EXPERT SYSTEMS FOR AT&T SWITCHED NETWORK MAINTENANCE

Paul H. Callahan

AT&T TECHNICAL JOURNAL

Paul H. Callahan *is supervisor of the Expert Systems Application Group in the Network Operations Systems Planning Department at AT&T Bell Laboratories (Red Hill) in Middletown, New Jersey. He works on artificial-intelligence applications for network operations, and for systems engineering and software development technology. Mr. Callahan joined the company in 1982. He has a B.A. in physics from the State University of New York at Geneseo, and both an M.S. and Ph.D. in physics from the University of Illinois at Urbana-Champaign.*

As AT&T's telecommunications network becomes increasingly sophisticated and complex, advanced operations systems capabilities are needed to provision, manage, and maintain the network. To meet this need, artificial intelligence concepts are being incorporated into operations systems. This paper focuses on an expert system for switched circuit operations that builds on AT&T's Testing Operations Provisioning Administration System (TOPAS). The TOPAS Expert System (TOPAS-ES) is a real-time, distributed, multitasking expert system for switched circuit network maintenance that uses the diagnostic knowledge of experts in network maintenance to analyze network problems and locate faults.

Background and Perspective

Over the last 10 to 15 years, major advances in switching, signaling, and transmission—e.g., AT&T's 4ESS™ switch and common channel interoffice signaling (CCIS)—have dramatically increased the sophistication and capabilities of AT&T's network. (Panel 1 lists the acronyms used in this paper.) Coupled with this growth in capability and complexity, there was an increasing need for advanced *operations systems* capabilities to provision, manage, and maintain the network.

In response to the demand for increased operations capabilities, AT&T Bell Laboratories has developed a powerful array of computer-based operations systems that use state-of-the-art advances in software and hardware technology such as data-management, data-communications, transaction processing, microprocessors, and digital-signal processing. Today's embedded base of operations systems provides a sophisticated platform for network operations.

To build on that platform and further enhance operations capabilities, research and development in operations systems have been investigating artificial intelligence to provide increasingly powerful operations systems.

Panel 1. Acronyms in This Paper

CCIS	common channel interoffice signaling
CHA	chronic history analysis
CMS-1	Circuit Maintenance System No. 1
CPU	central-processing unit
CTMS	Carrier Transmission Maintenance System
ES	expert system
ESTA	expert system trouble analyzer
ESTS	expert system trouble sectionalizer
GTP	general telemetry processor
LTE	line terminating equipment
MUX	multiplexer
OS	operations system
RMS-Dx	Remote Monitoring System—Digital
TOPAS	Testing Operations Provisioning Administration System
TOPAS-ES	Testing Operations Provisioning Administration System—Expert System

Tomorrow's Network Operations. Tomorrow's network managers, provisioners, and maintainers will function in an environment that has a wide variety of service offerings, customer control, and dynamic provisioning and routing. Network technicians must respond to alarm patterns and other network data to anticipate network changes. Analytical solutions for dealing with this information load will be increasingly difficult to engineer in operations systems. Instead, the knowledge of network experts will be analyzed and captured for use in intelligent operations systems.

Customers seeking control of their critical communications needs expect robust, "user-friendly" interfaces to network operations systems. Network knowledge bases will be able to support explanation-based query of network information. This—combined with speech recognition, natural-language processing, and advanced display technology—will give customers a natural, understandable interface to the AT&T network.

Expert Systems Today. The ever-increasing importance of AT&T operations capabilities and advances in artificial-intelligence technology go hand-in-hand. This paper focuses specifically on the subfield of artificial intelligence known as *expert systems*. Expert-systems technology has importance for network operations in general and network maintenance in particular.

Network-maintenance methods possess a heuristic type of knowledge that lends itself to an expert-systems architecture and paradigm for implementation. Equally important, expert-systems technology is "ready." Rule-based programming technology is mature enough that rule-based shells can be interfaced with data-management systems and other application-programming languages. (A shell is a special program that guides actions.) Some shells run in the UNIX® operating system environment and thereby obtain all the benefits of UNIX system process communications and the UNIX system software development and test environment.

Virtually all expert-system applications to date are stand-alone or batch-processing advisor systems. The Testing Operations Provisioning Administration System—Expert System (TOPAS-ES) for AT&T network maintenance, described in this paper, is among the first of a generation of real-time, distributed, multitasking expert systems for effecting specific actions in the network.

Switched Network Maintenance Overview

This section gives an overview of maintenance for AT&T's switched network.

