Machine Intelligence for ATC Equipment Maintenance

Harold M. Heggestad

The normal duties of ATC equipment maintenance are to observe performance indicators and measurements, infer from these data whether the equipment is in good health, and, if a problem is found, diagnose its cause and perform the necessary repairs. Machine intelligence (MI) technology offers a means for automating much of this work load. Indeed, Lincoln Laboratory has developed and fielded several expert systems for similar applications involving the maintenance and control of military communications systems. This article describes three such expert systems, and then discusses opportunities for MI to automate the monitoring and maintenance of ATC equipment. A high-level design of an MI approach to the remote maintenance monitoring of the ASR-9 airport surveillance radar is presented, along with conceptual descriptions of other MI applications for the Federal Aviation Administration (FAA).

In the maintenance of ATC equipment, normal duties include observing performance indicators and measurements, inferring from these data whether the equipment is in good health, and, if a problem has been detected, diagnosing the problem's cause and performing the necessary repairs. Traditionally, these tasks have been performed by personnel at numerous well-staffed work centers located near the equipment sites. This established practice, however, has recently undergone change because of tightening budgets and increasing equipment complexities. Indeed, technicians today must handle increasing numbers and types of equipment, and they must usually cover larger territories. The situation has been aggravated by the fact that expert technicians are retiring and their successors have been fewer in number and less well trained. To alleviate these growing pressures, machine intelligence (MI) technology could be used to automate much of a technician's work load.

The development of MI has thus far been a mix of solid accomplishments and unrealized expectations. One area of MI that has achieved considerable practical success is the expert system, which undertakes to capture and reproduce the knowledge of human experts in solving complex but relatively well-defined problems. Lincoln Laboratory has recently implemented three expert systems to automate the control of military communications networks [1–7]. The three systems are described in the following sections of this article. The succeeding section builds on this work to describe an expert-system approach for the Federal Aviation Administration (FAA)—the modernization of the Remote Monitoring System (RMS) of the ASR-9 airport surveillance radar. The next section discusses expert-system technology as a critical factor in a consolidated FAA Operations Control Center (OCC) concept.

In the remainder of this section, we give a brief introduction to expert systems by discussing the problem domains where expert systems can be valuable, the capture and encoding of problem-solving knowledge, and the practical alternatives available for building expert-system software systems. We will see that these factors are well matched to the military communications applications already in hand, as well as to a variety of FAA problems—in particular, to modernizing the ASR-9 RMS.

One generic problem domain that is particularly well suited for expert-system application is a skill-intensive repetitive multistep process, such as observing and interpreting customer complaints or complex alarm patterns, doing methodical cause-and-effect fault diagnosis, or determining the appropriate action to correct diagnosed faults. Examples of this type of application are the Machine Intelligent Technical Control (MITEC) and the Transmission Monitoring and Control (TRAMCON) Adjunct Processor (TAP), two of the military communications systems defined and described in the following sections.

Another class of problems suitable for expert-system application involves the vigilant monitoring of large volumes of mostly unremarkable system status data to look for complex but recognizable patterns and events that indicate possible trouble in the target system. An example of this type of application is the Network Management Expert System (NMES), another Lincoln Laboratory–developed military communications system that is defined and described in a succeeding section.

A third class of problems involves the preventive maintenance of systems by the continual monitoring and evaluating of performance parameters and built-in test results, both to verify satisfactory performance and to predict failures, thereby permitting fix-before-break maintenance actions. All three military systems described in this article include preventive maintenance functionality.

Central to the implementation of an expert system is the process of *knowledge engineering,* which is the acquisition and implementation of the problem-solving knowledge base for the target domain. Such expert knowledge can take various forms. One familiar category of knowledge is the lifetime treasure of learned expertise and empirical rules possessed by a highly skilled individual who is about to retire. This knowledge category was the object of a number of pioneering expert-system development efforts. Another form of expert knowledge is an extensive compendium of textbook-like rules or explanations that an expert needs to know and be able to retrieve to solve every problem or combination of problems that can occur in the target domain. A third form of expert knowledge is the body of behavioral relationships of the components of a target domain. These relationships may be causal or stochastic or a mixture of the two, and they can be expressed in terms of a model or a simulation. Such a model could be used for further elaboration of knowledge about the domain, and the model itself could be made an integral part of a problem-solving expert system.

The real work of creating an expert system is the knowledge engineering. In the early phases of a project, a commercial off-the-shelf (COTS) expert-system shell can be very useful because it allows developers to focus most of their energies on the knowledge engineering and, with minimal programming effort, to convert the growing knowledge base into a series of executable prototypes. After the knowledge base is well in hand, the emphasis may shift to achieving high performance in field testing and to the potential deployment of the expert system. This shift in emphasis tends to lead the developer to use custom software for reimplementing the latest prototype, thereby avoiding the performance-degrading generality and overhead of the COTS shell. In fact, this cycle was followed in developing all three of the military expert systems reported in this article.

The Communications Control Environment of the Department of Defense

Figure 1 illustrates the control hierarchy for the Defense Information Systems Network (DISN)—the worldwide communications system for the Department of Defense (DoD). The superscripts in the figure indicate the target locations of Lincoln Laboratory–developed expert systems, as described in the following sections of this article.

The Continental United States (CONUS) Operations Center, located in the headquarters of the Defense Information Systems Agency (DISA) in Washington, D.C., is concerned primarily with long-range planning and administration of the DISN. In CONUS, active control of DISN operations is conducted by the commercial contractors that provide leased DISN services.

Overseas, there are two manned military theater operations centers where day-to-day control of the DISN is actively exercised. The European facility is located at Patch Barracks in Stuttgart, Germany, and the Pacific theater center is located at Wheeler Army Air Field near Honolulu, Hawaii.

The bottom of Figure 1 shows the communications equipment and facilities distributed throughout each theater. These facilities, where the nuts-and-bolts maintenance and troubleshooting are performed, are manned overseas by military personnel, in many cases with sizable around-the-clock staffs. Typically, however, the facilities are undermanned to a considerable degree, skilled experts are rare, and many of the personnel are trainees. Thus the facilities are excellent candidates for the implementation of expert systems.

Machine Intelligent Technical Control (MITEC)

The Tech Control Facilities (TCF), illustrated in Figure 1, share the primary mission of maintaining reliable high-quality operation of about 61,000 dedicated military communications circuits worldwide. These circuits include hot lines to critical overseas commanders, data links between computer centers, and trunk connections among the switches in military voice and data networks. The DoD operates about 400 TCFs, each of which is a major junction point where as many as 1000 circuits converge, typically in multiplexed groups on carriers such as fiber optics, telephone circuitry over land, satellite channels, and radio links. In a TCF, the circuits pass through various communications equipment, including patch panels that allow rerouting and test access for troubleshooting.

Each dedicated military circuit is assigned a restoral priority (RP) appropriate to its mission criticality. A circuit typically passes through several TCFs along its route from source to destination. When a circuit degrades or fails, Tech Control personnel at the TCFs collaborate in locating the faulty element and restoring service by using any available spares. If there are no spares, a failed circuit can preempt assets from a circuit of lower RP. The time allowed for this repair may be minutes to hours, depending on the criticality of the circuit.

