# NASA Flight Controllers Become AI Pioneers

*NASA's mission control center, which has not been fundamentally upgraded since Apollo, is now leaping straight into the 1990s*

FOR THE BETTER PART of three decades now, first for Apollo and more recently for the space shuttle, flight controllers at the National Aeronautics and Space Administration's (NASA's) Johnson Space Center in Houston have been staring into a nightmare of user unfriendliness: monochrome computer screens covered with column after column of cryptically shifting acronyms. This mainframe-processed telemetry data is supposed to tell them whether or not their particular subsystem of the shuttle is headed for disaster—and even the old-timers have been known to miss things.

All that is beginning to change, however. Starting with the shuttle's first post-Challenger launch in September 1988, the normally cautious controllers have been stepping straight from the technology of the 1960s to the technology of the 1990s: high-powered workstations, high-resolution color graphics, the works. They are even taking the plunge into artificial intelligence (AI), a software technology that still has more of a reputation for hype than for practical accomplishment.

"It's been tremendous fun," says John F. Muratore, a veteran of 17 shuttle missions as chief communications officer and now the prime mover behind the changeover. The color graphics alone make an enormous difference, he says. Instead of watching inscrutable alphanumerics, the controller sees a color-coded schematic diagram of, say, the shuttle's communication system or its main engines. If anything goes wrong, the AI-based software displays a variety of diagnostic messages and causes the affected components to light up red. The problem is virtually impossible to miss.

While only a handful of the older consoles have been replaced so far, says Muratore—each piece of new software first has to be tested and compared with its older counterpart by operating in parallel on a real flight—the response from the controllers themselves has been very positive. "The pace of implementation is actually accelerating, not slowing down," he says. "We have more requests for development of applications [in different areas of shuttle flight control] than we can fill."

To create these applications, the Houston controllers draw heavily on a type of AI program known as an expert system, in which knowledge derived from human experts is expressed as a series of rules: "If *this* is the case, then do *that*." They are hardly unique in that effort. Expert systems have been widely touted as AI's first genuinely useful spinoff, with applications ranging from medical diagnosis to the interpretation of seismic data. But their deployment in the shuttle mission control center marks one of the first times that anyone has used the output of an expert system to make split-second, life-or-death decisions.

On the other hand, says Muratore, these expert systems are emphatically *not* the kind of mythic, all-knowing oracles that have so often been portrayed in the popular press. In particular, they are not replacing the human experts. "People have gone out and tried to build these 'Grand Oracles' for routine use, and I think they've failed miserably," he says. The software techniques available to program such an oracle are not nearly mature enough, and the commercial hardware available to run it is not nearly powerful enough. Instead, Muratore and his colleagues are integrating the expert systems programming techniques into a set of conventional programs with the aim of creating intelligent assistants: software applications that can give the controllers relief by turning the flood of telemetry data into something meaningful.

Imagine, for example, that the stream of telemetry coming down from the shuttle's main engines started to show a pattern of anomalies. The control center's current generation mainframe computers would be able to detect such a pattern very quickly, says Muratore. And indeed, if it were just a matter of diagnosing the symptoms of isolated faults in isolated components of the engines, they could even figure out what the pattern meant. The problem is that the engines' multitudinous pumps and valves and fuel lines are not isolated; a failure in one place can easily lead to multiple failures elsewhere—at the same time that multiply redundant backup systems may be kicking in. The possible number of patterns is astronomical, and the mainframes' conventional programming algorithms would quickly bog down in the overwhelming number of combinations. This is why the interpretation of the data has generally been left to the human controllers, says Muratore—and this is why even they could use some help.

Thus the appeal of expert systems, he says: the rules can incorporate an expert controller's knowledge about the physical structure of the engines and the way faults tend to propagate, which means that the program can reason its way through the various combinations of faults and reach a conclusion much

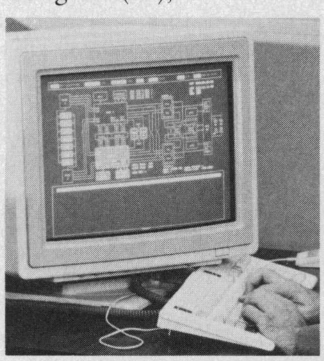

NASA

**The old and the new.** *Communications officer John F. Muratore (below) has led the charge to replace the space shuttle mission control center's old monochrome displays with high-resolution color graphics workstations (left).*

# AI Is The Able Assistant

After years of media hype and grand promises, artificial intelligence (AI) is finally finding some real-life applications—although they are seldom the kind of dramatic schemes being ballyhooed just a few years ago. The shift in emphasis was readily apparent at Stanford University recently during the first national conference on successful AI applications.* In the early 1980s, said conference cochair Alain Rappaport, president of a small expert systems company in Palo Alto, California, known as Neuron Data, millions of dollars were invested in "knowledge engineering" to extract the appropriate rules from human experts, and more millions were invested in special-purpose LISP machines just to run the programs. Yet today those efforts are beginning to look like fossils. With rare exceptions, he says, "the old, heavy AI" is simply not economically feasible.