Network Operations Functional Hierarchy. From an operations perspective, the AT&T switched network can be viewed on three functional levels (see Figure 1):

1. *Network elements*—provide interconnection, translation, and transportation services.
2. *Operations systems*—provide test access, measurement, transaction processing, and data management services for maintenance, provisioning, administration, and management. (The provisioning function defines new circuits or alters existing ones.) Operations systems also provide the human interface to the network elements for operations.
3. *Work centers*—analyze and test reported circuit problems (trouble analysis), analyze alarm and traffic patterns (correlation), identify the faulty section of a circuit (sectionalization), and prepare a repair order (referral). They also provision, administer, and manage the network.

Maintenance of the AT&T switched network is the sum of

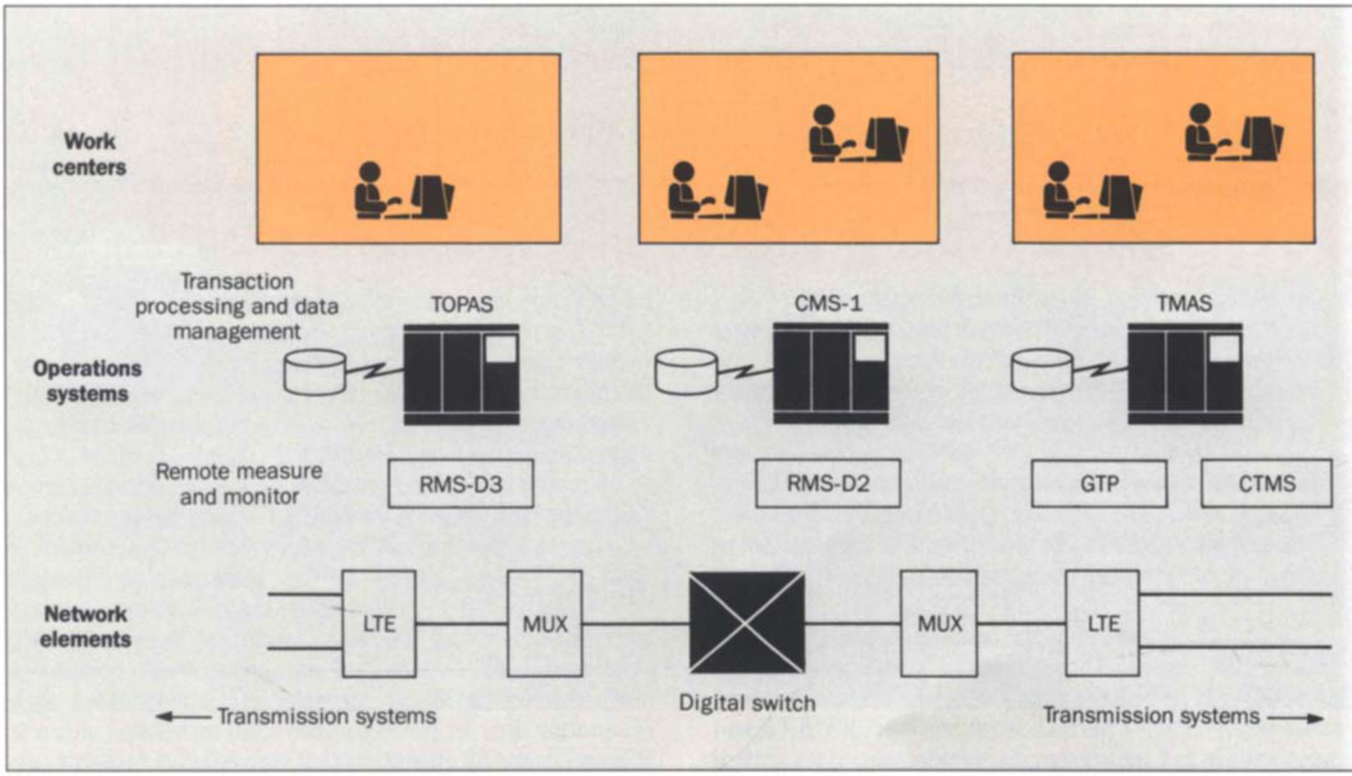

Figure 1. Network operations functional hierarchy for maintenance, provisioning, administration, and management. CMS-1 is the Circuit Maintenance System No. 1, RMS-D2 and RMS-D3 are digital remote measurement systems, TMAS is the Transport Management and Administration System, and TOPAS is the Testing Operations Provisioning Administration System. CTMS = Carrier Transmission Maintenance System; GTP = general telemetry processor; LTE = line terminating equipment; MUX = multiplexer.

the maintenance capabilities distributed across these functional levels.

