AI Magazine Volume 12 Number 1 (1991) (© AAAI)

The Use of Artificial Intelligence by the United States Navy: Case Study of a Failure

Captain Stephen B. Sloane, United States Navy, Retired

This article analyzes an attempt to use computing technology, including AI, to improve the combat readiness of a U.S. Navy aircraft carrier. The method of introducing new technology, as well as the reaction of the organization to the use of the technology, is examined to discern the reasons for the rejection by the carrier's personnel of a technically sophisticated attempt to increase mission capability. This effort to make advanced computing technology, such as expert systems, an integral part of the organizational environment and, thereby, to significantly alter traditional decision-making methods failed for two reasons: (1) the innovation of having users, as opposed to the navy research and development bureaucracy, perform the development function was in conflict with navy operational requirements and routines and (2) the technology itself was either inappropriate or perceived by operational experts to be inappropriate for the tasks of the organization. Finally, this article suggests those obstacles that must be overcome to successfully introduce state-of-the-art computing technology into any organization.

Organizations are adaptive systems that continually attempt to push the limits of their own effectiveness to approach perfection. This approach is true of the "mom and pop" store that is threatened by the growth of shopping malls. It is true of the gigantic corporation that is threatened by public regulation and private competition. It is particularly true of organizations that are confronted with complex tasks, the vagaries of uncertainty, and the high and visible costs of irreversible error.

The cause of organization ineffectiveness or, indeed, failure is often perceived to be human frailty (Perrow 1984). The rationality of human participation has been described as bounded, in large part because all human decisions have both a value and a factual premise (Simon 1957). Organization performance can be degraded when the value obscures the factual premise. As the complexity and scope of the organization task increase (imagine, for example, the sorts of tasks envisioned by the system proposed in the Strategic Defense Initiative), it is natural that organizations will try to adapt by unbounding the rationality of the organization process. In an earlier era, this adaptation involved developing and using scientific management techniques, where the one best way to objectively do the job would be sought out and implemented (Taylor 1911). Latter-day attempts to unbound the limits of organization effectiveness, for example, the planning programming and budgeting system (PPBS) or the use of telemetry to monitor the performance of technical systems, involve management control schemes designed to constrain the weaknesses of human decision making. Often, these schemes involve the use of computing technology in general and AI in particular.

What does it take to make computer-assisted attempts to approach perfect effectiveness successful? What are the reasons that such attempts can fail? Obviously, the answers to

0738-4602/91/$4.00 ©1991 AAAI

these questions must be searched for in a technical domain. The technical knowledge required to have machines help humans make more effective decisions is certainly necessary to achieve increased effectiveness; it is not, however, enough. In addition, we must explore ways to manage the introduction of computing technology into the complex environment of modern organizations. This article examines this process by exploring the case of a U.S. Navy aircraft carrier that was a pilot project designed to use computing technology, including AI, to optimize combat readiness. In this case, we see the organization reject a technically sophisticated attempt to solve the problem of the limits of organization effectiveness. This article is the case study of a failure. What lessons can be learned?

This article centers on the aircraft carrier USS *Starship*. Aircraft carriers are useful subjects to study if we are interested in the process of organizational adaptation to technological change (or to anything else for that matter). On an aircraft carrier, the behavior of participants is likely to be adaptive; at least, there is a strong incentive to do things right. The reason for this is that the result of maladaptive behavior, particularly decision behavior, is often unambiguously evident in the short run. These ships are part of a category of organization where although the probability of maladaptive behavior is small, the consequences of error are highly visible and can involve great cost. Aircraft carriers such as the *Starship* operate supersonic aircraft in the most demanding of circumstances. The ship itself is powered by a nuclear power plant. If the operational requirement arises, the ship stores, maintains, and deploys nuclear weapons. All these special functions notwithstanding, the sea is a generally unforgiving environment. Disastrous fire, collision, grounding, and flooding lurk around every corner.

The aircraft carrier, as an organization, has to adapt to its task. If it does not, failure is unambiguous. The narrow range of right choices and the clarity with which evidence of wrong choices becomes manifest result in an organization process that is likely to be adaptive. On board an aircraft carrier, lessons are more often than not written in blood. In addition, lessons learned at the cost of blood are likely to become part of the corporate memory whether they become part of the formal structure. Because of this and because the navy assigns only officers who have proven to be effective in highly demanding jobs to positions of significant authority aboard the carrier, we expect that the reaction of carrier professionals to research and development (R&D) projects will tell us something meaningful about the value of the specific project; the general usefulness of the fundamental concepts that are the foundation of the project; and, most important, the process of managing technical innovation.

The names of the people and the ship are fictional. The events are not fiction. In gathering the data, I entered the work environment as an observer in the role of social scientist and directly observed the events I describe. I conducted unstructured interviews that elicited both opinions and descriptions. I served as an active duty naval officer from 1953 to 1983 and experienced both the intrinsic and extrinsic positive and negative rewards of a naval career. Accordingly, the images described here, as well as the meaning that is ascribed to these images, were filtered through perceptual lenses tinted by previous experience and involvement. Nevertheless, I hope that whatever distortion results from having played the role of navy professional will be compensated by whatever intuitive advantage was gained as a result of the experience.

This article analyzes an attempt to use computing technology… to improve the combat readiness of a U.S. Navy aircraft carrier.