Nowadays, says Rappaport, successful AI applications are starting to follow the same pattern that NASA is trying in its shuttle flight control system (see accompanying story): that is, start with a set of conventional programs, and then use AI as a kind of "technological glue" to link them together—and to help the user with the inevitable avalanche of data and choices. "Here is a database that means nothing to me," he says. "But by using AI, I can collect and illuminate just the right data and enhance the decision-making process. I can enhance the activities that are inherently human."

In organizing the Stanford conference, Rappaport and his colleagues deliberately limited the talks to AI applications that were in productive, everyday use—as opposed to prototypes, laboratory demonstrations, and bright ideas. "If we're promising people miracles at Lourdes," said expert systems pioneer Bruce Buchanan of the University of Pittsburgh, "at least we ought to count the crutches."

Consider just two examples:

■ TARA—the Technical Analysis and Reasoning Assistant—is an expert system developed at the Manufacturers Hanover Trust Company in New York to aid foreign currency traders in making their decisions to buy, sell, or hold. This is a high-pressure environment, to say the least. Traders are constantly bombarded by pricing data on 15 major currencies and a host of minor ones. They communicate with each other via 120-line telephones. They always have their lunches brought in because they might lose a fortune if they ever took the time to go out. And even the best traders are pleased if their predictions are right as much as 60% of the time.

TARA is accordingly an effort to ease the level of frenzy while simultaneously improving the odds. The system takes in a stream of live financial data—with "dead" data being defined as older than 3 seconds. Then it automatically analyzes the data for trends, presents the user with an on-screen display of the essential facts that he or she needs to know (but no more), and makes recommendations for trades. In short, it serves as a high-level assistant, so that the trader does not have to monitor the situation constantly.

■ The Force Requirements Expert System (FRESH) is a naval battle management expert system developed by the Defense Advanced Research Projects Agency as part of its Strategic Computing initiative. It has been in regular use by the staff of the Commander in Chief of the U.S. Pacific Fleet since August 1987. The problem they face is to coordinate the movements of nearly 300 ships, including aircraft carriers, submarines, destroyers, and tankers, together with more than 2000 aircraft, spread over some 95 million square miles of ocean and 2450 ports of call. Any plan for deploying these forces has to take into account such factors as the time available to meet a crisis situation, and the combat readiness of each individual unit—which may be very different in reality than it is on paper. Devising such a plan by hand is inefficient, at best, and often results in important factors being ignored.

The FRESH system attempts to do better by using expert system techniques combined with existing databases and command and control systems. As a resource manager it monitors the current state of the fleet and evaluates the impacts of any changes in plan; reports that previously required hours can now be produced in 30 seconds, with accuracy improved from 90% to 99%. And as a planning assistant, FRESH can generate and evaluate some 100 options in just a few minutes, versus the hours it used to take to evaluate perhaps a dozen options by hand. ■ **M.M.W.**

*Conference on Innovative Applications of Artificial Intelligence, 28 to 30 March, Stanford University, Palo Alto, California.

more efficiently than a conventional algorithm can. The upshot is that the control center's new workstations can display not just data, but interpretations that help controllers make better decisions faster.

Interestingly enough, one of Muratore's strongest supporters is Eugene F. Kranz, the Johnson Space Center's Director of Mission Operations and one of the people who originally created the mission control center back in the early 1960s. "When you look at how John has been received by the other controllers, it's incredible," says Kranz. "They give him a problem and he can give them a solution in just a few weeks." Muratore has put together his own tiger team to implement the new technology, and has even had black silk jackets made up for the team members sporting the Mission Operations Directorate logo. "John's gotten too valuable to use in routine console operations," says Kranz.

"The only problem," he adds, "is that we have to grow out of the old system and into the new—and that means big bucks." Replacing the whole existing system means putting in workstations for more than 100 controllers at a time, not to mention all the necessary backups and redundancy. Now add in enough expanded capacity to handle ground control for the space station starting in the mid-1990s—assuming that the space station is really built—and the bill will come to several hundred million dollars. Given his current budgetary situation, says Kranz, he is going to have to string out that expenditure for the better part of a decade, although he is now trying to reprogram some funds to speed things up.

Lending urgency to this changeover is the two-tiered structure of control center demographics: one group of relatively young controllers who were brought in for the shuttle era, and one group of older, more experienced controllers who originally signed on in the Apollo era—and who are now starting to retire. "I'm losing the whole 45-and-older half of my organization," says Kranz. So if NASA really does manage to fly the shuttle at least a dozen times a year, as it hopes to, and especially if the agency gets the go-ahead for its space station, those remaining controllers are going to need a lot of help—which is precisely what Muratore and his artificially intelligent machines can provide.

As Muratore himself points out, the new systems will not only allow the control center to get by with fewer controllers per shift, but they should ease the training bottleneck. He believes that the time it takes to bring a novice controller up to speed could be cut from 2 years on the old system to about 18 months on the new—a savings of 25%. ■ **M. MITCHELL WALDROP**