The dedicated military circuits have been installed incrementally over decades; hence the circuits vary greatly in equipment type, age, and complexity. Because of this circuit variability, Tech Control personnel need extensive and sophisticated knowledge to perform their jobs, especially under the given time

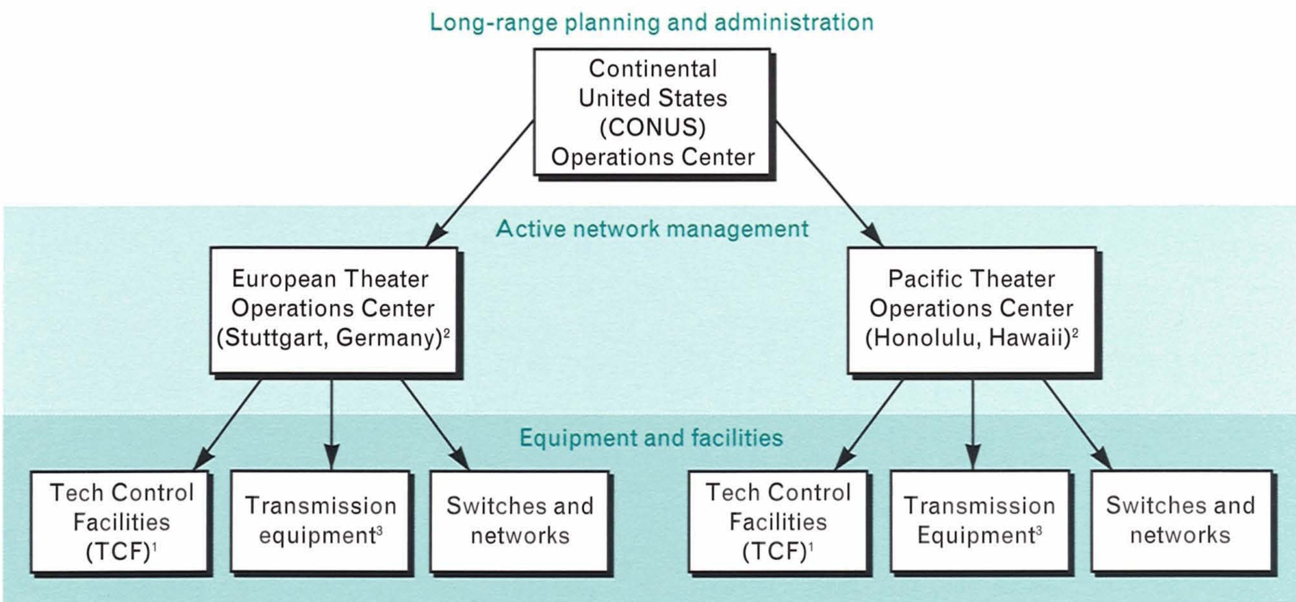

FIGURE 1. Control infrastructure for the Defense Information Systems Network (DISN)—the worldwide communications system for the Department of Defense. The different superscripts indicate the target locations of expert systems developed at Lincoln Laboratory. Superscript 1 indicates locations of the Machine Intelligent Technical Control (MITEC) expert system, superscript 2 indicates locations of the Network Management Expert System (NMES), and superscript 3 indicates locations of the expert system for the Transmission Monitoring and Control (TRAMCON) Adjunct Processor (TAP).

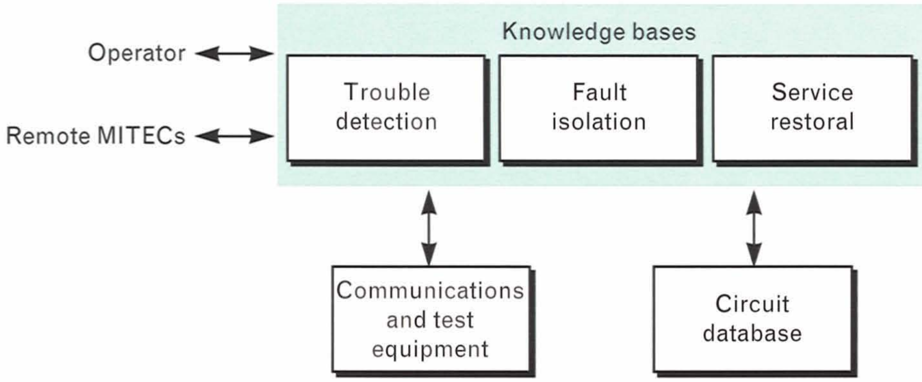

FIGURE 2. Top-level block diagram showing the architecture of the MITEC expert system. The knowledge bases for trouble detection, fault isolation, and service restoral are derived from a combination of textbook procedures and human expertise.

constraints. The situation is compounded by the familiar budget and training problems that have made it very difficult to grow and retain adequate staffs. Because of these factors, the U.S. Air Force Rome Laboratory sponsored a project at Lincoln Laboratory several years ago to develop an expert system for the TCFs, and the Machine Intelligent Technical Control (MITEC) system was created as a result.

The operation of MITEC can be explained conveniently in terms of the top-level block diagram in Figure 2. There are three knowledge bases, each derived from a combination of textbook procedures and human expertise. The trouble-detection function is triggered by alarm signals from communications equipment or by user complaints entered by an operator.

Upon identification of the affected circuit and recognition of the problem type, MITEC applies logical procedures from the fault-isolation knowledge base to locate the failed element in the circuit. These procedures include the electronic switching of test equipment to test points as appropriate, and can involve collaboration with MITECs at other TCFs to apply and measure test signals. Figure 3 shows the MITEC screen during a typical step in the fault-isolation process in which MITEC is about to measure the signal at the digital port of a modem in the path of the faulty circuit. The MITEC screen displays a diagram of the circuit with a dashed box highlighting the present focus of MITEC's attention, namely, a test point that is to be electronically switched to a digital oscilloscope. Figure 4 shows the clock and data waveforms that were read into MITEC from the oscilloscope, and the logic that analyzed the waveform to determine that the signal at that point was good. Using this result, MITEC now selects the next logical step in the process.

After the faulty element has been found, the service-restoral knowledge base shown in Figure 2, in consultation with the circuit database, searches for those spare or preemptible parts which can be electronically patched around the failed element. Thus the circuit is restored to operation.

Throughout the MITEC project, heavy emphasis was placed on involving the prospective users, namely, the military Tech Control personnel. Figure 5 shows an Air Force sergeant in the TCF at Andrews Air Force Base, Maryland. The sergeant is demonstrating the MITEC software on a computer that Lincoln Laboratory placed there for that purpose at the beginning of the project. The users repeatedly provided feedback about the current software to Lincoln Laboratory, and the next version of software sent back to the users incorporated that feedback. This interaction heightened user enthusiasm and facilitated the knowledge-engineering process.

At the completion of the MITEC project at Lincoln Laboratory, the technology was transferred to the newly created MITEC Program Management Office (PMO) in the Air Force Systems Center, located at Tinker Air Force Base, Oklahoma. The MITEC PMO is field-testing an early-release version of MITEC at Andrews Air Force Base, and is making fi-

