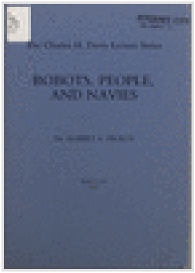

Robots, People, and Navies (1983)

Pages
41

Size
7 x 10

ISBN
0309310229

Frosch, Robert A.; The Naval Studies Board; National Research Council

The Charles H. Davis Lecture Series

Sixth Lecture

ROBOTS, PEOPLE, AND NAVIES

by
Dr. Robert A. Frosch
Vice President, General Motors Corporation
General Motors Research Laboratories

Presented Before the Students and Faculty
of the
Naval Postgraduate School
July 27, 1982

and

The Naval War College
September 7, 1982

Organized by

The Naval Studies Board
of the
National Research Council

Sponsored by

The Office of Naval Research

NATIONAL ACADEMY PRESS
Washington, D.C. 1982

CONTENTS

THE CHARLES H. DAVIS LECTURE SERIES

AT THE CLOSE of that greatest of all contests of men and machines, World War II, Theodore von Karman could say, with deep personal conviction, that ". . . scientific results cannot be used efficiently by soldiers and sailors who have no understanding of them, and scientists cannot produce results useful for warfare without an understanding of the operations." With such simple truths fresh on their minds, von Karman and his civilian and military colleagues proceeded to forge institutional links—such as the Office of Naval Research—through which they hoped to encourage an enduring partnership between the scientific and military communities. Though the intensity of the bond has fluctuated with the ebb and flow of international relations and internal affairs, the partnership has endured to produce a military capability but dimly perceived by those who established it. But the partnership is not self-sustaining; it requires the constant vigilance of those who have not forgotten the bitter lessons of the past, the outspoken dedication of those whose vision extends beyond the next procurement cycle, and, above all, it requires open communication between the partners. It is to this latter task that the Charles H. Davis Lecture Series is dedicated.

The lecture series is named in honor of Rear Admiral Charles Henry Davis (1807–1877) whose distinguished career as a naval officer and as a scientist so epitomizes the objectives of the series, and whose clear vision of the proper role of science in human affairs redounded to the betterment of all men. The topics and the speakers in the series are chosen by a Search Committee operating under the National Research Council of the National Academy of Sciences, and two lectures are presented each year before the students and faculty of both the Naval Postgraduate School in Monterey, California, and The Naval War College at Newport, Rhode Island. The series is sponsored by the Office of Naval Research.

REAR ADMIRAL
CHARLES H. DAVIS
(1807–1877)

CHARLES HENRY DAVIS was born January 16, 1807, in Boston, Massachusetts. His education consisted of preparation at the Boston Latin School followed by two years at Harvard University (1821–1823). In 1823, Davis was appointed midshipman and sailed (1824) on the UNITED STATES to the West Coast of South America where he transferred to the DOLPHIN for a cruise of the Pacific. Returning to Harvard he continued to work on a degree in mathematics and is listed with the graduating class of 1825.

In 1829 Davis became passed midshipman and was ordered to the

ONTARIO (1829–1832) of the Mediterranean squadron. Later, while serving aboard the VINCENNES (1833–1835), he was promoted to lieutenant. Aboard the INDEPENDENCE (1837–1841) Davis made a cruise to Russia and then to Brazil. Throughout these early years at sea Davis continued to study mathematics, astronomy and hydrology. During this period one of his superiors would write of him, "C. H. Davis is devoted to the improvement of his mind; and his country may expect much from him."

From 1842 to 1856 Davis undertook a number of special tasks and served on several commissions and boards. Notable among these was his participation in a survey of the New England coastal waters (1846–1849) during which he discovered several shoals that may have been responsible for a number of unexplained wrecks in the area. It was during this period in his career that Davis published "A Memoir upon the Geological Action of the Tidal and Other Currents of the Ocean" (1849) and "The Law of Deposit of the Flood Tide" (1852). He was also a prime mover in establishing the "America Ephemeris and Nautical Almanac" (1849) and supervising its publication at Cambridge, Massachusetts until 1855 and again from 1859 to 1862.

