ACE: AN EXPERT SYSTEM FOR TELEPHONE CABLE MAINTENANCE

Gregg T. Vesonder*, Salvatore J. Stolfo**, John E. Zielinski*,

Frederick D. Miller*, David H. Copp

* Bell Laboratories, Whippany, NJ.

** Columbia University, New York, NY

ABSTRACT

ACE, a system for Automated Cable Expertise, is a Knowledge-Based Expert System designed to provide trouble-shooting reports and management analyses for telephone cable maintenance. Design decisions faced during the construction of ACE were guided by recent successes in expert systems technology, most notably R1/XC0N, the Digital Equipment Corporation VAX configuration program. ACE departs from "standard" expert system architectures in its use of a conventional data base management system as its primary source of information. Its primary sources of knowledge are the users of the database system and primers on maintenance analysis strategies.

I. Introduction

In this paper we describe ACE, a knowledge-based system for Automated Cable Expertise. It was designed to provide support for directing telephone cable maintenance. The development of ACE, begun by Bell Telephone Laboratories two years ago, demonstrates that Artificial Intelligence techniques can be applied to significant and practical "real-world" problems in an efficient manner. The developers of ACE were able to implement, test, modify, and verify the performance of the system ahead of predicted schedules. Furthermore, ACE demonstrates a successful merger of two complementary yet independent technologies: database and knowledge base systems.

Unlike, for example, MYCIN [1] and M0LGEN[2] ACE is not a consultant. It is an automatic analysis System digesting hundreds of telephone cable maintenance reports daily. This data is provided by another computer system, CRAS (Cable Repair Administration System), a data management and report generation system. Telephone company engineers use CRAS to pinpoint trouble spots in the local telephone network, to predict future work force requirements and to drive plant rehabilitation decisions. ACE has subsumed some of these tasks and uses its knowledge of cable analysis to retrieve reports from CRAS and to provide "bottom line" information and recommendations concerning trouble spots.

Although the problem domain of ACE appears fundamentally different from that of R1/XC0N,[3] the Digital Equipment Corporation VAX computer configuration program reported by McDermott, there are enough common characteristics of both domains that suggested the organization of R1/XC0N is suitable for ACE's application. Consequently, many of the design decisions faced in the development of ACE were guided by the methods used in R1/XC0N.

Specifically, the ACE inference engine is a forward-chaining Production system executing the Match problem solving paradigm. The declarative nature of the 0PS4 production system notation[4] effected the rapid development of the ACE knowledge base. The team of knowledge engineers (cognitive psychologists), who interviewed several human experts, had little difficulty operationalizing the acquired knowledge directly in 0PS4. Subsequent cycles of test, debug and modify were carried out with relative ease. The end result was a working, useful prototype ahead of schedule. An advanced version of the system will be present in the telephone companies in the near future.

The following sections detail the problem domain and organization of ACE and describe how the development and installation of a prototype of this system in a "live" production environment was completed in a cost effective manner.

II. The Problem Domain

In normal operation the telephone company supports a telephone line from a residential or business site to telephone company switching machines in the central office. This line is called a cable pair. A collection of pairs are bundled to form the cables that hang from telephone poles or reside underground. A collection of cables form a wirecenter. These three levels form the bulk of the local telephone network and the cable maintenance force concentrate their efforts at all three levels. A more detailed description of the local telephone network can be found in a Bell System handbook.[5]

A variety of electrical faults and environmental conditions can cause failure in *one* or more cables or individual pairs. (Insect infestations in terminal boxes and gnawing rodents are perennial problems.) An important and expensive operation performed by the operating companies is the general maintenance and rehabilitation of these 1ines.

Customer generated maintenance reports provide information for identifying trouble spots within the local network. In a high density geographic area, the logging and tracking of failure reports has become an expensive and important data processing operation.

In order to identify trouble spots for repair and rehabilitation, many telephone companies use CRAS to monitor the repair activity of the local network on a daily basis. Trained analysts routinely examine impressive volumes of data and attempt to identify spots for maintenance to prevent further disruption of service to customers.