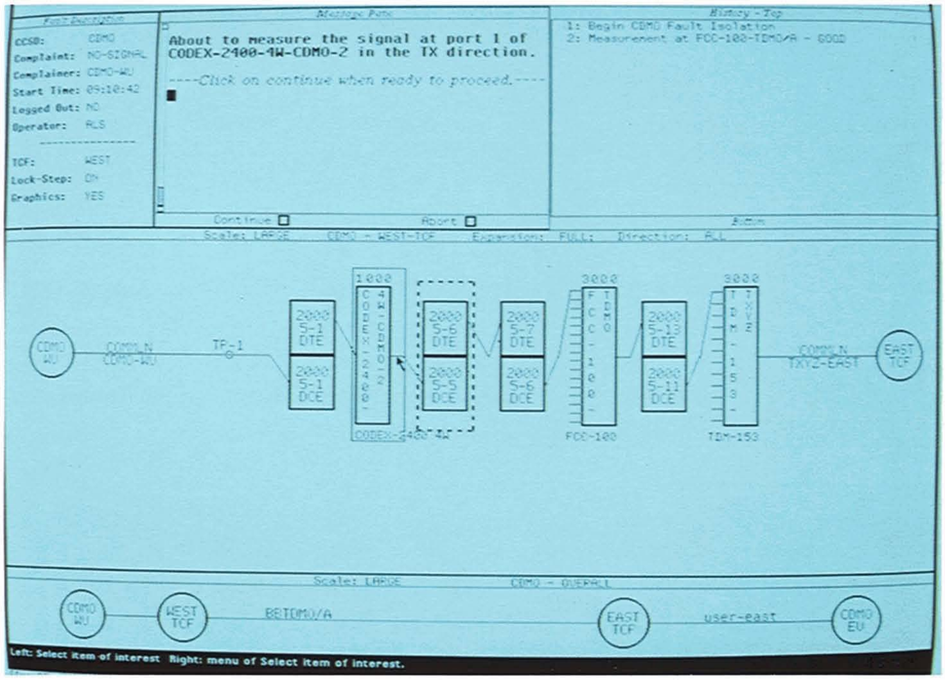

FIGURE 3. Example of MITEC display during troubleshooting. In the photo, MITEC is about to measure the signal at the digital port of a modem that is in the path of a faulty circuit. The computer screen shows a diagram of the circuit with a dashed box highlighting the present focus of MITEC's attention. This test point will be electronically switched to a digital oscilloscope.

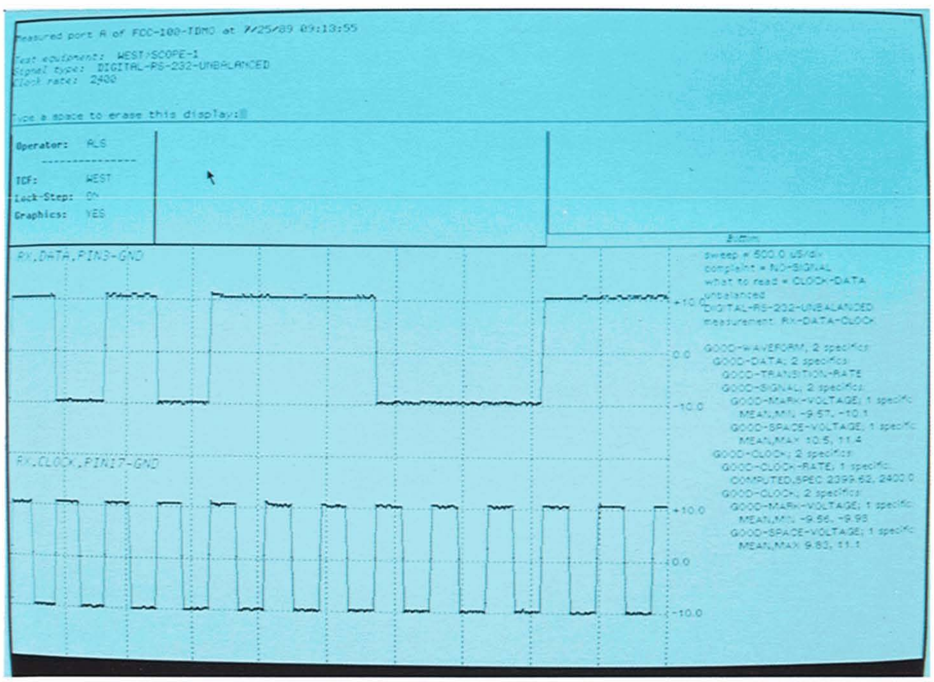

FIGURE 4. Example of MITEC display of waveform evaluation. The computer screen shows the clock and data waveforms that were read into MITEC from an oscilloscope for the example of Figure 3. The screen also displays the logic that analyzed the waveform to determine that the signal at that test point was good.

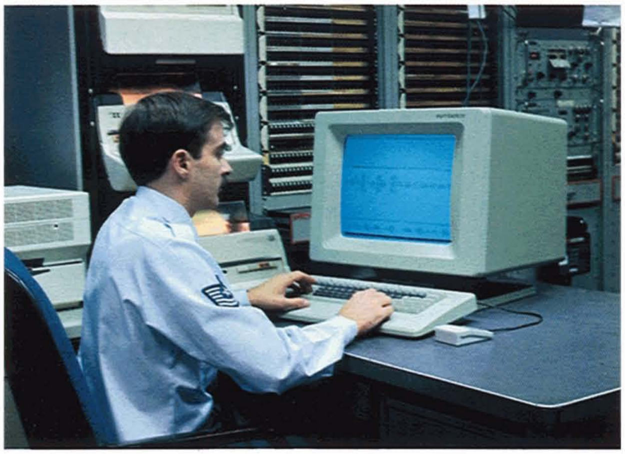

FIGURE 5. Demonstration of MITEC at Andrews Air Force Base, Maryland.

nal preparations to install there Release 1.0 of the software. Release 1.0, a reimplementation for field deployment of the final MITEC laboratory prototype, runs under UNIX on personal computers based on Intel 386 and 486 microprocessors. Written in the Ada language, Release 1.0 was developed primarily by a Lincoln Laboratory subcontractor in collaboration with the software professionals at the MITEC PMO, who have since assumed responsibility for the system.

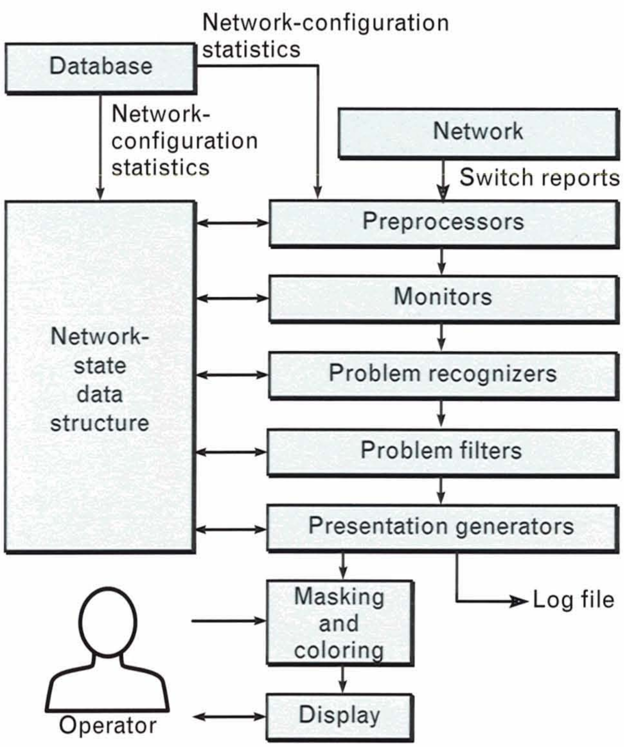

FIGURE 6. Architecture of NMES.

The MITEC PMO expects to install MITEC at numerous military locations as funding permits.

Network Management Expert System (NMES)

The Defense Information Systems Network (DISN) incorporates several voice and data subnetworks, as indicated in Figure 1. Management of these networks is conducted by operators at the operations centers in each of the two overseas theaters, and this is the appropriate location for the Network Management Expert System (NMES). The purpose of NMES is to leverage the performance of the manager on the largest and most challenging of the DISN subnetworks, namely, the theater-wide dial telephone system called the Defense Switched Network (DSN).