Current Maintenance Operations. There are several work centers for maintenance of the switched network. Figure 2 shows an example of switched-circuit maintenance in the 4ESS and 5ESS® digital switch environment. The TOPAS[1] and CMS-1[2] operations systems support switched-circuit provisioning and maintenance for 5ESS and 4ESS switches, respectively. The remote measurement systems, RMS-D3 and RMS-D2, provide TOPAS and CMS-1 with the capability for computer-controlled remote testing and monitoring of circuits that terminate on a switch.

During call processing, out-of-service indicators are reported to TOPAS and CMS-1 by the 5ESS and 4ESS switches, respectively. In addition, TOPAS and CMS-1 automatically test the entire switched-circuit network routinely, compare test results with network standards, and remove substandard circuits from service. For out-of-service indicators and substandard circuits, trouble reports are generated and reported to the work centers by TOPAS and CMS-1.

Many of the troubles reported to work centers are transient; that is, if the work center investigates the trouble, no trouble is found. To sort out transient troubles from "hard" troubles, work-center technicians perform trouble analysis. The transient troubles are cleared, and technicians then sectionalize the hard troubles and refer them to work groups for repair.

Other latent hard troubles are found among patterns of the transient troubles themselves. These troubles, while apparently transient, are nevertheless a result of faulty equipment in the network. Technicians use trouble-history logs, additional testing, and their own experience with certain circuit groups to identify such faulty circuits.

To sectionalize hard troubles, technicians execute

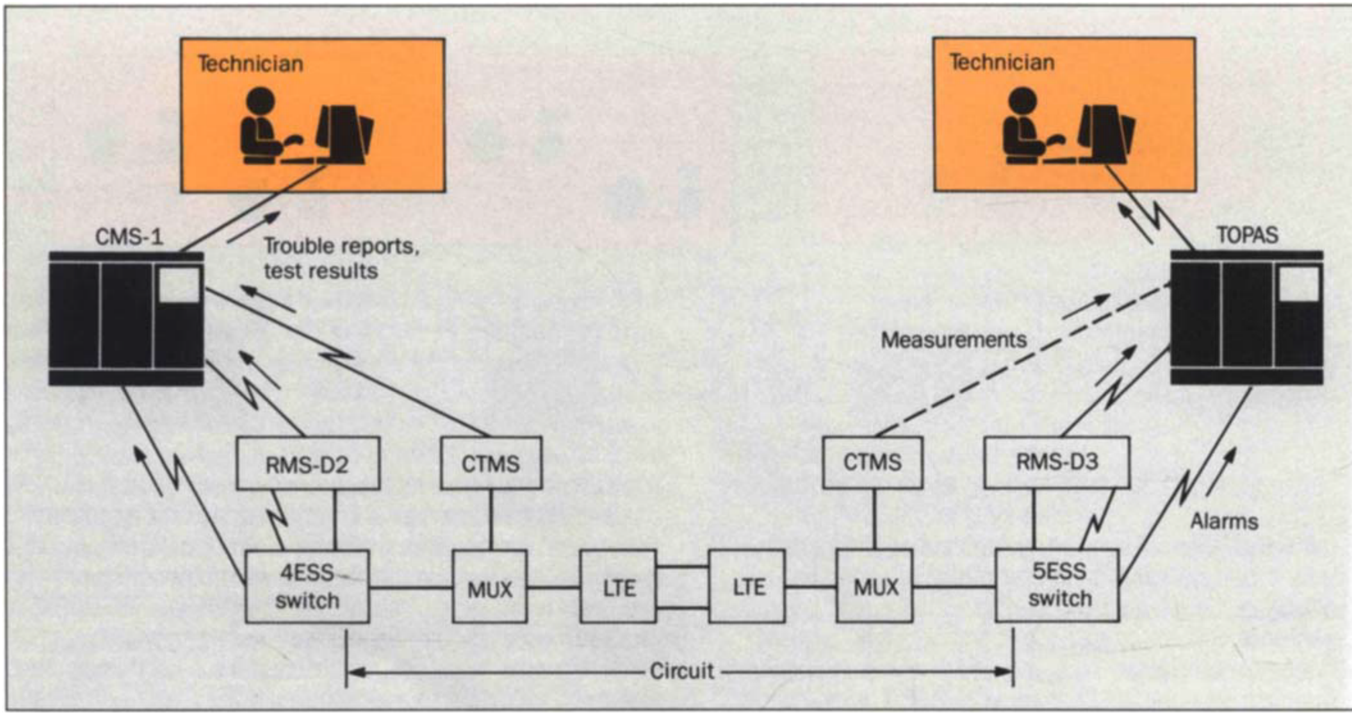

Figure 2. Switched circuit maintenance. CMS-1 is the Circuit Maintenance System No. 1, RMS-D2 and RMS-D3 are digital remote measurement systems, and TOPAS is the Testing Operations Provisioning Administration System. CTMS = Carrier Transmission Maintenance System; LTE = line terminating equipment; MUX = multiplexer.

various testing and measurement strategies on the faulty circuit. The strategies include:

- Test from both ends of the circuit.
- Test in both the transmit and receive directions.
- Take measurements at various test points available along the circuit.