The Premise

The idea for using computing technology to replace many of the more traditional methods of organization process aboard *Starship* was that of its first commanding officer, Captain Smart. In a previous tour of duty as the executive officer (second in command) of another ship, Smart had been frustrated by the administrative complexity of the organization and the difficulty of controlling the many paperwork chores that were his responsibility. Smart's cure for this administrative frustration was to use a personal computer that he had installed in his stateroom to follow up on the many administrative actions that he needed to control. This database of information and schedules did not significantly change the reality of the working environment aboard the carrier. It did change Smart's perception of his ability to maintain control of the flow of information. This perceptual change made his situation at least seem less problematic and, therefore, more tolerable. Even if Smart could not control events, at least he could control the flow of information to himself.

Years later, when assigned as the first commanding officer of the newly constructed USS *Starship* (his role as commanding officer started while the ship was still under construction), he again attempted to use computers to

...we must explore ways to manage the introduction of computing technology into the complex environment of modern organizations.

gain control of a complex organization. Smart believed that effectiveness could be optimized if everything—or nearly everything—could be structured to specify the best way to do all jobs. He wanted to control the organization in a way that would eliminate, as much as possible, the variation between the promises of planning and the realities of implementation.

The use of computing technology as an organization control mechanism aboard *Starship* might approach satisfying the conceptual aspects an ideal Weberian bureaucracy. The goal of near-total specification of structure would be the eventual elimination or, at least, limitation of nonstandard or ineffective behavior. If this goal could be achieved, the perceived randomness of the organizational process might be eliminated. With the random impact of individual behaviors controlled, the organizational process would be much more effective and much more predictable. At least, this was the wish.

The USS *Starship*

The USS *Starship* is an aircraft carrier whose energy for propulsion is derived from nuclear reactors. The *Ship's Organization and Regulations Manual* (SORM), an official document unique to each navy ship that specifies the design of the formal organization structure, indicates that the ship "is to be ready to employ its power anywhere in the world, as directed by the president of the United States." The missions pertinent to this general statement of purpose include antiaircraft warfare, antisubmarine warfare, and warfare against surface ships. *Starship's* mission also includes projection of air power into land areas from the sea. The airwing that flies from the ship, commanded by a carrier airwing commander (CAG), consists of about 90 airplanes of various types that are specialized in carrying out these missions.

The *Starship's* SORM was written by Captain Smart while the ship was under construction. The first few pages reveal that this particular organization design was intended to be unique. SORM indicated that organizational relationships, job descriptions, responsibilities, and tasks of nearly all decision makers aboard *Starship* were to be programmed in ZOG, one software element of the ship's various computer systems. By referring to SORM, we can start to appreciate the intent, with respect to organization design, of putting ZOG aboard a ship.

SORM was the written specification of organization design. It was the hypothetical prescription of Captain Smart concerning how the organization should work. What is revealing here is the desire at the top of the organization, that is, the desire of Captain Smart, to create the predictability specified by the bureaucratic ideal type, an ideal where all direction comes from hierarchy, all action is guided by standardization, and all tasks are accounted for by specialists. The manifestation of this desire does not necessarily require the availability of computing technology. The availability of computing technology can, however, stimulate the executive's understandable desire for predictability. With the use of computing technology, information about what should be done, as well as information about what is being done, can be readily managed. It is in this sense that here, Captain Smart viewed computing technology as just another, albeit technically sophisticated, management control mechanism.

Captain Smart pointed out in SORM that the primary responsibility of the ship was to "maintain the highest state of readiness to perform assigned missions." SORM then went on to say that to maintain readiness, it is necessary to "understand the ship's compartments, systems, personnel organization, and task assignment; ...operate these systems; and ...evaluate the ship's performance when necessary."

As one reads on, what starts out as obvious and, perhaps, a bit mundane rapidly turns esoteric: "[T]he entire hierarchical net of the ship's responsibilities required to do this [that is, to get the job done] is presented electronically and is named the Ship's Planning, Plan Implementation, and Evaluation/Maintenance [PIE] Procedures for Performing Ship's Missions. PIE, according to SORM, "makes use of ZOG, a means to structure complex information accompanied by a simple menu selection interface to access the information. Information is displayed as a frame; each frame is a logical collection of facts. For example, a frame might describe a billet (job), an individual, or a task responsibility. To access the additional information, a person must select using the keyboard, the appropriate number or letter listed beside the information desired. . . . To accomplish a task in the net, its subtasks must first be read and performed in a depth-first manner; that is, the first subtask of the task must be followed to its bottom subtask. By accomplishing all the bottom subtasks in this manner, the top task is also accomplished." If there was a point in the subtask chain where a human decision was necessary, the computer program would so indicate.

The Technology Transfer Program

The ZOG system was part of what was known as the Technology Transfer Program, a program unique to *Starship* that was intended by Captain Smart to take available state-of-the-art computing technology from the corporate and academic world that would be useful to the management of *Starship*. In addition to ZOG, the program included the construction of a spatial data management system. This element would assist in the management of information and decision making in a combat situation. Another important element of the program was Airplan, an expert system designed to assist in the planning and management of air operations. Other portions of the program were less exotic: a computerized checklist for getting under way and arriving in port, an information and visual display system that assisted in the training and maintenance effort on the ship's weapons elevators, and a method of internal communication based on computerized office-to-office memos. The Technology Transfer Program was to integrate new technology with computing technology that was already in the scheme of things for *Starship*, such as the system that assisted logistics officers in keeping the ship supplied with everything from fresh milk to spare parts for the airplanes.