DSN was created to provide instant voice connectivity among critical military users and commanders in CONUS and overseas in the event of an emergency, and to serve reliably the more routine telephone networking needs of the DoD at all other times. To this end, dedicated DSN telephone switches are widely distributed at U.S. military bases and centers, and are linked with each other by DoD-owned or -leased trunk circuits. Five levels of precedence have been implemented so that when it is necessary a high-level user can automatically preempt a lower-level call in progress.

The primary goals of DSN network management are to ensure the availability and reliability of DSN services. These goals are accomplished by polling every switch in the theater every few minutes for a statistics report that summarizes all the activity at the switch (for example, call attempts, failures, preemptions, and successes), monitoring these data to recognize and diagnose problems, and performing corrective actions as necessary. Obviously, ensuring the availability and reliability of DSN services is a very difficult and expertise-intensive job. In spite of software tools that gather and display the data automatically, network-management operators have difficulty keeping up because the information comes in at a high rate and the diagnosis problem domain is very complex. Operator training is a serious concern, especially because many of the operators are military personnel who tend to be transferred and replaced from time to time. In dealing with these problems, the Defense Information Systems Agency (DISA) sponsored a project at Lincoln Laboratory to develop NMES.

The architecture of NMES is illustrated in Figure 6. Incoming switch reports are preprocessed for presentation to a bank of monitors, each a relatively simple module that looks for a particular aberration or discrepancy in the data. Next, problem-recognizer modules look for patterns of monitor outputs to decide whether network problems exist. Problem filters then determine where to direct the operator's attention. For example, when a detected primary problem causes a number of other detected problems as side effects, the filters will allow the operator to ignore the side effects and address only the primary problem. Finally, presentation generators format (for presentation to the operator) the problem-diagnosis information as well as instructions for actions that the operator should take.

Figure 7 contains an actual NMES display based on archived data recorded live at the Pacific Theater Operations Center on 24 February 1992. The rectangular nodes in the diagram represent all of the DSN switches in the theater, labeled with three-letter mnemonics that are familiar to the operators. For example, SCS (near the center of the diagram) is the switch at Schofield Barracks on Oahu, Hawaii, and FGB (below and slightly to the right of SCS) represents the switch at Finegayan Bay in Guam. The lines linking the switches represent multicircuit trunk groups. For the different switches and trunk groups, color is used to highlight important factors. For example, a red trunk group indicates that faults have been diagnosed on that trunk, a red top half of a switch icon indicates that faults have been diagnosed within that switch, and a white square in the lower left corner of a switch icon indicates that NMES has information to report about that particular location. By clicking the computer mouse on a switch that has a white icon, the operator can bring up one or more recommendations windows listing the details of problems at that location, as diagnosed by NMES.

Figure 8 shows one of the recommendations windows that was brought up by clicking a mouse on the FGB switch icon. Here NMES is announcing possible hits, i.e., intermittent outages, on trunks connected to FGB. The "Actions" section tells the user what to do about the problem. In this instance, the recommendation was to take no action unless the problem recurred frequently. The "Observations" section substantiates the diagnosis by describing, in terms familiar to DSN network managers, the current symptoms and the reasoning that led up to the diagnosis. The lower half of the window describes another problem that NMES diagnosed at this switch, namely, exces-

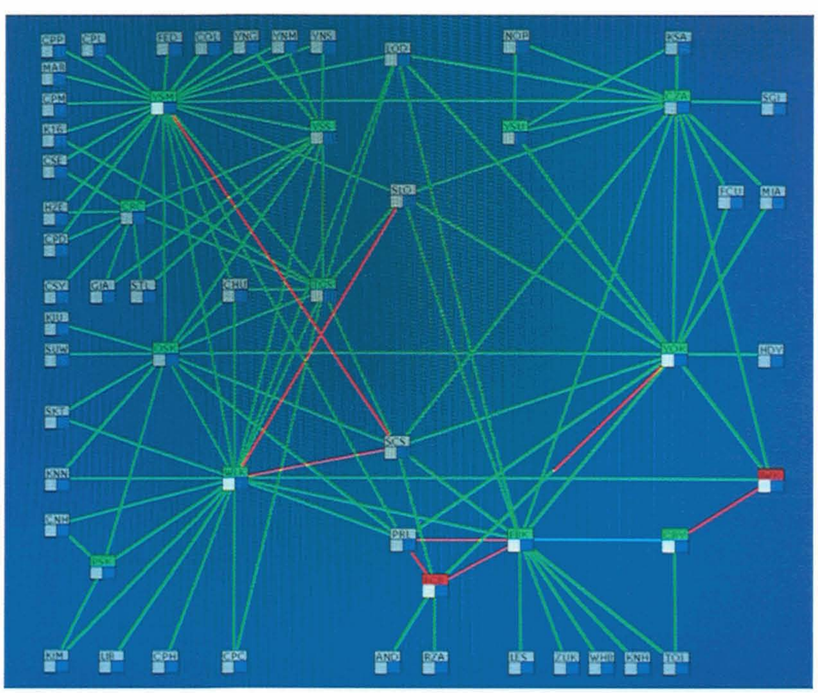

FIGURE 7. Example of NMES display. The nodes in the diagram represent all of the Defense Switched Network (DSN) switches in the theater (labeled with three-letter mnemonics, e.g., SCS for Schofield Barracks on Oahu, Hawaii), and the lines linking the switches represent multicircuit trunk groups. Color is used to highlight important factors. For example, a red trunk group indicates that faults have been diagnosed on that trunk, a red top half of a switch icon indicates that faults have been diagnosed within that switch, and a white square in the lower left corner of a switch icon indicates that NMES has information to report about that location. By clicking the computer mouse on a switch that has a white icon, the operator can bring up one or more recommendations windows listing the details of problems at that location, as diagnosed by NMES.

```
IWUI:IWES Recommendations - Site: FGB (920224.0330)

Problem:
 Possible TG hit.
 [PR5 - TG Hit?]

Actions:
 No action is indicated unless situation recurs frequently.

Observations:
 Symptoms of a possible TG hit are present.
 MF RCVR overflows = 1876.0 per hr.
 RADR delays exceeding the Lower threshold = 16.0 per hr
 RADR delays also exceeded the Higher threshold at 8.0 per hr
 Permanent signal conditions were observed at 140.0 per hr
 The following trunk(s) show symptoms that may be of interest:
   PRLIN1 ( 207) InFails = 104.0 per hr, InCalls = 780.0 per hr, HoldingTime = 18.0 sec.
   PRLIN1 ( 207) SBUsage accounted for 5.6% of possible trunk usage.
   YOKIN1 ( 201) InFails = 348.0 per hr, InCalls = 1456.0 per hr, HoldingTime = 17.1 sec.
   SCSIN1 ( 205) InFails = 44.0 per hr, InCalls = 152.0 per hr, HoldingTime = 65.9 sec.
   SCSIN1 ( 205) SBUsage accounted for 2.8% of possible trunk usage.
   FBKIN1 ( 204) InFails = 228.0 per hr, InCalls = 1428.0 per hr, HoldingTime = 12.3 sec.

 Since 1992 Feb 24 02:00, TG hits have been observed at this switch 1.0 times.

Problem:
 Some Mf Receiver Overflow

 [PR12 - Mf Receiver Overflow]

Actions:
 Check with site if this persists

Observations:
 PR12 - Current Mf Receiver Overflow rate is 1876.0 per hr.

History:
 PR12 - There has been Mf Receiver Overflow for 15.0 min.
 PR12 - RADR:RadLdlypMf has been < 24.0 for 105.0 min.
 PR5 - Possible TG hit.
 Mf Rec Holding Times in secs. Mean 1.4 Min 1.0 Max 2.0 Current 2.0 Nr of samples 13.0
 Mon1 - CCB Seizures: min = 1044.0, max = 5580.0 at 1992 Feb. 24 03:15, currently 5580.0 per hr.
```

FIGURE 8. Example of NMES diagnostic information for the switch located at Finegayan Bay in Guam (FGB). This example is a representation of a computer screen display that was brought up by clicking a mouse on the white square below the red FGB rectangle in Figure 7. Note that NMES has listed two problems at the FGB switch site: (1) a possible hit, i.e., intermittent outage, on a trunk connected to FGB, and (2) excessive demand for multifrequency tone-dialing receivers.

sive demand for multifrequency tone-dialing receivers, which is another symptom of noisy circuits. For this particular case, NMES has listed the history of earlier evidence that further substantiates the current diagnosis.