CRAS provides a set of report generating programs each producing a specialized summary of various aspects of customer generated and employee generated repair tasks. An individual record maintained by CRAS consists of numerous fields detailing a repair task reported by a customer or employee. This is *referred* to as a trouble in telephone company jargon. Many distinct CRAS records representing troubles may refer to the same pair or cable, indicating a potential chronic problem.

The large volume of detailed information permits the human expert to make an informed selection of candidates for rehabilitative maintenance. However, the limited number of specialists available, and the size of the database inhibits the timely analysis and reporting of persistent problem areas that require rehabilitation. The backlog delays the assessment of future work force needs and the selection of prospective areas for maintenance. The approach of installing an expert system as an adjunct to the CRAS system, assisting management decision making, was proposed as a solution to the problem of timely and accurate selection of areas for rehabilitation.

III. The System Organization

Knowledge-based expert systems have constructed, typically from two loosely coupled modules that form the problem solving engine. (see figure 1)

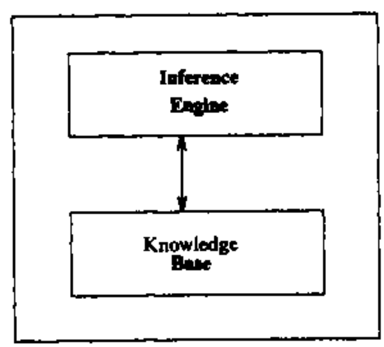

Figure 1: Organization of the Problem Solving Engine of a Typical Expert System.

The knowledge base contains all of the relevant domain-specific information permitting the program to act as a specialized problem solver. Within ACE the knowledge base is represented in rule form. The inference engine controls the deductive process. This control strategy is also represented in rule form within ACE.

In ACE a third component has been added: a database. Over the past decade database technology has progressed from dull-witted, efficient systems to systems offering abstract data models to facilitate semantically based retrieval functions. However, few, if any, offer the ability to deduce new data from old. At best aggregate data functions can be applied but deductions from raw data are not possible. On the other hand, Knowledge based expert systems have been implemented as consultants to specialists in narrow domains. These systems are designed as problem solvers. Expert systems rarely operate with a large amount of data.

ACE attempts to merge the two technologies and this merger may be useful in a variety of applications. Kellog[6] reports on the Knowledge Management I (KM-I) system and its organization is similar in scope to that represented by ACE. However, KM-I is designed as an intelligent front-end, interfacing the database to the user by an English-like query facility. This makes the database easier to use by making the database access language closer to the language of the application area.

In contrast, in ACE the expert system is the user! Tailored to solve a single class of problems, the knowledge base component of ACE automates the tasks the database supports. Thus the expert system not only answers the questions but also poses them.

A. Details of the Implementation

The knowledge base and inference engine of ACE are written in Franz LISP[7] and the 0PS4 Production System language running on a DEC VAX 11/780. Supporting routines for ACE, including CRAS interfacing and electronic mail facilities (to send ACE output to selected user mailboxes) are written in resident UNIX* software. The gross system level organization is shown in figure 2.

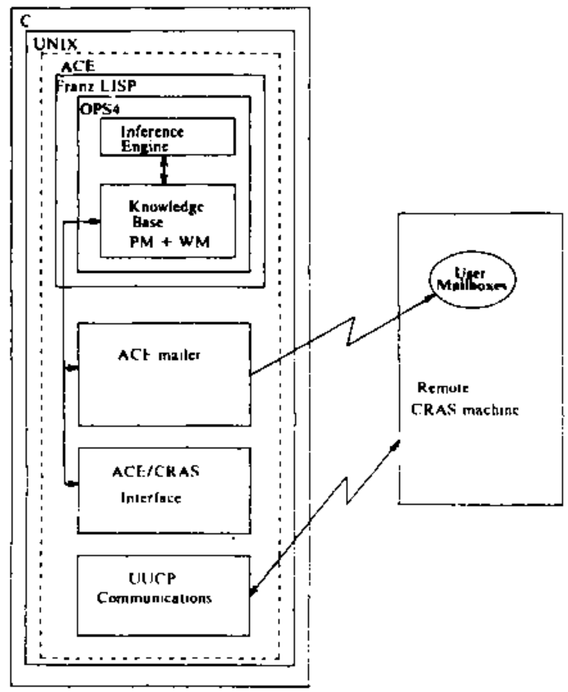

Figure 2: The Organization of ACE and Supporting Systems.