The installation of ZOG, as well as other advanced computer-based technology, was not part of the ship design originally envisioned by higher U.S. Navy headquarters. Rather, it was the result of the persistence of Captain Smart and his considerable skill at getting unusual and innovative projects approved. In this case, the management of the introduction of technology was just as innovative as the technology itself. The normal process of research, development, test, and evaluation (RDT&E) would have been a long, drawn-out bureaucratic ritual. Those who developed the original technology would be separated and, indeed, buffered from those who would use the technology by a massive RDT&E formal bureaucracy as well as by the informal vagaries of institutional politics. However, Captain Smart was intent on speeding up the process of transferring state-of-the-art technology into *Starship*. Moreover, he was determined to let those who use the technology perform the functions of development, test, and evaluation rather than the navy bureaucracy and field units that specifically existed for these purposes. Accordingly, he (as prospective commanding officer of the ship) went personally and directly to the source of the technology at Carnegie-Mellon University and elsewhere. He was able to convince civilian scientists who had been developing ZOG that the building and commissioning of *Starship* was a window of experimental opportunity for trying out nascent computing technology. He was able to convince navy officials to use the newly commissioned *Starship*, a ship that would be performing operational navy duties, as a laboratory for RDT&E.

When *Starship* put to sea on its first operational mission (an around-the-world cruise that would change her base of operations from the Atlantic to the Pacific Coast), ZOG had been installed but not yet programmed with functional knowledge. The scientists at Carnegie-Mellon had provided the framework for an expert system. The crew of *Starship* would provide the expertise. At the onset, ZOG was much like a library with empty shelves. The idea was that job incumbents would transfer their relevant knowledge to the computer program. Each of the major organizational participants was to teach ZOG what he knew. Then, as the participant learned more, this learning would be transferred to ZOG. Eventually, ZOG would become the guide for the behavior of subsequent incumbents. In this way, as jobs turned over, which they did every few years, a new participant would be at the same point on the learning curve as the departing person. Smart hoped that as time went on, and these programs were debugged and implemented, the ZOG program would eventually capture all

In this way, as jobs turned over… a new participant would be at the same point on the learning curve as the departing person.

aspects of the ship's tasks and become the recorded corporate memory and corporate expertise. ZOG could then be used for decision making or decision assistance. At least, this was the plan.

The Technology Transfer Program placed *Starship* at sea in the capacity of an R&D platform. Her immediate mission was to act as a seaborne carrier ready to conduct combat operations. However, no room and no assets existed in the overall scheme of carrier employment for the R&D role. Navy command officials ashore and, indeed, the formal structure of the navy RDT&E system made a clear distinction between the process of operations and the process of R&D. Captain Smart's concept for the implementation of the Technology Transfer Program, however, would blur the distinction between RDT&E on the one hand and operations on the other.

There were two genres of innovation at the heart of Captain Smart's Technology Transfer Program, and both of these were antiethical to the navy's way of thinking and doing business. The first type of innovation was an impingement on day-to-day operations and readiness by the process of RDT&E. According to traditional navy norms, once the navy went operational with a unit such as the USS *Starship,* operations would be the heart and soul of the functions to be performed. *Starship* would be an asset at sea, ready to take on any task in peace or war. *Starship* would be controlled by operational staffs and would use resources that were appropriated for operational use.

The second sort of innovation of the Technology Transfer Program that went against the traditional navy grain was the thought that human professional expertise might not be sufficient to achieve operational effectiveness. The idea of technology as such is not unattractive to the navy; the navy is, indeed, a "high-tech" organization. However, the navy has resisted technology that implies the superiority of machine over man. For example, in the years following the Civil War, for a time the navy rejected mechanical systems for propelling ships (as opposed to sails). It was feared that officers would lose control of shipboard operations if the people needed to operate propulsion plants were located below deck out of sight of the officers responsible for maneuvering a ship (Morison 1986). The navy also resisted using surface-to-surface antiship missiles because it was believed that manned aircraft were inherently a better system for sinking ships than unmanned missiles. The essential element of navy shipboard organization design and norms is the concept of individual, that is, human, accountability. The use of computing technology to restructure a shipboard organization design to place greater emphasis on computer-assisted control of the organization process and less reliance on human professional ability to use technical, operational, and managerial expertise is in this sense antithetic to navy tradition.

Both before and during her maiden voyage, *Starship* had to complete a litany of operational readiness inspections and operate as a normal carrier. The admirals ashore in Type Commander (aviation forces) and Fleet Headquarters (high command) were skeptical of Smart's ideas. They made no allowances in the operational schedule for the development of the innovative technical systems that had been installed in the ship. Personnel headquarters in Washington, D.C., had made no assignments to *Starship* based on her need for extra people to implement the Technology Transfer Program. Thus, a Management Department was formed by Smart to create an organizational home for the program. Enlisted men with technical skills who were assigned by personnel headquarters to other departments were internally transferred into the Management Department. A few officers who had experienced relevant graduate training were internally shifted to duties that would help Smart accomplish his RDT&E objectives.

The success of ZOG depended on the willingness and ability of officers to conceptualize their own skills and program these into the ZOG software framework. There was a predominately negative reaction to ZOG among the officers aboard *Starship*. Because most officers felt that they had neither the time nor the expertise to play with ZOG and also perform their operational functions, the process of filling in the frames of ZOG with the knowledge base of the incumbent officers never significantly progressed.