Currently, NMES is operationally deployed at the European Theater Operations Center in Stuttgart, Germany, and the Pacific Theater Operations Center at Wheeler Army Air Field in Hawaii. At those sites, NMES is incorporated in the DSN Integrated Management Support System (DIMSS)—a networked set of workstations and software that has replaced an old and very limited PC-based network-management support tool used earlier. Developed by DISA's field support contractor GTE Government Systems over the past three years, DIMSS provides a range of essential manually operated services. In addition to the fully automated functions of NMES as described above, these services include the maintaining of network-configuration databases, the polling of telephone switches, and the displaying and archiving of data. (Note: Figures 7 and 8 are actually from DIMSS displays.)

NMES has since come to be called the Integrated Workstation Expert System (IWES). At the present time IWES is being turned over to DISA for maintenance and support, and Lincoln Laboratory will no longer be working on the system.

An important point about IWES is that practicing experts in the user community worked with Lincoln Laboratory to prescribe everything—problem selection, diagnostic logic, and even the wording of the

FIGURE 9. Simulation of the TRAMCON system for the Digital European Backbone II-A (DEB II-A). Used by the Department of Defense as its mainstay transmission system in Europe, DEB II-A is a microwave radio network with strings of mountaintop repeater stations providing backbone interswitch trunk circuits for voice and data networks. The simulation shown here is for a segment of DEB II-A in Germany, beginning at Donnersberg (DON), continuing through Langerkopf (LKF) and Stuttgart (SGT), and ending at Reese-Augsburg (RAG).

recommendations text. Lincoln Laboratory programmers implemented all of this information as given and, as a result, users have been very receptive to the system. A recent new release of IWES, called the User-Programmable IWES, brings the system even closer to the user community by allowing users in the field to modify or add problem categories. In fact, the User-Programmable IWES enables users, even those with few or no software skills, to implement and experiment with new problem diagnostic rules. When a user has a new diagnosis fully tested and working, the status of the diagnosis can be changed from experimental to permanent without the risk of altering or crashing the software.

TRAMCON Adjunct Processor (TAP)

The transmission facilities indicated at the bottom of Figure 1 include numerous DoD-owned and -leased media. In Europe, the mainstay transmission system is a microwave radio network—the Digital European Backbone II-A (DEB II-A)—with strings of mountaintop repeater stations throughout the theater providing backbone interswitch trunk circuits for the voice and data networks. This connectivity is critical to the reliability of the networks; hence, whenever any key equipment item in the DEB II-A stations fails, control automatically switches over to a redundant backup system.

Nevertheless, operator monitoring of DEB II-A stations is required because there are many nonredundant systems in each station. To facilitate this monitoring, the DoD developed the Transmission Monitoring and Control (TRAMCON) system some years ago to permit unmanned operation of remotely located DEB II-A repeater stations. In TRAMCON, each station is provided with a data-logging device connected to 50 or more signals and parameters for the monitoring of, for example, equipment alarms, received power levels, the on/off statuses of tower lights, and the fuel levels of diesel generators. A TRAMCON master station at the end of each 10-to-15-station network segment polls all the stations in the segment every few minutes, gathering hundreds of data points that an operator must then evaluate.

The problem with the existing (and, in fact, quite old) TRAMCON system is that even a simple failure at one station can cause a bewildering array of primary and sympathetic alarms at the trouble site as well as at numerous other sites. For example, loss of a T1 data signal (a standard 1.544 Mb/sec signal from 24 telephone circuits multiplexed together) at one station can cause loss-of-synchronization alarms to pop up all over the network. Only a highly skilled operator can deal with this flood of data.

To allow a less skilled operator to control many TRAMCON stations, Lincoln Laboratory is currently developing the TRAMCON Adjunct Processor (TAP) Expert System. As a side benefit of this work, TAP will replace the TRAMCON processor, an aging HP1000 computer for which repair parts are no longer obtainable.

To support the development of TAP, we created a TRAMCON simulator that reproduces the alarm-propagation behavior of the real DEB II-A network. Figure 9 is a top-level TRAMCON simulator display of a segment of the DEB II-A system in Germany,

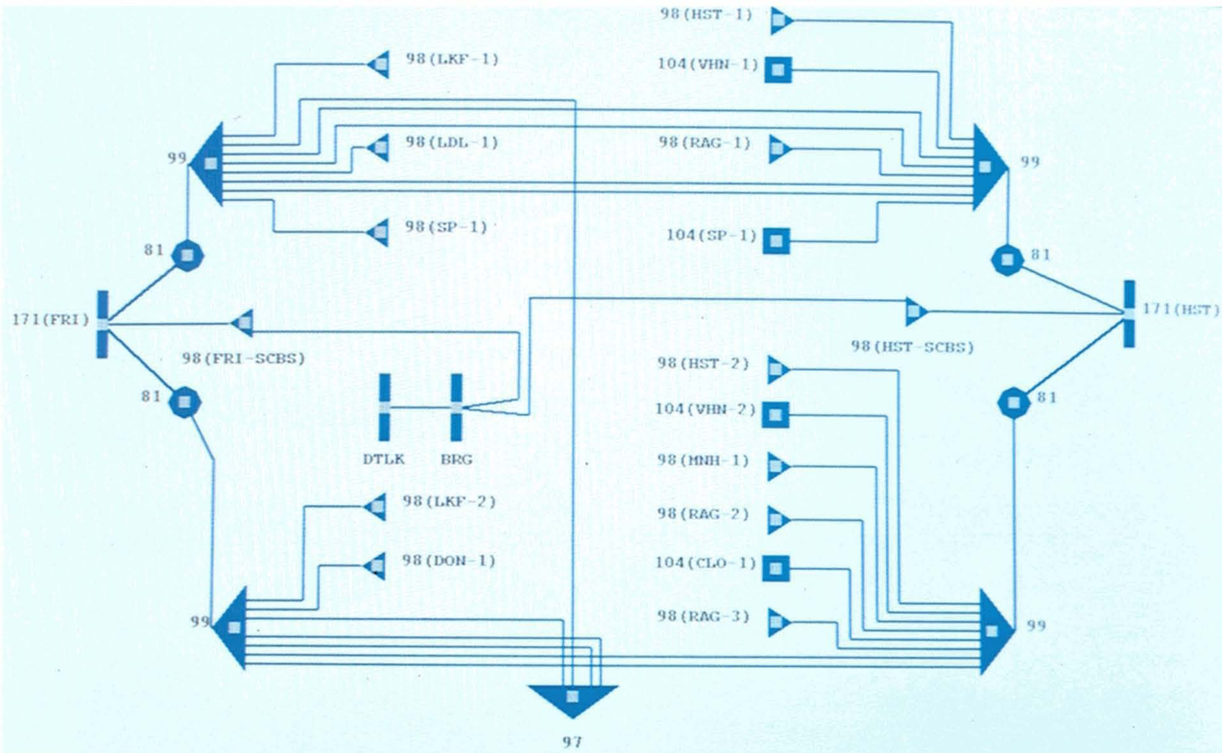

FIGURE 10. DEB II-A equipment suite at Stuttgart, Germany. This computer screen display was brought up by clicking a mouse on the SGT icon in Figure 9.

fault ∇

- PORT-PCB
- PORT-STRAPPING
- RX-DATA-INT
- TX-DATA-INT
- RX-CLOCK-INT
- DEMUX-STRAPPING
- MUX-STRAPPING
- PERIPH-STRAPPING
- EXT-POWER
- PERIPH-PCB
- PWR-SUPPLY-PCB
- MUX-PCB
- IN-LOOP
- DEMUX-PCB
- TIMING-SOURCE

FIGURE 11. Possible fault conditions for the AN/FCC-99 multiplexer. This menu was brought up by clicking a mouse on the AN/FCC-99 icon in the upper left of Figure 10.