Note that a user is not necessary to initiate ACE on a daily basis. ACE is standalone, awakened nightly by UNIX to study the day's transactions. The only record an ACE user sees is an electronic mail message on trouble spots it discovered. (see figure 3)

* UNIX is a registered trademark of Bell Telephone laboratories.

From ace Thu Apr 14 15:42:27 1983

ACE 1.0 output message follows:

***************ACE REPEAT MESSAGE***************
In cable 47 wirecenter 999550 1 cable troubles were reported yesterday. There were 15 cable troubles in the cable in the past thirty days. 4 pairs were in the complement range from 600 to 700. 75 percent of the addresses were williams.

The most common disposition code was 0436. This indicates the presence of a problem at one or more crossboxes in this cable. In all likelihood it is one crossbox. This crossbox should be cleaned, rehabilitated and properly closed up, to prevent further troubles, or it should be replaced if it has deteriorated to a point where that is necessary.
The detail data for cable 47 is:

pair	disp	ccode	ctime	cttn	trob-address
627*	436	0	01/02	50034	17 williams
648*	436	0	01/29	50175	17 williams
655*	400	0	01/14	53021	17 williams
659*	436	0	01/30	50184	4 watson
838	436	0	01/11	50128	b box williams st
879	431	0	01/22	50161	austin @ kenwood
891	431	0	01/07	50113	p2 9 kenwood
1006	400	0	01/12	53010	17 williams st
1031	436	0	01/14	65002	87 pearl
1101	401	0	01/24	25036	pr18 upland
1509	431	0	01/08	50116	35 calender
1607	431	0	02/01	50193	calender 1st fr putn
1607	431	0	01/07	50110	131 huron ave
1695	433	0	01/08	25003	p@ 35 calendar
2513	401	0	01/18	25021	p 19 pleasant st

End of ACE transmission.

Figure 3: A Sample ACE Message. Produced by Real Data.

B. Details of the Problem Solving Engine

Within ACE, the body of Knowledge about wirecenters, CRAS data and commands and analysis strategies is embodied by an 0PS4 program.

1. *Production Systems* In 0PS4, data elements in Working Memory have the form of arbitrary LISP list structures. The user can also define other condition and action predicates, the former testing for a variety of conditions in Working Memory and the latter adding functions beyond modification of Working Memory. An English language equivalent of an ACE production rule is presented in figure 4.

IF a range of pairs within a
    cable have generated a large
    number of customer reports
ANDIF a majority of the work on those
    pairs was done in the terminal
    block
THEN look for a common address for those
    repairs.

Figure 4: English Language Equivalent of an ACE Production Rule.

2. *Characterizing the Problem Domain* The forward-chaining, data-driven approach typified by R1/XC0N was chosen for ACE since cable analysis is primarily a bottom-up, data-driven task.

Using the taxonomy of problem domains described in[8] we classify our domain in the following way: (1) The data are temporally based. The analyses produced by ACE are dependent on the frequencies of failures over time. Thus the analyses done by ACE are data-driven. (2) The data are reliable for the task. The CRAS data base provides most of the necessary information about the cable plant. (3) The knowledge of the domain is reliable. (4) The search of the CRAS database is exhaustive, but the reliability of the knowledge permits significant pruning of the space of possible conclusions. (5) The main focus of the problem-solver is the quality of the inferences it makes. The ACE design was driven by the necessity of producing the same analyses as its human counterparts.

Because of these characteristics the Match [9] prbblem-solving strategy was chosen for ACE. As in R1/XC0N, ACE requires no backtracking search of a large problem space, rather its primary task is divided into a fixed sequence of subproblems. These sets of productions are described in the next section.