In the early stages of working up for the first cruise and during the early phases of the cruise itself, *Starship* was visibly less than effective in performing her required mission-related tasks. In the opinion of most of the department heads aboard *Starship*, the reasons for this ineffectiveness were (1) the failure of Smart to use more traditional ways of managing and leading the crew, (2) the fact that the Technology Transfer Program was not in all instances appropriate to decision-making requirements, (3) the lack of time to perform both the developmental functions of the Technology Transfer Program and operational tasks, and (4) the technical problems encountered in various aspects of the program.

By the time *Starship* reached the Indian Ocean halfway through the cruise, she was perceived by commanders ashore to be a ready ship fully capable of conducting combat operations. Captain Smart had shifted his priorities somewhere between the Atlantic Coast and the Indian Ocean from ZOG R&D (and the Technology Transfer Program in general) to short-term operational proficiency. The test-bed mission for the Technology Transfer Program was clearly now of lesser importance than near-term operational readiness. The atmosphere of operational professionalism and the requirements of the tasks at sea had started to invalidate, even before the end of Smart's tour as commanding officer, the Technology Transfer Program and its unique method of introducing innovation. It was at this point that Smart's tour of duty ended. He was relieved in the normal course of events by an officer whose style and outlook were much more conventional.

With the departure of Captain Smart, the driving force behind the Technology Transfer Program was no longer present. This program had been "pulled" by Smart but would not now be "pushed" from within the navy hierarchy because, as was noted, there was no enthusiasm in the command structure ashore for this project. By the time Captain Smart left *Starship,* the process of using the ship as an RDT&E test bed had a low priority. Most of the ship's officers considered the project a dead issue. Within the operating forces, operations (albeit peacetime operations) would continue to be the unambiguously highest priority.

Smart's legacy to his relief, Captain Good, was the hardware and software that had been installed in the early fruitful months of the project. Various elements of the Technology Transfer Program were available for use during Captain Good's tour of duty. These were decision-assistance and information-processing elements that could be used to help run shipboard operations concerning navigation and tactical control from the bridge, air operations in general and the control of air traffic in the vicinity of the ship in particular, control of procedures for getting the ship under way and arriving in port, and maintenance procedures with respect to the ship's weapons elevators and control of logistics (spare parts and so on).

An Assessment of the Technology Transfer Program

Smart's first objective, the synthesis of RDT&E with operations was not achieved. A look at some of the details shows that Smart's second objective, that is, placing greater emphasis on computer-assisted control of the organization process and less reliance on individual ability to use technical and managerial expertise, would also wither.

The element of technology transfer devoted to the conduct of flight operations aboard *Starship* was Airplan. It had the capacity to specify flight operations cycles, schedule events, and help make decisions during flight operations. The Airplan element is particularly interesting because flight operations are the heart and soul of carrier operations and because of the complexity that is inherent in such operations.

Elements of the Technology Transfer Program generated a planning schedule output based on the mission requirements identified by the ship's Operations Department. The aircraft squadron schedulers used this output to manually prepare specific aircrew and aircraft assignments. Once flight operations commenced, computer terminals located on the bridge and in Air Operations, the tower, and the Carrier Air Traffic Control Center (CATCC) could be used to perform various functions. Airplan was based on AI programs that were designed to be capable of solving some of the complex problems involving the planning and implementation of flight operations. The expert system was programmed to provide information concerning aircraft fuel remaining and projected landing time, gather statistics on air operations, and alert relevant personnel for the necessity of a human decision. The system could help make decisions concerning the necessity to refuel aircraft from tanker planes while they were airborne and predict the effects on planned landing and takeoff times in the cycle of various events, for example, unexpectedly landing an aircraft with mechanical difficulty or launching interceptor aircraft to investigate airborne intruders.

...decision makers that were responsible for the effetiveness of the air control function... preferred to do their job without computer decision assistance.

Night flight operations were one of the more demanding functions performed by *Starship*. The observation of these operations further questions the use of computing systems. *Starship's* pilots were required to land at night aboard the ship at least every seven days to stay qualified. Thus, the ship scheduled night flying at every available opportunity. When they operated in the vicinity of the ship at night, all aircraft were controlled by radar air traffic controllers. These controllers worked in CATCC. They gave directions and glidepath information to pilots over the radio until the landing aircraft were in the close vicinity of the ship. At this point, the pilot took over, visually guided by a beam of light projected from a mirror on the deck. The objective was to catch the tailhook of the aircraft on one of four crossdeck wires.

In CATCC, air traffic control information was displayed by writing backwards on the back of lighted plexiglass boards with a grease pencil. A great deal of information had to be displayed to be immediately available to the CATCC officer who was responsible for controlling traffic in the vicinity of the ship: identification of each airborne plane, pilot's name, fuel remaining in each airplane, number of landings completed, weather condition, status and locations of airfields ashore to which a plane could be diverted (Bingo airfields), status of tanker (air-to-air refueling) aircraft that circled above the landing pattern ready to provide fuel as necessary.