Alarms for DEBIIA

SGT MAJOR FRC-171-SGT-FRI RADIO MBS-1 XMT Failed [A or B]
SGT MINOR FRC-171-SGT-FRI Radio B Side Failure
SGT MINOR FRC-171-SGT-FRI Radio A Side Failure
SGT MAJOR FRC-171-SGT-FRI Radio A and B Failed
SGT MINOR FCC-99-SGT-LKF-MBS-1 TDM XMT MBS Data Loss
SGT MINOR FCC-98-SGT-LKF-1 44CNP1 SGT 1-1 to LKF 1-1 1st Level Mux
LKF MINOR FCC-98-LKF-SGT-1 44CNP1 LKF 1-1 to SGT 1-1 1st Level Mux
LKF MAJOR KG-81-LKF-SGT-MBS-1 CRYPTO Failed
LKF MAJOR FRC-171-LKF-FRI RADIO MBS-1 RCV Failed [A or B]
LKF MINOR FRC-171-LKF-FRI Radio B Side Failure
LKF MINOR FRC-171-LKF-FRI Radio A Side Failure
LKF MAJOR FRC-171-LKF-FRI Radio A and B Failed
LKF MINOR FCC-99-LKF-SGT-MBS-1 TDM RCV MBS Data Loss
LKF MINOR FCC-99-LKF-SGT-MBS-1 TDM Frame Loss
LKF MAJOR FCC-99-LKF-SGT-MBS-1 TDM Frame Error Seconds Red
LKF MAJOR FCC-99-LKF-SGT-MBS-1 TDM Frame Error Count Red
LKF MINOR FCC-99-LKF-SGT-MBS-1 TDM Output Port Loss - B Side
LKF MINOR FCC-99-LKF-SGT-MBS-1 TDM Output Port Loss - A Side
LKF MAJOR FCC-99-LKF-SGT-MBS-1 MBS Demux A and B Output Loss
LKF MINOR FCC-98-LKF-HST-1 44CNZ2 LKF 1-7 to HST 1-7 1st Level Mux
HST MINOR FCC-98-HST-LKF-1 44CNZ2 HST 1-7 to LKF 1-7 1st Level Mux
FRI MAJOR FRC-171-FRI-SGT RADIO MBS-1 RCV Failed [A or B]
FRI MINOR FRC-171-FRI-SGT Radio B Side Failure
FRI MINOR FRC-171-FRI-SGT Radio A Side Failure
FRI MAJOR FRC-171-FRI-SGT Radio A and B Failed
FRI MAJOR FRC-171-FRI-LKF RADIO MBS-1 XMT Failed [A or B]
FRI MINOR FRC-171-FRI-LKF Radio A Side Failure
FRI MINOR FRC-171-FRI-LKF Radio B Side Failure
FRI MAJOR FRC-171-FRI-LKF Radio A and B Failed

FIGURE 12. Alarm list caused by the false loopback condition "IN-LOOP," which was highlighted in Figure 11.

beginning at Donnersberg, continuing through Langerkopf and Stuttgart, and ending at Reese-Augsburg. Using a mouse to select one of the high-level icons in Figure 9 (e.g., "SGT" for Stuttgart) brings up a window showing a detailed diagram of the equipment that is actually installed at that station (Figure 10), including microwave radios, encryption/decryption units, and second- and first-level multiplexers. Similarly, using a mouse to select any of the equipment icons in Figure 10 (e.g., the AN/FCC-99 multiplexer in the upper left) brings up a menu of all of the fault conditions that can occur in that unit (Figure 11). One of the fault conditions has been selected in Figure 11, as indicated by the shading of "IN-LOOP," which stands for the erroneous operation of the internal loopback switch, i.e., the accidental disconnection of the multiplexer from the network. For this particular fault condition, the simulator produces a list of 29 primary and sympathetic alarm signals (Figure 12) spread over four DEB II-A stations, and all caused by that false internal loopback at Stuttgart. With the old TRAMCON system, an operator would need great skill and experience to recognize the cause of this excessively complex alarm pattern.

At this point we stop the simulator and pass the alarm list to the TAP Expert System just as though TAP were receiving the list live at a TRAMCON master station. TAP then does exhaustive hypothesis testing, identifying all possible faults on the DEB II-A segment that could have caused the observed alarm pattern. In this case TAP identifies four candidate faults, as shown in Figure 13. Note that all of the faults are in the correct AN/FCC-99 multiplexer at the correct location, namely Stuttgart, and one of them is the false internal loopback that was the actual cause of the alarms. The overall effect of this TAP functionality is that a relatively unskilled TRAMCON operator can dispatch one technician to the correct site with a small set of repair parts and tests to run, one of which will correct the problem (as opposed to the operator requesting several repair teams to investigate a multitude of alarms at four repeater stations).

TAP is a more recent project than MITEC and IWES, and is currently in an intermediate stage of development. The data-acquisition front end of

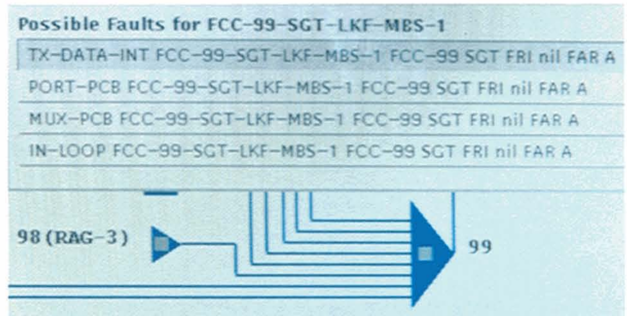

FIGURE 13. Candidate faults identified by the TAP Expert System for the alarm list of Figure 12. Note that one of the faults is the true cause of the alarm list, namely, the false internal loopback ("IN-LOOP") in the AN/FCC-99 multiplexer at the Stuttgart station.

TAP (which collects a copy of the data being received by a TRAMCON master station) was successfully demonstrated on DEB II-A in September 1993 at the U.S. Army 5th Signal Command in Worms, Germany. A proof-of-concept model of TAP is currently being developed at Lincoln Laboratory and is intended for field testing in Europe.

Expert System for the ASR-9 Airport Surveillance Radar

Lincoln Laboratory has begun a new expert-system initiative for the ASR-9 airport surveillance radar, which the FAA has installed at all large airports in the United States (Figure 14). Because continuous and reliable surveillance radar data are critical to terminal ATC operations, the ASR-9 has a variety of robustness and reliability features. For example, the radar has redundant backups that automatically switch in for key subsystems when a failure occurs in the primary channel. The ASR-9's reliability is further enhanced by a built-in Remote Monitoring System (RMS) that provides technicians access to status indicators, performance parameters, and built-in test results gathered from all the cabinets and subsystems of the radar.