The strategies vary according to the technology, type of trouble, and test results. After the fault is located, the trouble is referred to the appropriate work group for repair.

In any large switched network, the volume of calls, number of circuits in service, and array of trouble detecting and reporting capabilities generate many trouble reports per month. (A switching system completes millions of calls per day and transmission systems continuously monitor more than a million scan points.) The trouble reports must be resolved by technicians in the work centers.

Applicability of Expert-Systems Technology

Here, we discuss the applicability of expert-systems technology for network maintenance. Expert-systems technology is ready, and appropriate, to tackle the complexity of work-center maintenance tasks.

Domain Characteristics of Network Maintenance. The key attributes of the knowledge and tasks in the network-maintenance domain that make it ideal for an expert-system application are:

- Those who do the tasks use diagnostic problem solving.
- The tasks have the right attributes; they lend themselves to expert-system techniques.
- Expertise exists.

Diagnostic problem solving. Network maintenance is an excellent example of diagnostic problem solving; that is, inferring possible system malfunctions from observables or symptoms. Diagnostic-problem-solving knowledge is often associational and heuristic, as opposed to, for example, reasoning from first principles.

As an example of the difference, consider a circuit with an echo problem. A network maintenance expert's view of the problem is: "If a measurement confirms the presence of echo, then the channel unit needs to be balanced." An electrical engineer's view of the same problem is: "If a signal encounters an impedance mismatch, then signal energy will be reflected." The first is diagnostic problem solving, the second, reasoning from first principles.

The network expert doesn't need to understand electrical impedance to diagnose the problem. And, while the electrical engineer knows why an echo problem could occur, he or she doesn't know the tests, test points, strategies, and operational procedures to infer the source of trouble on the circuit. The inference chains formed in diagnostic problem solving lend themselves to an expert-system implementation.

Task attributes. Typical tasks in a network-operations work center are: provisioning, trouble analysis, correlation, sectionalization, and referral. To accomplish these tasks, technicians often use known procedures, but also rely heavily on their own experiences with various network technologies and trouble types. Although many of the procedures are standard, they are nonetheless complex. Whether using standard procedures or experience-based alternatives, technicians must exercise decision making and judgment based on qualitative and quantitative information.

Plotted in Figure 3 are problem attributes that suggest a problem is well suited to an expert-systems solution.[3,4] The checked boxes represent the problem attributes, according to knowledge engineering studies, of the typical network-operations work-center tasks. The qualitative result is that at least 60 percent of the work-center maintenance functionality can be represented and implemented in an expert system.

Existence of expertise. The existence and profile of expertise in the domain of interest are important criteria in selecting a realistic expert-system application.

In the network maintenance domain, technicians exhibit a high level of competence at their jobs. They have learned the procedures and can resolve common troubles that occur. Some of the technicians have become truly expert in network trouble-shooting; they know the network in detail and, therefore, know shortcuts and effective alternatives. These are the technicians that others look to when they encounter a problem they cannot resolve.

The existence of and close collaboration with such experts is essential to successful expert-system construction. Expert systems also provide a way to preserve their valuable trouble-shooting expertise.

Representing and Controlling the Knowledge. The expert knowledge in the domain must lend itself to a formalized representation. In an expert-system architecture that separates inference from knowledge, the formalized representation is implemented in a "knowledge base." The control of inference is handled by the "inference engine."

Knowledge base. Common knowledge-representation techniques include semantic networks, frames, and rules. The diagnostic expertise in the network maintenance domain is well represented as rules (IF *condition* THEN *action*) and facts. During a diagnostic session, a technician has knowledge of a particular set of symptoms, a set of test results, and a particular network state. The technician, using experience-based and procedural knowledge, then decides what action to take next. In terms of rules, this implies that one can specify the conditions of a rule's invocation and the associated action (commonly called the *recognize-act* cycle).

For example, consider the following rule for a technician to start a transmission test on a faulty circuit:

IF: a transmission test has been selected
and a test access circuit has not been selected
THEN: select a test access circuit

This rule illustrates two key aspects of expert-system rules for the domain.

- The conditions of its invocation are uniquely clear.
- It encapsulates an *item of expert knowledge.* That is, the domain expert can understand directly the conditions and actions that the rule represents.