The AI and information-processing element of the Airplan program could handle and display much of this information. However, the system was not used, nor was any other computer system that was available. The decision makers that were responsible for the effectiveness of the air control function in CATCC preferred to do their job without computer decision assistance. The expressed reasons for this preference were (1) a fear that massive system failure might suddenly transfer the total control of airborne aircraft to the human decision process and that such an event would place traffic control personnel in a situation that they could not handle; (2) a feeling that the computer system, even operating at peak performance, was not technically capable of providing decisions or assistance in making decisions that would increase effectiveness; and (3) a total personal involvement with the task being performed that resulted in a sense of responsibility that mitigated against placing even part of the decision process out of the hands of human practitioners. In short, for reasons that were both technical and emotional, expert practitioners did not believe that the Airplan expert system could increase effectiveness. Although there is little doubt that some computing functions, such as the display of information, could have increased effectiveness, it was the sense of responsibility felt by all CATCC personnel that made them want to keep direct control in human hands. These reactions were certainly in accordance with traditional navy notions concerning the singular importance of the (human) professional as well as norms concerning the accountability of personnel.

The level of effectiveness of CATCC, as measured by the rate of the aircraft landing in day and night conditions, was high. Accordingly, there was no incentive to take risks by changing the organizational process. Moreover, peak effectiveness periodically required the disregard of standard procedures. By disregarding the rules from time to time concerning minimum aircraft separation, controllers could shorten the time between landings, thus demonstrating an increase in effectiveness. CATCC personnel were willing to take responsibility for cheating on the rules to do a better job. The decrease in minimum separation of aircraft called for by standard procedures had become part of the normal CATCC routine. It was felt that there was little room in the air control function for preformed decisions that might be prescribed by the Airplan expert system.

Commander Boss, head of the Air Department, was responsible for the actual landing, takeoff, and deck handling of aircraft aboard the ship's flight deck. Boss's attitude concerning the Technology Transfer Program was reflected in this statement: "The computers are not very useful to me. They cause people to create even more paperwork. Oh, they are all right for helping with enlisted evaluation reports, officer fitness reports, and aspects of personnel and administration functions that we have to perform in the department. But they're really only good for that sort of word processing stuff. They do not help me run the show or make decisions."

Although the tower, which was Boss's perch for observing and controlling the operation of his flight deck, was equipped to take advantage of computerized information and decision-assistance tools, they were not used. Here, as in the case of CATCC using the grease pencil, plexiglass and the judgment of human experts reigned supreme. The observed rejection of the use of computerized expert systems in CATCC and the tower was replicated with the control of the tactical or strategic situation of the ship from the bridge. Here, the spatial data management system had been installed to give the commanding officer a real-time display of geography, intelligence, disposition of forces, and so on. However, Captain Good, although able to use this equipment without leaving his seat on the bridge, never turned it on and indicated that he would never use it. Captain Good also rejected the use of electronic memos that would allow him to communicate with his subordinates. He preferred face-to-face confrontation.

The Decision to Use Human Expertise

The success of the Technology Transfer Program in general—and the ZOG element in particular—rested on the following assumptions: (1) Centralized decision making was more effective than decentralized decision making. The program of knowledge that was the ultimate objective of ZOG and the use of this knowledge at all levels of a hierarchy are the analogue of a knowledgeable superior explicitly directing or advising subordinates. (2) The nature of human expertise—in particular, the highly developed expertise that can cope with extreme complexity—is analytic rather than intuitive. (3) Human beings can translate their expert knowledge into computer language. (4) Computing technology is capable of accepting and using the highest levels of human expertise.

The decision processes observed, for example, in CATCC, demonstrate that these assumptions were either incorrect in an objective sense or were perceived to be incorrect by practicing experts. Furthermore, observations in CATCC demonstrate that at least in one respect, this organization varied from the ideal Weberian bureaucracy, the implicit model for the Technology Transfer Program. A Weberian bureaucracy relies on standard decisions. As they controlled the flow of airborne traffic, experts in CATCC increased the effectiveness of the organization by practicing the art of nonstandard decisions.

Weberian bureaucracy also calls for hierarchical control. To the contrary, the organization process aboard *Starship* relied, particularly during complex and demanding contingencies, on the devolution of decisions to human experts who used their own intuition rather than analysis. With respect to the use of human expertise in complex contingencies, the decision process in the event of an airborne emergency is typical of what went on aboard *Starship*.

To illustrate, one night a pilot reported that he had the cockpit indication of a landing gear malfunction. Either one of his landing wheels would not drop to the position for landing, or he had a false indication of this problem. The options were to let him try to land aboard ship or divert him to an airfield ashore. The fuel situation was the critical factor. When he heard the pilot radio the ship of the problem, in the tower, Commander Boss reached up and pushed a button that resulted in the display of a computer-recommended airfield to which to divert the plane along with appropriate course, distance, and fuel computation.

Before Commander Boss could reflect on this computer output or pass it along, the word from CATCC came through on the intercommunication squawk box: "We're going to Bingo [divert] him to Kadena airfield." CAG and the commander of the aircraft squadron of the troubled airplane, both physically located in CATCC, had conferred. All necessary calculations were made in their heads. They knew that Kadena was the closest field because they could read this fact on one of the plexiglass boards. They also knew it was important not to delay the decision by analyzing the nature of the landing gear malfunction indication.

The decision-making factors concerning the probability of success if the aircraft were sent ashore were analytic. However, those factors that influenced the trade-offs between the risk of sending the aircraft ashore or the risk of attempting to land it aboard the ship were neither analytic nor clear. The crux of the decision at hand was not "Can we send the aircraft ashore?" It was "Would it be preferable to send the aircraft ashore with or without an attempt to conduct airborne refueling, or would it be preferable to attempt a shipboard landing?" In analytic terms, the relative merits of these choices were not clear. In an analytic context, therefore, it might have been desirable to take time to obtain more information concerning the nature of the landing wheel malfunction indication. Was it real? Was it false?