The ASR-9 RMS has elements of all three generic problem domains described earlier as amenable to the application of an expert system: complex, repetitive, and highly skill-intensive operations; vigilant monitoring of status and alarm information; and ongoing verification that performance is within allowable ranges. Recently, two serious ASR-9 problems have

FIGURE 14. The ASR-9 airport surveillance radar, which the FAA has installed at all large U.S. airports.

come to light: (1) the loss of logistics support for the outdated microcomputer that hosts the RMS, and (2) human-factors deficiencies in the operator interface of the RMS. These two problems have triggered a new project to replace the RMS with an expert system.

The logistics problem exists because the ASR-9 design dates back a number of years. At that time, the computer selected to host the RMS was the Intel 310—a 1970s multiboard microprocessor system that was developed before the invention of the personal computer. Since then, Intel has stopped manufacturing repair parts for the 310, and FAA stocks of parts have been dwindling. An outage in the RMS processor is a serious matter because this processor is not merely a passive observer of alarms and status; it is also the primary conduit for operator control of the radar. A modern personal computer or small workstation hosting an RMS replacement in the form of an expert system as described below would solve the Intel 310 logistics problem.

The human-factors deficiencies in the ASR-9 RMS are also attributable to the antiquity of the present design. In its day the menu-driven operator interface of the RMS was quite advanced; by today's standards, however, it is primitive and user unfriendly. Furthermore, the interface is highly expertise intensive, requiring technicians to possess uncommonly sophisticated skills to use it successfully. Specifically, the interface has about 238 menus in nests up to seven layers deep, and the menus contain curt, obscure language but no help functions. Consequently, to solve a typical fault-diagnosis problem, a technician must know the corresponding route of navigation up and down the menu trees to gather various pieces of information, and the technician must know how to correlate and analyze this information. For very unusual problems, an expert ASR-9 technician may need to look directly at the raw status and alarm data, which the RMS stores in a 256-word first-in/first-out (FIFO) memory. Direct use of the raw data requires another kind of arcane knowledge because the data items are interleaved with each other, have neither time stamps nor references, and are stored as 4-character non-mnemonic codes that the technician must either know from memory or look up in a loose-leaf notebook. To make matters worse, access to these raw data can be gained only by stopping the processor and putting it into debug mode, during which time all incoming RMS data will be lost.

The special knowledge of technicians who can handle all the above tasks is rare, and this knowledge is disappearing as experienced technicians age and retire. Not surprisingly, the incoming generation of technicians with normal training and skill levels have great difficulty using the RMS effectively. Often these technicians can only guess at the general area of a suspected fault, and must then swap spares until they find a combination that works. This inefficient process has a costly side effect: every swapped spare that does not correct the problem is suspect, and must be shipped to the depot, tested, and recertified before it can be returned to the stock of ready spares.

To address these concerns, the FAA has asked Lincoln Laboratory to develop an expert system for the ASR-9 RMS. The first step of this large project is Phase I: a twelve-month effort to produce and demonstrate a proof-of-concept model. Working with FAA engineers and technicians, the Phase I developers will select a set of ASR-9 diagnostic and trend-analysis functions, suitably sized for the twelve-month effort and capable of clear and compelling demonstration. Figure 15 illustrates the planned configuration for this initial work. Note that the system will run on a freestanding workstation that nonintrusively taps into the RMS data stream of an unmodified ASR-9, thus allowing risk-free demonstration on an operational radar. The system development will be done on an engineering testbed ASR-9 that is operated in the field by Lincoln Laboratory. During the Phase I period, a statement of objectives, a plan, and a schedule for a Phase II effort will also be developed.

The general goal of Phase II will be to build a well-designed and fully documented model of an ASR-9 RMS expert system that would implement and demonstrate the features and functionality required for a deployable replacement for the RMS in operational ASR-9s. Thus the expert system would need to provide all of the functions of the present RMS, in addition to implementing automated fault diagnosis, trend analysis, and prescription of corrective action for an extensive repertoire of ASR-9 problem conditions. Figure 16 is a high-level view of the probable configuration for the Phase II model running on commercial off-the-shelf (COTS) hardware. During Phase I the necessary ASR-9 RMS knowledge and system experience will be acquired to permit the preliminary identification of suitable COTS hardware items.

Early in Phase II, the FAA Technical Center will be consulted to develop a test plan for a stringent field demonstration and evaluation exercise to which the Phase II system will be subjected on completion. This test plan will play a significant role in guiding the design and implementation of the Phase II system because the results of the field tests will no doubt strongly influence FAA decisions on how to proceed further. The desirable final outcome would be for the FAA to assume maintenance and distribution responsibility for the software, modify it as needed, then deploy it in commissioned ASR-9s in COTS hardware procured by the FAA for that purpose.

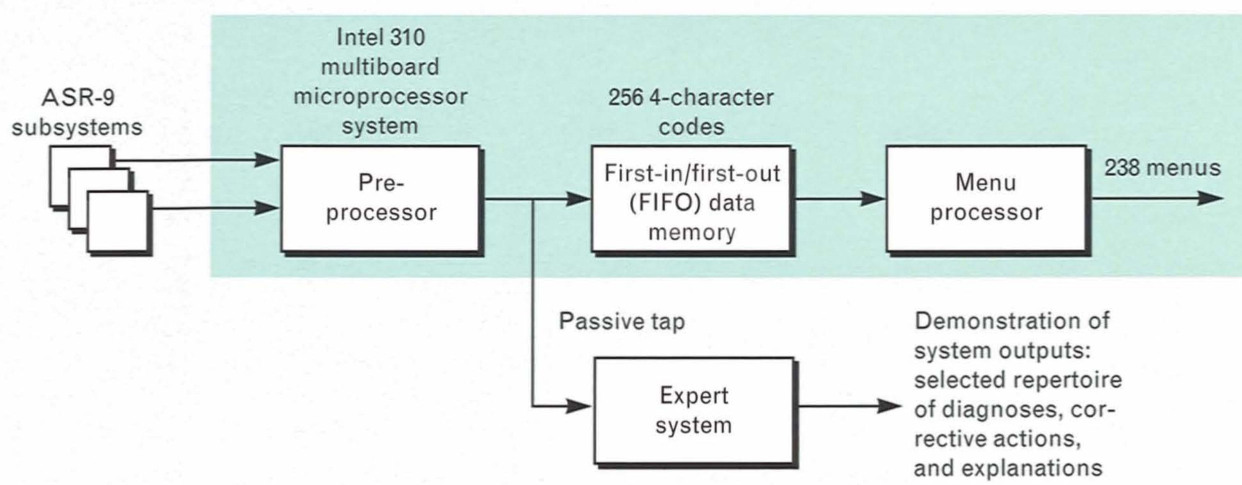

FIGURE 15. Phase I proof-of-concept configuration for an expert system for the ASR-9 Remote Monitoring System (RMS).