This rule can be "debugged" in isolation, separate from the rest of the program.

Also, a large number of facts associated with the domain affect both the conditions of rule invocation and the choice of action selected. If these facts are represented

Expert system attributes →	Provisioning	Trouble analysis	Correlation	Sectionalization	Referral
Combinatorial complexity			√		
Nonalgorithmic		√		√	
Extensive symbolic reasoning	√	√	√	√	√
Multiplicity of facts	√	√	√	√	

Network expert tasks →

Figure 3. Problem attributes of work-center tasks. A check means that expert systems offer a solution to the problem.

separately and explicitly in the knowledge base, they can be engineered separately from the rules. As examples of facts from the circuit testing domain for the 4ESS switch, consider:

FACT: The default level of the 1004-Hz test is –16.0 db.
FACT: The frequency of a continuity check test is 2010 Hz.

To construct the knowledge base of rules and facts, a knowledge engineer peruses existing documentation and carefully interviews experts to find out how they trouble-shoot network problems.

Inference engine. For most rule-based systems, the inference strategy selected is either "forward chaining" (data driven) or "backward chaining" (goal driven). In a forward-chaining system, given a current set of facts that are true, the system deduces what other conditions now become satisfied. In a backward-chaining system, a question is asked and the system uses rules to find facts from which it can deduce an answer to the question.[5]

When an exhaustive search of possible answers to a question is neither feasible nor desirable, but data relevant to answering the question can be obtained, then forward chaining is the more natural choice. Here, the data are a more reliable guide to a decision than is the final goal. The control of inference over the rules and facts in the network-maintenance domain is best handled by forward chaining. For example, the data:

"There is a noise trouble report on this circuit."

is a more specific guide to a technician's actions than the question a technician seeks to answer:

"What equipment on this circuit is faulty?"

Technicians are data driven in their approach to trouble-shooting.

Conflict resolution. One common complication in expert systems is what to do when, under a given set of conditions and facts, the IF parts of more than one rule are satisfied. This "conflict set," if not resolved, may imply taking simultaneous, but different, actions as the next step. For example, it would be likely in the circuit-testing domain for conditions to arise where the actions, "measure noise" and "measure loss," apply at the same time.

To resolve such a conflict, the domain expert must tell the knowledge engineer what action, given a certain trouble type, is most likely to converge first. Then—by using a resolution strategy known as *data ordering*—the best action will be inferred first. Where data ordering leaves unresolved conflicts, the *specificity* strategy will exercise the rule with the most preconditions. If there are still rules in the conflict set, *recency* will execute the rule with the most recently updated working-memory elements.[6] (This is equivalent to assuming that the most recent data is the most reliable guide for taking action.)

This conflict-resolution strategy is important in this domain because maintenance actions can be time and resource intensive; that is, they may require accessing circuits, using operations system (OS) and network element resources. Using this conflict-resolution strategy, an expert system will attempt actions that are likely to lead most efficiently to the source of trouble.

The advent of expert-systems technology has greatly increased the feasibility of formalizing diagnostic knowledge. The knowledge-engineering methodology, which requires collaboration between knowledge engineers and expert technicians, can capture, implement, and refine diagnostic knowledge in an executable form. The type of knowledge and the existence of network experts make network maintenance an ideal application domain for expert systems.

The TOPAS Expert System

TOPAS-ES is an expert-system application for trouble analysis, sectionalization, and referral for network troubles in the 4ESS and 5ESS switched-circuit environment. TOPAS-ES is an off-board system (see Appendix A) that interfaces to the TOPAS and CMS-1 operations systems. Figure 4 shows the TOPAS-ES configuration in the 4ESS and 5ESS switch domain.

The TOPAS-ES architecture is based on a set of concurrent, communicating processes and its highest level has three main subsystems:

- Knowledge base and inference engine subsystem
- Communications and system interface subsystem
- User interface subsystem.

Each major subsystem consists of one or more UNIX system processes. The knowledge in TOPAS-ES is represented as rules and facts, and the inference engine uses forward chaining. The main TOPAS-ES processes are "alive" at all times.

TOPAS-ES—a real-time, multi-user, expert system—must keep up with the volume of trouble reports, perform its analysis, and continuously assert the proper trouble resolution and circuit disposition. TOPAS-ES runs several expert-system processes simultaneously—each with its own separate working-memory state—to cover all circuits that terminate on the switch. (The real-time requirement is intrinsic to this maintenance task. The multi-user requirement is an artifact of the off-board architecture and the way the host OS assigns circuit maintenance responsibility to technicians.)