In the final analysis, CAG knew that the important element of an effective decision in this situation was not the decision itself but rather the process of making a decision and communicating it to the pilot in trouble. Most of all, CAG feared vacillation in the decision process. He could remember losing planes because of vacillation in similar situations throughout his years as a naval aviator. His intuitive capacity was a function of the large number of unique cases that he had experienced. What CAG recognized was a pattern of stimulus that caused him to react as much out of intuition as analytic calculation. Rather than conforming to a pattern of analy-

…with other department heads… the same sort of decision process and aversion to cutting the human expert from the decision loop was evident.

sis that he might have learned in training or on the job, one that could be applied to a particular class of problems by a human or machine expert, he was using his intuition to treat this contingency as another unique case. CAG had learned to recognize the need to avoid vacillation in this sort of contingency in much the same way that a baby learns to recognize its mother's face, not by recognizing the parts of the whole (nose or mouth with the baby; type of aircraft or fuel on board with the situation at hand) and not by analyzing the relationships among parts of the whole (Dreyfus and Dreyfus 1986). CAG's expertise was the capacity to recognize a situation in terms of an integrated whole, fit the situation into a relevant category, and act to satisfy the requirements of a particular type of contingency.

We do not know whether a computer would have recognized the same pattern and reacted the same way. It is clear, however, that officers aboard the ship did not think so. Neither do we know how CAG could have told ZOG or the Airplan expert system why it was so important not to vacillate in this particular type of circumstance. In some types of emergencies, it might be a good idea to vacillate. In others, it might not. It is also not clear whether some future CAG, when confronted with the choice of using the expertise programmed into a computing system by the current CAG, would have been willing to take such machine-generated expert advice. Who would be accountable for decisions in such a scenario?

An interview with CAG concerning this decision-making scenario sheds further light on the nature of his expertise and its contribution to the effectiveness of the organization.

CAG: Even though I do not know all the nitty-gritty about all the [aircraft technical] systems, I have a pretty good gut feel. All aircraft are basically the same when it comes down to it . . . so I can more or less gather the data a lot quicker than other people that do not have as much experience as me. Then I can tell the Captain. … Last night, even though we had the fuel airborne [in a tanker], you know that you only have 4000 pounds to give to the guy, and you know he is going to take 1500 pounds just to join up on the tanker. And then I had to look at the divert field… 120 miles. I computed the gas he needed… 5600 pounds to get to the divert field. So I said, 'Well, he's got 6000 pounds.' I said that we need to make a decision here in about two minutes. So that's what I did… just came up and made the decision. 'Captain, I recommend that we send him to the beach because of this and that.' He said, 'OK, fine'. A lot of people will get themselves into trouble because what they will do is say well… we got to do this. … OK, let's do that. And then somebody else will say let's do this. So they will be going back and forth. Send the guy to the beach. So the guy will start heading for the beach. Then, oh no, bring him back here. So the guy will turn around and come back. I can't tell you how many airplanes they have lost. Some of them have gone right into the water because they are just jerking the pilot around, back and forth, and nobody was able to make a decision. You've got to make the decision. You've got to know your procedures and stick with it. Nine times out of ten it's going to be the right one.

Interviewer: Last night, did you use any of the computer systems aboard to help you make decisions?

CAG: None whatsoever! All you do is figure how much gas it takes him to get to 120 miles away. Then you make sure he's got it. And then you launch him. The only thing that was maybe computer assisted… I wanted to know the winds aloft to make sure that he did not have like 100-knot headwinds where it's really going to get him. And also I checked on the weather to give the guy. You know the pilot is extremely concerned in this circumstance. That's why the guy on the deck should be able to give him all the information that he needs to hear, to calm him down. … Make sure he has the big picture.

The Opinions of Other Department Heads

In discussions with other department heads aboard *Starship,* the same sort of decision process and aversion to cutting the human expert from the decision loop was evident. When asked how he felt about the use of computers aboard ship, the reactors officer (department head in charge of nuclear reactors) said, "When I first got here, I used ZOG to help me run the department, to plan and control what goes on. Now, as you can see, I do not even have a PERQ [hardware]. The thing is worthless because it's so user unfriendly. In concept, it was useful. In practice, it was not useful. … We need to be able to handle things here in the Reactor Department on a real-time basis. If we were to use the computers to help us operate the plant, we would always be behind. We can't rely on the thing being properly programmed or on the correct data being put into data banks. A major difference between commercial and naval nuclear plants is that they use computers to help with both normal and emergency procedures, and we do not. Three Mile Island is a good example of decisions being left to computers. In that accident, the computer was 30 minutes behind the problem."

The officer in charge of nuclear reactors explains that his task environment is as dynamic as those of the air controllers in CATCC or as CAG. In all these tasks, the rejection of the use of computers is significant for more than just technical reasons. The preprogrammed routines of decision or decision assistance in expert computing systems sets the organization in a static design. The ZOG concept itself specified a fixed structure for the organization in that all missions and tasks were to be broken into a hierarchy of subtasks to be performed by specific individuals occupying specific places in the structure. However, the working environment of *Starship* is sufficiently dynamic to require an almost infinite variety of designs. The locus of decision making, that is, the flow of information and the designation of a decision maker, must vary from situation to situation for this complex organization to be effective.

