Lisp environments available from such companies as Lucid and Gold Hill. In time, the gap between these two environments will also diminish.

### **VLSI Lisp Machine Implementations Are Coming**

The primary response of the traditional Lisp machine vendors to the erosion of their future markets has been the development of very large scale integration (VLSI) versions of their symbolic processors. Texas Instruments (TI) will provide the Defense Advanced Research Projects Agency (DARPA) with evaluation systems of its Compact Lisp Machine (CLM) in spring 1987. The system is built around a 550,000 transistor VLSI microprocessor implementation of its Explorer Lisp machine processor that contains 60% of the logic of that processor. According to TI, the CLM provides five times the power of current Explorer models and is small enough to fit in a shoe box. Systems such as TI's CLM promise future symbolic computing systems that could offer substantially increased performance over today's Lisp machine products, board-level integration with standard computing environments, and board-level pricing.

Ultimately, such products either in workstation packages or as add-in boards with Lisp software environments are likely to cost between \$10,000 and \$25,000 and will probably assume the high-end segment of the Lisp system market. The rest of the market will be based on portable Lisp environments running on PC workstations and other corporate computing environments. In single-board format, VLSI-based LISP machines might also open up new markets as embedded intelligent controllers and diagnostic devices. With their small size and high performance, these units will facilitate applications that currently are not feasible.

At AAI-86, TI was not the only vendor talking about VLSI implementations of Lisp processors. Xerox also publicly discussed its program for VLSI symbolic microprocessors. The Xerox Common Lisp Processor (CLP) is the first result of a five-year Xerox Palo Alto Research Center (PARC) project to build a family of high-performance, low-cost, custom VLSI-based Lisp systems. Xerox describes the processor as a "specialized RISC machine with the small, fixed instruction set and many fast registers of the classical RISC machine but with some quite complex execution sequences for a small number of instructions that are known to be critical to Lisp performance." The system features a tagged architecture with very tight processor-to-memory binding so that memory looks like more like medium-speed cache than a traditional main-memory system. Xerox estimates that the most conservative implementations of the architecture will achieve five times the performance of the current Xerox 1186 workstation at a lower cost. Xerox is looking for internal versions of CLP during the first half of 1987.

Lisp machine leader Symbolics has been quiet about VLSI processor developments since it mentioned some effort in this direction in its 1983 initial public offering literature. When questioned about its current direction in VLSI processors, the company replied only that its competitors' VLSI offerings seemed to be based on the original Lisp machine designs from MIT and that Symbolics is now several generations beyond such an architecture. Because the company stands to lose a great deal from preannouncements of such products, we can understand why Symbolics is loathe to discuss the matter. Nonetheless, the response I received to my questions indicates that Symbolics is not going to let its competition gain market share through superiority in VLSI.

Symbolics introduced new versions of several of its models that are based on the single-board VLSI gate-array processor implementation which the company introduced in early 1986 in its 3640AE product. The net result is 20 to 40% higher performance for Symbolics entry-level and midrange systems. Additionally, the new systems require much less space than earlier models. An entry-level system with 4MB RAM, 190MB disk, Ethernet interface, keyboard, mouse, and monochrome console is now priced at \$49,900. Symbolics reduced the minimum price of its high-end product from \$107,900 to \$98,900.

### A New Lisp Machine Vendor?

A new system vendor, Integrated Inference Machines (IIM) of Anaheim, California, exhibited its symbolic computing system at AAI-86. The company has developed a proprietary, 40-bit (32 data, 8 tag), pipelined, symbolic architecture with fast hardware tag decoding that it claims is four to five times faster than the industry standard Symbolic 3670 system. The Symbolics 3670 uses a 36-bit, nonpipelined architecture. An interesting facet of IIM's machine is its 80286-based I-O processor that operates independently of the symbolic processor. The I-O processor runs MS-DOS applications and supports network access through an Ethernet interface. In the future, IIM expects to run Microsoft Windows on the I-O processor, providing the user with a graphic front-end for symbolic processing applications, independent MS-DOS applications running in parallel with symbolic applications, and remote network terminal operation (windows across a network).

A basic IIM system, including 10MB RAM, a 33MB hard disk, a 5.25-inch diskette drive, serial-parallel ports, 1280 x 800 monochrome display, mouse, keyboard, Common Lisp interpreter-compiler, and EMACS-like text editor, will sell for \$40,000. The company planned to ship systems to Beta-test sites in February 1987. With life getting tougher for the established Lisp machine vendors who enjoy the advantages of owning mature development environments, I can only imagine that IIM will have difficulty finding a home in the Lisp system market with yet another proprietary architecture.

### Prolog-Lisp Co-Processor

While IIM was busy describing its challenge to the Lisp machine vendors, another new vendor, Xenologic, unveiled an *inferencing co-processor for Sun workstations*. Xenologic's X-1 inferencing engine is based on DARPA's Aquarius research into the use of fifth generation computer technology to develop AI supercomputers for inferencing and logic programming. The Aquarius project has been led by Dr. Al Despain at the University of California at Berkeley since 1983 and has focused on Prolog-oriented systems. Despain is also a founder of Xenologic, chairman of the Xenologic board, and director of research for the company.