TOPAS-ES is a distributed expert system in that one TOPAS-ES, located at the control office for a faulty circuit (the office with maintenance responsibility), will communicate with another TOPAS-ES, located at the other end of the circuit, and cooperate in sectionalization. TOPAS-ES is *bimodal*; it has a director mode and a responder mode. While TOPAS-ES is sectionalizing a control-office trouble (director mode), it can simultaneously accept an interrupt from another TOPAS-ES and help it sectionalize another circuit (responder mode).

TOPAS-ES consists of two expert-system subsystems that work cooperatively together. The first is the expert system trouble analyzer (ESTA), and the second is the expert system trouble sectionalizer (ESTS).

Expert System Trouble Analyzer. ESTA performs three main functions: trouble ticket analysis, chronic history analysis (CHA), and referral.

Trouble ticket analysis. As described earlier, many reported troubles are transient, and only a few trouble reports reflect hard faults in a piece of network equipment that will need to be sectionalized and referred. So, it is inefficient for technicians to spend time verifying that most reported troubles are transient.

The ESTA system uses heuristics based on the experience of expert technicians to do an analysis that differentiates between hard and transient troubles. For example:

IF: the outage report is "local carrier alarm"
and the trouble has not cleared during initial monitoring
and the trouble has occurred on a chronic circuit
THEN: keep the circuit out of service
and perform extended monitoring

The ESTA knowledge consists of testing and soak strategies (holding a trouble for a predetermined interval to see if the switch will clear it) based on, and varying

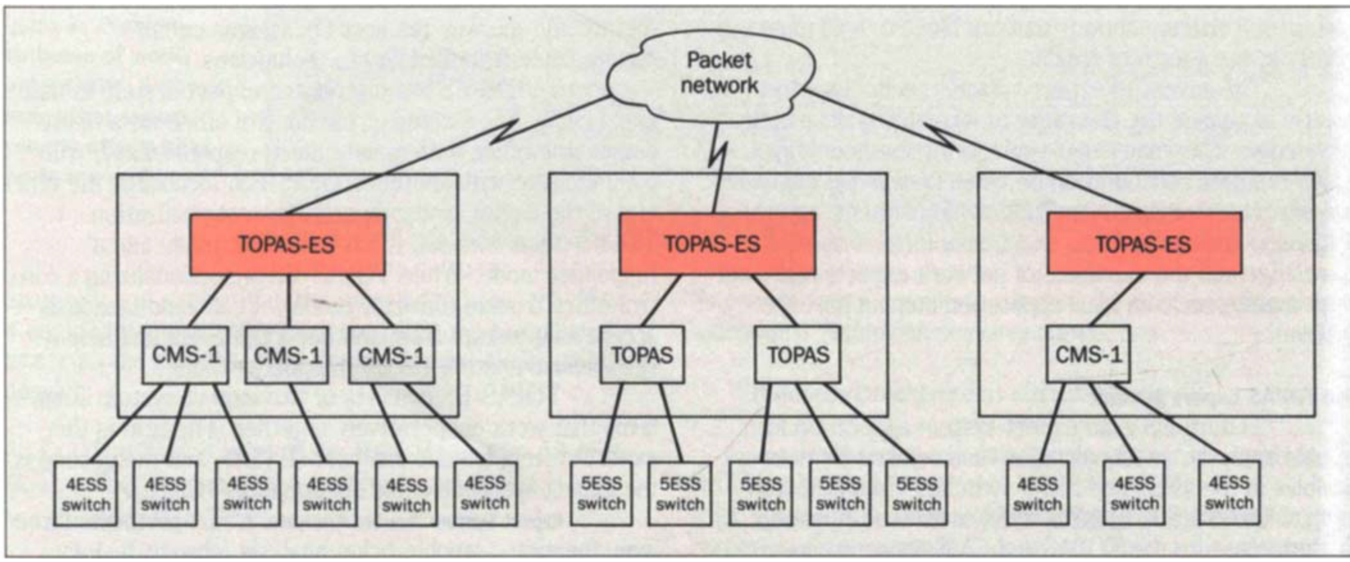

according to, the outage-code type. This feature of ESTA is able to clear most transient troubles.

Chronic history analysis. The CHA, which is ESTA's second primary function, identifies faulty circuits in transient trouble patterns. To do CHA, ESTA applies expert rules to a statistical distribution of outages that it maintains. For example:

IF: the signal type is CCIS
and more than 5 reports occurred in less than 30 days
THEN: the circuit is a chronic circuit

Because transient circuits were often returned to service, CHA-type troubles previously remained in the network undetected and continued to affect adversely customer perception of quality.