Sometimes, the reactors officer has nearly complete authority to make decisions with respect to his area of responsibility, and the captain of the ship has little or no information regarding reactor status. At other times, however, it would be necessary for the reactors officer to coordinate decision making with others with respect to reactors and to let them make decisions that affect reactor maintenance and operations. With the ship in a tactical combat situation, for example, the decision to operate reactors at reduced capacity for maintenance or safety reasons would necessarily involve the judgment of the ship's captain and the admiral who commands the combat task force. With the ship in close proximity to dense population centers, reaction to a nuclear power plant problem would result in a more conservative assessment of risks than in combat or during peacetime exercises at sea.

One of the primary command functions aboard the ship is the determination of the design of locus of decisions and information flow based on the substance of the decision and the contingency at hand. The captain's leadership effectiveness is not determined by the capacity to make decisions. Rather, it is determined by the capacity to decide how a decision will be made in a given circumstance. The function of command is to orchestrate a decision process rather than operate a decision process.

The professionals aboard *Starship* were not rejecting computing technology as such. They are generally fond of advanced technology. What they were rejecting was the static organization design suggested by the bureaucratic ideal type and specified by the Technology Transfer Program. They did so because this design—indeed, any static notion of how an organization should run—would be dysfunctional to the nature of carrier operations. For the computer to capture such a dynamic design capacity, individual experts would have to program the computer with their capacity to recognize an almost infinite variety of contingencies, categorize specific contingencies in a useful way, and select design prescriptions that would be efficacious in a particular contingency category. This was essentially the way that Captain Good and his senior subordinates ran the ship.

It is important to understand that these officers, although good at their jobs, were not able to articulate the manner in which they were constantly modifying the design of the organization, particularly the locus of decision making, authority, responsibility, and accountability. Neither could they articulate their intuitive feel for the specific variations of design that were employed aboard *Starship,* variations on which the effectiveness of the organization depended. The articulation of these various modes of decision making, however, is what the original ZOG project required.

Logistics is one area that we might think would be friendly to the use of computing.

The professionals aboard Starship were not rejecting computing technology as such. …they were rejecting… the static organization design…

The problem of keeping the ship supplied with everything from beans to fuel for the aircraft to sophisticated spare parts for aircraft and the ship seems at first glance to be a matter of information control, data processing, and communication. However, even here, the work environment was perceived to be more dynamic than could be handled by computing systems. The department head in charge of logistics commented that "we have the database to track all our parts moving through AIMD [maintenance] and supply. When the system goes down, we come to all stop. Even when the system is up, when you ask it the question, 'Don't we need more of these items because I see a higher demand?' The answer is, 'The allowance is two'. Then I say [to the computer], 'I didn't ask you what the allowance is. Should we have more of these because we are using more?' And so that second class aviation storekeeper [enlisted supply technician] would say, 'Hey, I can tell you anything that you want to know'. 'How many did we use on our last deployment?' I know that—three'. I say, 'That's irrelevant data. We used six so far on this deployment. Do I need more for the rest of my deployment?' But now the mind has abrogated all that stuff to the computer."

The supply officer continued by explaining that computers are rejected as a method of controlling the operations of his department: "The goal is to have 70 percent of your aircraft mission capable. You can fly them each day. That now has been generally exceeded by most carriers. Mission capability can be degraded by either the need for parts or the need for maintenance. That is why it is necessary to have a very close relationship between supply and maintenance. …Each morning we have an 8 o'clock meeting with the CAG and the head of AIMD [aircraft maintenance]. We review each squadron's readiness statistics for that day. There is a daily message that goes out to the world from CNO [Chief of Naval Operations] on down which says *Starship* is mission capable to a certain extent; this is the number of outstanding requisitions; this is the number of planes down for supply; here is a list of all the areas which the carrier feels are critical, requirements we have to get filled in the next couple of days. We are a deployed carrier, and we should be getting first priority on the parts. … I make just about all the decisions down here. When you first come on board the carrier and join the organization, you have to find out what his [the commanding officer's] management style is and what he is emphasizing. I have a huge database down here, and I can give him the whole piece of the pie, a sliver of it, or whatever. … I found on the aviation side the interest is in what are the trends and how quickly you are going to get well when you have a problem. We have nine squadrons, and if all nine went down today for different reasons, [Captain Good would say], 'How soon are you going to get the material? When do you see it reaching us by COD [airborne delivery to the ship]?' Initially, my thought process was, 'Here is the status on the requisitions. Then if you had 400 requisitions outstanding, you'd have to find out how many have been shipped. Then I would mentally extrapolate. If I had 130 that have shipping status, maybe 50 are key items. Then how many of the 50 will arrive on board?' He wants that measured into a nice neat ball, to say, 'Captain, we will be able to get 10 more requisitions for 10 more planes'. What I am doing here is setting the goals for myself. Your credibility in the world of logistics is built on how close to giving good data you are because obviously his report card from the admiral is based on the mission-capable figure. That is what your warfare commander wants to know about. He doesn't want to know how things are now as much as he wants to know how things are going to be tomorrow. I am not really thinking about what the captain wants. I'm thinking about what the admiral wants from him."