3. *The structure of the Knowledge Base* Although no structure is provided (or imposed) by 0PS4, the set of about 100 productions and 50 related condition and action functions in ACE's knowledge base can loosely into sets of related rules that collectively do the analysis.

A set of productions does short term analysis by daily examining the flow of maintenance reports. If troubles are reported for a cable that has no previous history of chronic troubles then information is retained that indicates the cable may soon require attention. When new failures are reported for a cable with a history of persistent problems, ACE requests further detailed reports from CRAS. This information is used to deduce: (1) whether the repair task done on that cable suggests that preventive maintenance will reduce future repairs; (2) if preventive maintenance is required then what type is likely to be effective; and (3) if possible, where the rehabilitation should be done. Thus ACE not only identifies trouble spots, but also suggests how to repair them.

Each of these aspects of short term analysis requires substantial deductive power. For example, locating the physical spot where rehabilitation should be done is a difficult problem because the employee reporting the repair task is permitted to enter the site of the failure in a free textual format. Subsequently, the CRAS record of the task may contain an entry that not only may be inaccurate but also subject to typos and capricious abbreviations. As an example, "WASH 5" might refer to the same location as "WASHINGTON AND FIFTH" or "CALENDAR" might refer to the same location as "CALANDER STREET." Thus ACE contains productions and associated LISP functions encoding heuristics to estimate whether several addresses refer to the same general location.

The main sources of knowledge for the short term analysis were: (1) textbook knowledge from primers on telephone cable analysis; (2) expert advise from the developers of CRAS; (3) expert advice from theoreticians of cable analysis both in Bell Telephone Laboratories and in the local operating companies; and (4) local analysts from the operating companies and users of CRAS performing the actual analyses.

Another portion of the ACE PM contains a set of productions that know how to access CRAS. Based on requests for additional information generated by other productions, these productions assemble CRAS commands and then access the resulting data stream retrieved from CRAS. The actual transmission of requests and data is handled by a UNIX communications program accessible from 0PS4.

Finally, a set of rules assemble the appropriate message about the day's events recognized by the system and call on the UNIX mail facilities to deliver them to the appropriate users. ACE knows the target of each message based on the importance of the message to the user. This information is contained in WM and is user modifiable.

IV. Testing ACE

ACE has been field tested continually since the spring of 1982 by a long distance connection to remote CRAS systems. The local analysts report that they are very satisfied with ACE's performance. Although ACE has not discovered any problems unknown to human users, it did discover them quicker and missed neither the obvious nor the subtle. The praise of the analysts was outweighed only by their enthusiasm for permanent installation of the software.

The execution performance characteristics of the system are encouraging. Each nightly run for a large metropolitan area consisting of about 400,000 lines averaged only about one hour of VAX 11/780 CPU time (which can confidently be improved). On average, 5000 production invocations are sufficient for a complete run. New analysis tasks are under development to expand the scope of ACE's grasp of cable maintenance.

References

1. Shortliffe, E.H. 1976. Computer-based medical consultations: MYCIN. New York: American Elsevier

2. Stefik, M. 1980. Planning with constraints Rep. No. STAN-CS-80-784, Computer Science Dept., Stanford Universi ty.

3. McDermott, D.V. 1981. R1: The formative years. AI Magazine 2:21-29.

4. Forgy, C. and McDermott, J. 1977. OPS, a domain-independent production system language. IJCAI 5, 933-939.

5. Bell Telephone Laboratories. 1977. Engineering and operations in the Bell System.

6. Kellog C. 1982. Knowledge management: A practical amalgam of knowledge base and database technology, AAAI-82, 306-309.

7. Foderaro, U.K. 1979. The FRANZ LISP manual. Computer Science Dept., University of California, Berkeley.

8. Stefik, M., Aikins, J., Balzer, R., Benoit, J., Birnbaum, L., Hayes-Roth, F. and Sacerdoti, E. 1982. The organization of expert systems: A prescriptive tutorial. Xerox Corporation Report, Palo Alto Research Center.

9. Newell, A. 1969. Heuristic programming: ill-structured problems. In J. S. Aronofsky [Ed.], Progress in Operations Research. John Wiley and Sons. 361-414.