Xenologic builds on Aquarius Prolog findings with the inclusion of Lisp-specific hardware enhancements to the basic inferencing mechanisms. According to Xenologic, the X-1's inferencing capabilities are far in advance of its competition:

**Table 1. Thousands of Logical Inferences per Second.**

Sun 2	14
VAX 11/780	15
PSICOT	30
DEC 2060	43
Symbolics 3675	53
Xenologic X-13	305

The company's first product is a pair of VME boards that implement the X-1 inferencing co-processor for Sun II/III workstations. Complete with Prolog compiler and assembler, the board set was planned to be available in January for \$30,000 (additional boards sets at \$20,000 a set). The company also planned Apollo and MicroVAX versions of the product and hoped to have its Lisp compiler ready by February.

Ultimately, Xenologic is shooting for a VLSI version of its processor that will be packaged as a single board co-processor to be available in late 1987. According to Richard Mirabella, Xenologic marketing director, this product will offer better performance and might be priced around \$10,000, including Prolog and Lisp software; VME and PC-AT (or AT-successor) versions are likely. If Xenologic can come through with most of what it promises, it is likely to find a lot of willing partners at its doorstep.

### Summary and Conclusions

The vendors and users participating in AAI-86 confirmed that AI has moved out of the research laboratories and is finding a home in mainstream corporate American life. Even IBM, which until now has maintained a low profile with respect to AI, unveiled an ambitious, wide-ranging program to promote the use of AI technology within its own ranks and by its customers.

AI technology vendors have been quick to respond to the requirements of business through the integration of AI software and hardware into standard information systems:

**Software** Delivery of applications on standard corporate computing platforms from PCs through mainframe computers,

Movement away from Lisp toward C and other languages that are commonly used by business concerns,

A growing capability to embed knowledge-based technology directly in horizontal and vertical business applications.

**Hardware** Development of VLSI versions of symbolic processing systems that will increase performance, lower price, and support tighter coupling with business information systems.

Along with the assimilation of AI technology into business information systems, we see a diminishing role for special-purpose symbolic architecture hardware. Although such systems currently enjoy an expanding market due to their powerful application-development environments, the only application-delivery systems that will be based on such hardware will be applications which constantly change, that is, perpetual prototypes. Eventually, the portable Lisp environments and the higher-level tools running in standard computing environments will consume the market now enjoyed by specialized hardware vendors.

The new generation of business workstations and office systems based on Intel's 80386 processor is not going to help the plight of the Lisp machines. As these systems appear over the next few years, they will be able to provide equiva-

lent power of today's symbolic workstations at a fraction of the cost. Lisp machine vendors will respond in a year or two with VLSI versions of the symbolic processors that will satisfy those applications requiring top-end performance. Offered in add-in-board form, these special-purpose systems will be configurable in standard business systems, thereby providing much higher levels of integration than are achieved with today's symbolic processors.

In the realm of expert system software tools, it appears that object-oriented programming of one variety or another will join rule-based development as a standard facility. Programming by rules alone might suffice for small problems, but for larger problems it is not efficient. This change and the normal evolution of technology means that PC expert system software is getting more sophisticated. Until now, full-function PC expert system shells have been priced as high as \$5,000—much higher than most other PC software. In the future, I expect to see prices of even the most sophisticated PC expert system tools falling under \$1,000.

AAAI-86 showed us that the basic ground rules for bringing knowledge-based systems to business have been established. Some of the smoke and the that has surrounded this transition is fading. We are now entering a period where it will be possible for real advances in AI technology to be seeing such advances—many of them very significant—during the next few years.



**Halbrecht Associates, Inc.**

*Executive Search Consultants*

## **SERVING THE ARTIFICIAL INTELLIGENCE COMMUNITY**

***Halbrecht Associates, founded in 1957, was the first search firm in the United States to specialize in Information Systems Management, Operations Research, Management Sciences, Telecommunications, Computer Sciences and Decision Support Systems. We are now also in the forefront of the Artificial Intelligence field. Because we have been so extensively involved in this emerging technology, our portfolio of people includes extensive numbers of AI professionals at all levels, from the United States and abroad.***

***Our client list includes companies coast to coast, large and small, including start-ups with equity opportunities. Accordingly, depending on your needs, we are in a unique position to assist you rapidly with any of your requirements.***

***Our current assignments range downward from several Chief Scientists, or those who can provide the technical leadership to head up major AI programs, to Senior LISP Programmer. If you would like to be considered to be part of our AI data base, please send us your credentials. This will enable us to give you first notice of appropriate opportunities. All inquiries are in strictest confidence.***

**Contact: Daryl Furno, Senior Associate**

**Herbert Halbrecht, President**

*1200 Summer Street • Stamford, CT 06905 • 203-327-5630*