CAG, the supply officer, and the reactors officer were confronted with different types of tasks. All three rejected the standardized and centralized type of decision assistance that computers can provide. The reactors officer prefers to rely on human judgment because it can be used to implement procedure on a real-time basis and because he can rely on the programming of his people to a greater extent than he can rely on the efficacy of a computer program. The way that CAG calculated the variables necessary to ensure the safe arrival of his aircraft and the way that the supply officer calculates the number of requisitions expected to arrive is similar. They both use their own experience to develop rules of thumb that can be applied to get a solution that is good enough rather than a solution that is "perfect." However, these rules of thumb are by no means standardized. These officers knew that "perfect" can be the enemy of "good enough." In the case of CAG, if he had procrastinated long enough to get a perfect answer, the plane undoubtedly would have run out of fuel. The supply officer knew that if he had communicated to the commanding officer that the question about arrival times of the requisitioned items could not be answered with perfect confidence, he would lose the confidence of Captain Good. If he loses this confidence, he might lose the opportunity to use his own discretion in most decision-making situations.

These experts did not operate using a process of analysis. Even their rules of thumb were not readily subjected to a process of analysis. If they always relied on an analytic decision process (if, for example, the supply officer took the analytic advice of his enlisted storekeeper), he would never rise above the level of mere competence (Dreyfus and Dreyfus 1986), that is, above the level of analytic capacity. Similarly, if these officers relied on the mechanical expert systems in all circumstances, their effectiveness would be limited by the inability of the computer systems of the Technology Transfer Program to make the transition from analysis to pattern recognition, that is, intuition.

Conclusion

It is clear that Captain Smart's attempt to use computing technology to improve the effectiveness of *Starship* failed. Some of the reasons for this failure have to do with his attempt to modify the procedures for using RDT&E to effect innovation. Some of the reasons for failure have to do with the inappropriateness, real or perceived, of the technology as such to the organization and its functions. The description and analysis of this case leads us to the following conclusions concerning the reasons for failure: First, the original sponsor, Captain Smart, did not stay with the project. The driving force of the project was the intellect and the leadership of one man. In this sense, Smart's role is analogous to the role of Admiral Rickover with respect to the use of nuclear-powered ships. One difference is that Rickover was able to stay with his project, and Captain Smart was not. Similarly, the developers of the project, civilian scientists and engineers as well as the original members of Captain Smart's Management Department, were not continuing members of the project team. Second, there was no clear distinction made in the design of the project between decision automation and decision assistance. Third, the operational staff, that is, the crew of the ship, was not given an incentive to develop the project. Because operational requirements were of the highest priority, the operational accomplishments of the crew would be the accomplishments that were rewarded. Fourth, the project strategy was to implement a large and comprehensive technical system, which obviously required great organizational change. No attempt was made to effect such change in an incremental way so that the organization could adapt to the use of the technology. Thus, when the project and its objectives came into conflict with organizational norms and the values of key actors, the project was put aside. Fifth, no attempt was made to demonstrate the technical capacity of the project to accomplish the desired objectives before the start of testing and evaluation. Thus, there was no way to predict the extent to which testing and evaluation might adversely affect operations.

For those who would attempt to develop the state of the art of AI to more closely approach optimum organizational effectiveness, the description and analysis of the case of the USS *Starship* suggest the following questions: (1) How can expert machine systems distinguish between contingencies where it is likely that analytic decision processes will be sufficient from contingencies that require more intuition or pattern recognition than technology can deliver? (2) How can expert machine systems determine when it is functional to cheat on preprogrammed decision criteria and rules? (3) How can organization members teach expert machine systems what they have learned by way of pattern recognition? (4) How can human operators perform operational functions and, at the same time, maintain the currency of machine expert system programs so that machine expert decisions are appropriate to a real-time environment? (5) How can organizations that use expert machine systems deal with the problem of accountability? Is a human being that relies on machine assistance accountable if the machine fails? Is a human being to be rewarded in some way if the machine succeeds? (6) How can expert machine systems whose focus is necessarily on a narrow range of expertise coordinate a wide range of expert activities that are required to bring about organizational effectiveness? (7) How can organization members be encouraged to accept the functional aspects of machine expert systems when the organization relies on value-oriented, as well as factual, decision criteria?

References

Dreyfus, H. L., and Stuart, E. 1986. *Mind over Machines: The Power of Human Expertise in the Era of the Computer.* New York: Free Press.

Morison, E. E. 1966. *Men, Machines, and Modern Times*. Cambridge, Mass.: MIT Press.

Perrow, C. 1984. *Normal Accidents*. New York: Basic.

Simon, H. A. 1957. *Administrative Behavior.* New York: Harper.

Taylor, F. 1911. *The Principles of Scientific Management.* New York: Norton.

Stephen Sloane was commissioned an officer in the United States Navy when he graduated from the U.S. Naval Academy in 1957. On completion of flight training and designation as a naval aviator, he served on active duty, specializing in airborne antisubmarine warfare operations as well as the study of strategy and administration. During his career, he served as a patrol plane commander, aircraft squadron commanding officer, aircraft carrier department head, and aircraft carrier officer of the deck and command duty officer and as well as a qualified carrier pilot. He is a graduate of the Harvard University John F. Kennedy Graduate School of Public Administration (M.P.A.) and the University of California at Berkeley (M.A. and Ph.D. in political science). He retired from active duty in 1983 at the rank of captain. He is currently a faculty member of the Government Department at Saint Mary's College of California.