A switch is most likely to detect a trouble and remove a circuit from service during peak traffic hours. But without ESTA, technicians sort through these troubles during the day. Therefore, the least capacity is likely to be available when the need for network capacity is the greatest. ESTA eliminates "think time" in restoring good circuits to service, thereby reducing blockage and alternate routing.

Figure 4. TOPAS-ES for switched-service maintenance. CMS-1 is the Circuit Maintenance System No. 1, TOPAS is the Testing Operations Provisioning Administration System, and TOPAS-ES is the expert system co-processor for TOPAS.

Referral. ESTA's third main function is to refer hard troubles for sectionalization. It will post all relevant results that are collected on a trouble ticket, along with an initial disposition that tells if the trouble is associated with a chronic circuit.

Expert System Trouble Sectionalizer. The ESTS system performs two main functions: sectionalization and referral. It sectionalizes hard troubles to isolate the faulty network element on the circuit in question. ESTS assumes that troubles it receives are confirmed hard troubles. Time and resources would be wasted if sectionalization were used to confirm "no trouble found" cases.

Sectionalization. As Figure 2 shows, the typical sectionalization session requires two technicians, one at each end of the circuit. Sectionalization—depending on the available test points, data, and outage type—may range from a simple and direct task to a very complex one.

The ESTS system uses rules that are based on the sectionalization strategies of expert technicians. For example:

IF: the far-end master group's transmit deviation is good
and an adjacent circuit on the same facility is good
THEN: the far-end channel unit is bad

The ESTS knowledge base contains sectionalization strategies that vary according to the underlying type of fault: either digital or analog, signaling or transmission. Another body of heuristics "connects" these segments to the most probable outage codes associated with the fault. Because sectionalization can be time and resource intensive, ESTS attempts the most likely scenarios first, but will try all its scenarios before giving up.

Referral. ESTS's second main function is to refer the located fault to the appropriate work group for repair. Once the source of the fault is located, it is straightforward—from an algorithmic viewpoint—to analyze the circuit layout data, determine what equipment should be replaced, and who has the repair responsibility.

TOPAS-ES can continuously execute trouble analysis and sectionalization in unattended work centers (i.e., nights and weekends). For new offices, TOPAS-ES "imports" expert technician capability where there would not ordinarily be an expert, thus improving the office's start-up quality.

Taken together, the ESTA and ESTS subsystems in TOPAS-ES provide powerful capabilities for automated trouble analysis, sectionalization, and referral covering the entire AT&T switched network.

Summary

Global information movement and management requires ever increasing network-operations capabilities. To meet this challenge, AT&T is incorporating artificial intelligence into advanced network-operations systems. In particular, expert-systems technology is being integrated into operations systems for network maintenance. The functionality of many network-maintenance tasks, the expert-systems paradigm, and the maturity of expert-systems technology are well matched.

The TOPAS Expert System is an expert-system application for trouble analysis, sectionalization and referral of network troubles in the 4ESS and 5ESS switched-circuit environment. TOPAS-ES is an off-board system that interfaces to the TOPAS and CMS-1 operations systems. TOPAS-ES is among the first real-time, distributed, multi-user expert systems and it does network maintenance without human intervention or consultation.

Acknowledgments

I gratefully acknowledge the valuable comments that Raj Dube, Allen Ginsberg, Robert Mallon, Alan McMichael, and Jason Tsay made on this paper. I thank Laverne George for his support of this project.

References

1. R. Merski and D. Parrish, "Operations Systems Technology for New AT&T Network and Service Capabilities," *AT&T Technical Journal*, Vol. 66, No. 3, May/June 1987, pp. 64-72.
2. J. Giunta et al., "Data/Trunk Administration and Maintenance," *The Bell System Technical Journal*, Vol. 56, No. 7, September 1977, pp. 1203-1238.
3. D. A. Waterman, *A Guide to Expert Systems*, Addison-Wesley, Reading, Massachusetts, 1986, pp. 24-32.
4. P. Harmon and D. King, *Expert Systems*, John Wiley & Sons, New York, 1985, pg. 8.
5. N. Nilsson, *Principles of Artificial Intelligence*, Tioga Publishing Co., Tioga, New York, 1980, pp. 32-34.
6. B. G. Buchanan and E. H. Shortliffe, *Rule-Based Expert Systems*, Addison-Wesley, Reading, Massachusetts, May 1985, pp. 48-50.