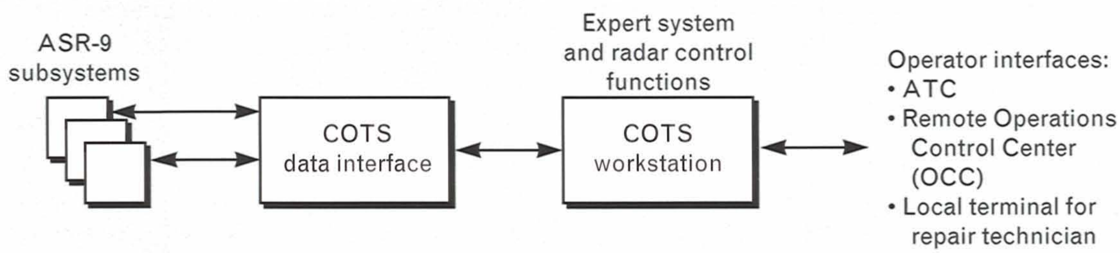

FIGURE 16. Phase II replacement-system configuration for an expert system for the ASR-9 RMS. Note the use of commercial off-the-shelf (COTS) hardware for key components of the system.

FAA Operations Control Center (OCC) Issues

Airports, ATC centers, and other FAA airway facilities throughout the United States deploy a multitude of equipment, including airport and air-route surveillance radars, navigation aids, and communications equipment and networks. The continual availability and proper operation of these FAA installations are critical for the safety of the flying public. In the past, staffs of specialist technicians at numerous work centers near the equipment sites have ensured reliable performance and high availability of the ATC equipment. In recent years, however, such maintenance has become increasingly difficult because of budget reductions, loss of skilled technicians due to retirement, and growing equipment complexity.

For some time the FAA has been moving toward a Remote Maintenance Monitoring System (RMMS) as a cost- and personnel-reduction alternative. In the RMMS structure, a smaller technician force is concentrated at much fewer consolidated work centers, each remotely servicing ATC equipment over a larger geographic area. The ASR-9 RMS was originally designed as an element of the much larger RMMS concept. Unfortunately, difficult problems have arisen in the implementation of RMMS. For example, increasing the number of ATC equipment items remotely monitored by a consolidated work center will lead quickly to an unsupportable work load in terms of the huge amount of incoming information an operator must examine, evaluate, and diagnose. As another example, many ATC equipment items (particularly the older models) were not designed with adequate provisions for providing an off-site technician with all of the essential control and monitoring data necessary to manage the equipment. Thus such equipment items require a retrofitting with more capable remote-access provisions. A third example is the need to train each operator at a consolidated work center to a high level of competence and knowledge of many more categories of equipment than an individual technician is typically trained on in today's environment. Such problems can be relieved through judicious application of expert-system technology. For instance, a suitably designed expert system at a consolidated work center can vigilantly monitor high volumes of status and alarm information by scanning the data for patterns and indicators of trouble and alerting the work-center operator only when a problem exists. Additional operator work-load reduction can be achieved by automating such tasks as the maintenance of equipment and status databases, the fault-tracing procedure in multi-subsystem networks, and the filtering process that separates primary problems from secondary side effects.

For each type of existing ATC equipment that does not now provide adequate remote access, an expert system could be built that is generically similar to the ASR-9 RMS expert system discussed earlier. The expert system could reside in a small workstation in the equipment shelter and, from that location, could examine and diagnose the full set of on-site technician-accessible status and performance data and forward the diagnostic results to the consolidated work center. The training crisis caused by the burgeoning of equipment types could be relieved by embedding much of the specialized equipment knowledge into expert systems, thereby effectively enabling an operator to produce expert-level diagnoses without having been a trained expert on the subject equipment.

Currently, the FAA is interested in a new concept called the Operations Control Center (OCC) to host the consolidated work-center functions for achieving the cost- and personnel-reduction goals stated earlier. In view of the potential expert-system roles in an OCC and in the remote equipment serviced by the OCC, features will be designed into the ASR-9 RMS expert system to make it a model of compatibility with a remote OCC. For example, provisions will be made in the design of the ASR-9 RMS expert system for convenient and effective communication of the system's results and recommendations to another computer. Moreover, the interfaces of the ASR-9 RMS expert system will conform to current FAA interface control documents for remote maintenance monitoring. Thus, at the FAA's option, the Phase II ASR-9 RMS expert system could be interfaced and integrated with a prototype OCC when and if such a prototype becomes available.

Summary

Machine intelligence (MI) technology offers a means for automating much of the work load of an ATC equipment technician, thereby alleviating such growing problems as the economic pressures to reduce costs, the retirement of experienced technicians, the bottlenecks in training new technicians, and the growing complexity of the ATC equipment. The practical feasibility of using MI technology in applications for the Federal Aviation Administration (FAA) has been demonstrated by Lincoln Laboratory in developing and fielding three expert systems that have addressed these problems in a generically similar domain, namely, the maintenance and control of military communications networks.

This article described the three military applications of expert systems, focusing attention on the common threads of the three applications: the capturing and preserving of technician knowledge, the facilitating of staff performance well above staff skill levels, the leveraging of staff productivity, and the freeing of staff from routine tasks so that more difficult problems could be targeted. This article then discussed particular problems that have arisen in the Remote Monitoring System (RMS) of the ASR-9 airport surveillance radar, and described a new FAA-sponsored Lincoln Laboratory program to address those problems by implementing an expert system for the ASR-9 RMS. The relationships among the FAA's older Remote Maintenance Monitoring System (RMMS) development efforts, the more recent plans for an Operations Control Center (OCC) prototype, and the expert-system technology applications were then described.

Acknowledgments

This work was sponsored by the U.S. Air Force and the Defense Information Systems Agency (DISA).

REFERENCES

1. *Lincoln Laboratory Annual Report—Knowledge-Based System Analysis and Control,* ESD-TR-90-103 (30 Sept. 1989), DTIC #AD-A230735.
2. *Lincoln Laboratory Annual Report—Knowledge-Based System Analysis and Control, Defense Switched Network Task Areas,* ESD-TR-90-141 (30 Sept. 1989), DTIC #AD-A230735.
3. *Lincoln Laboratory Annual Report—Knowledge-Based System Analysis and Control,* ESD-TR-90-125 (30 Sept. 1990), DTIC #AD-B152362L.
4. *Lincoln Laboratory Annual Report—Knowledge-Based System Analysis and Control, Defense Switched Network Task Areas,* ESD-TR-90-151 (30 Sept. 1990), DTIC #AD-A234181.
5. *Lincoln Laboratory Annual Report—Knowledge-Based System Analysis and Control,* ESD-TR-91-243 (30 Sept. 1991), DTIC #AD-A249611.
6. *Lincoln Laboratory Annual Report—Knowledge-Based System Analysis and Control, Defense Switched Network Task Areas,* ESD-TR-91-247 (30 Sept. 1991), DTIC #AD-249611.
7. H.M. Heggestad, "Expert Systems in Communications System Control," *1991 IEEE Military Communications Conf., 4–7 Nov. 1991,* p. 0863.

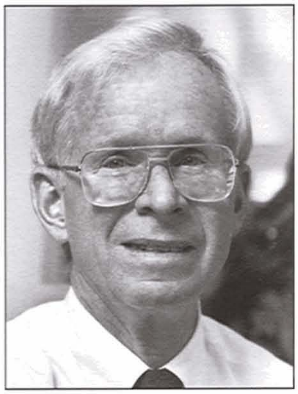

HAROLD M. HEGGESTAD is an associate leader in the Machine Intelligence Technology Group, where his focus of research has been on the analysis and control of knowledge-based systems, specifically, the application of expert-system technology to the monitoring and control of complex systems. His previous staff assignments have been in radar systems and satellite communications.

Hal received S.B., S.M., and Sc.D. degrees from MIT in electrical engineering, specializing in communications theory. He is a senior member of the IEEE, an associate fellow of the AIAA, and a member of Tau Beta Pi, Eta Kappa Nu, and Sigma Xi. His interests include bicycling, cross-country skiing, fishing, and wood carving.