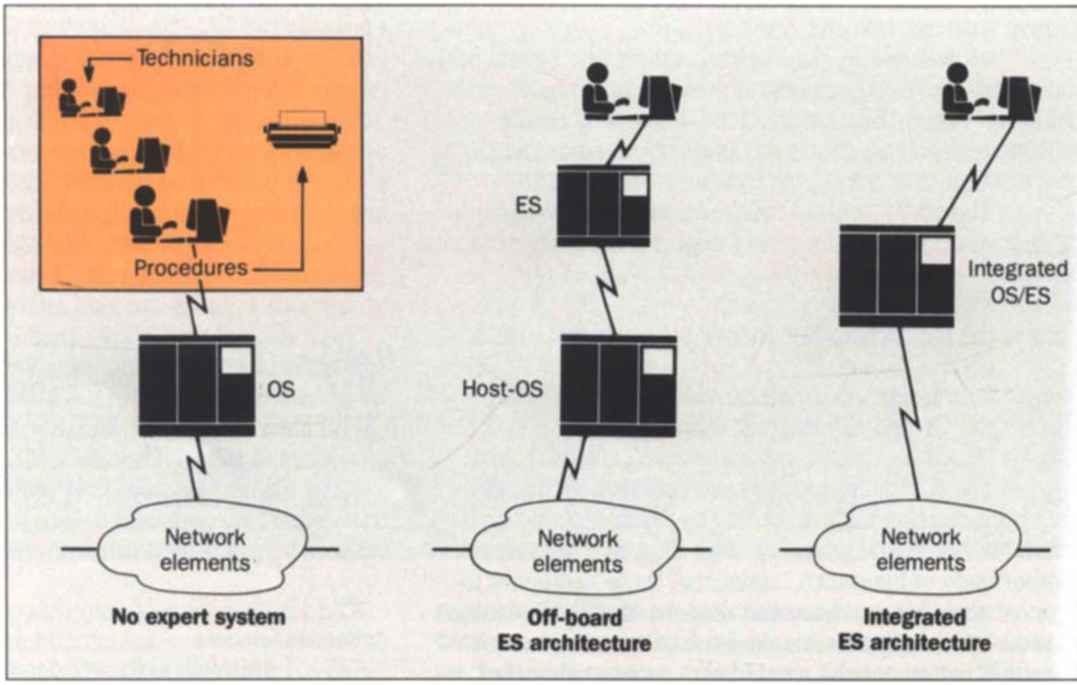

Figure A-1. Expert system (ES) and operations system (OS) integration.

Appendix A. ES/OS Integration

An expert system for network maintenance must integrate with the embedded base of operations systems. Two alternatives for integrating expert systems with the embedded base are (see Figure A-1):

- Deploy an "off-board" system (runs on a separate processor) that *interfaces* to existing OSs.
- Build an "on-board" system (runs on the same processor) that executes as another host-OS *application.*

Each alternative has certain advantages and disadvantages.

Off-Board Architecture. In this architecture, the expert system interfaces externally to the embedded base of operations systems and to expert technicians. The advantages of this architecture are:

1. With an off-board architecture, expert-system builders can move to the field for iterative implementation, testing, and refinement of the knowledge base.
2. The off-board system can implement other OS system interfaces that may not be available via the host OS or would be too costly to develop in the host.
3. The off-board architecture allows the expert system's project management to be decoupled from the host operations system's project management.

Some disadvantages of an off-board system are:

1. A proliferation of ES/OS interfaces will be difficult to maintain.
2. Additional, redundant host-OS/ES interface development is required whenever the host-OS will support a new or enhanced interface to another system (if the expert system is to use the new system's capabilities).
3. The extra computer increases hardware costs and systems administration.

On-Board Architecture. In this architecture, the expert system is integrated with the host OS as an application process. Some advantages to the integrated architecture are:

1. As the OS architecture of AT&T's network evolves, distributed-network data stores will be deployed. Host OSs—and, therefore, integrated ES applications—will have access to such data stores.
2. Any database and networking capabilities that the host OS supports will greatly increase the potential power and scope of the integrated expert system.
3. Integrated ES applications allow increased interaction complexity between knowledge subsystems and other applications. The knowledge subsystem can even function as a server to other applications (which is not generally true in an off-board architecture).

Some disadvantages of the on-board system are:

1. Today's ES development environment often includes software, and sometimes hardware, that is not supported in a "standard" OS development environment.
2. Depending on the design and scope of the integrated ES, it may negatively affect host-OS performance by being memory and CPU-cycle intensive. (A CPU is a computer's central-processing unit.)

Summary. Most existing expert systems in many application areas use an off-board architecture, but the first integrated applications are on the horizon. In the long run, whether an expert system is on-board or off-board will depend more on the architectural and product-evolution considerations as a whole than on any particular difficulties with expert-systems technology.

(Manuscript received October 28, 1987)