Promoted to commander in 1854, Davis resumed sea duty in command of the ST. MARYS in the Pacific (1856–1859). While he was captain of the ST. MARYS he was instrumental in securing the release of the adventurer William Walker and his followers who were beseiged at Rivas, Nicaragua.

With the outbreak of the Civil War Davis was immediately appointed to a number of important positions. He became the executive head of the new Bureau of Detail for selecting and assigning officers. He was one of three officers appointed by Secretary Gideon Welles to the Ironclad Board which passed judgment on the plans and specifications for the MONITOR and other ironclads. Promoted to captain in November 1861, Davis participated in the development of plans for blockading the Atlantic Coast, planning the operation against Hatteras Inlet and Port Royal Channel, and the early naval strategy of the war.

During the operations against Port Royal, Davis served as captain of the fleet and Chief of Staff to Admiral Samuel F. Du Pont. He shares with Du Pont a great deal of the credit for the excellent plan of attack carried out on November 7, 1861. Later, as flag officer of the Mississippi Flotilla, Davis led successful engagements against the Confederate fleet which contributed to the abandonment of Fort Pillow and the surrender of Memphis. He was promoted to commodore in July 1862, and to rear admiral on February 7, 1863.

In late 1862 Davis returned to Washington to head the newly established Bureau of Navigation. From this position he worked closely with such distinguished scientists as Joseph Henry and Alexander Bache to establish a "Permanent Commission" to advise the government on inventions and other scientific proposals which were being stimulated by the war. The Permanent Commission was established by the Secretary of the Navy on February 11, 1863 with Davis, Bache and Henry as members. However, Davis and his colleagues saw a wider need for cooperation between science and government and worked diligently for the establishment of the National Academy of Sciences. Their efforts were successful; President Abraham Lincoln signed a bill authorizing the establishment of the Academy on March 3, 1863.

In 1865, Admiral Davis was appointed superintendent of the Naval Observatory in Washington. In 1867 he returned to sea in command of the South Atlantic Squadron. Back in Washington in 1869 he was made a member of the Lighthouse Board and commander of the Norfolk Navy Yard. He later resumed his post as superintendent of the Naval Observatory where he served until his death on February 18, 1877.

DEDICATION

by

Dr. Robert A. Frosch

to

Dr. William B. McLean

I dedicate this Charles H. Davis Lecture to the memory of the late Bill McLean, inventor of the Sidewinder missile, technical director of the Naval Ordnance Test Station, now the Naval Weapons Center, China Lake, California, and later of the Naval Undersea Center, San Diego, California. Bill was an inspired technologist and innovator, one of the earliest advocates of, and experimenters with, the human extension ideas that I will discuss today.

Dr. WILLIAM B. McLEAN

(1914–1976)

On June 29, 1974, after 33 years of distinguished federal service, Dr. William B. McLean retired from his position as technical director of the Naval Undersea Center in San Diego, California. At his retirement ceremony, the assistant secretary of the Navy (R&D), Dr. David S. Potter, referred to Dr. McLean as "the greatest scientist of our decade in the Civil Service." Dr. McLean's remarkably successful career was largely due to three talents rarely found in a single individual: he was a creative and innovative scientist-engineer who always preferred the simple rather than the complex solution; his interest, his dedication, and his persistence spanned the range from idea to operational hardware; and he was able to pursue his goals efficiently and effectively within the federal bureaucracy.

William Burdette McLean, the son of the Reverend and Mrs. Robert N. McLean, was born in Portland, Oregon, on May 21, 1914. His early interest in science led to enrollment in the California Institute of Technology, where he received his B.S. (1935), M.S. (1937), and Ph.D. (1939). He then spent two years as a postdoctoral fellow at the University of Iowa, where he studied nuclear physics. Completing his studies in 1941, Dr. McLean accepted a position as research physicist at the Bureau of Standards in Washington, D.C., where he worked on the design and production of proximity fuses.

In 1945, Dr. McLean began his long and productive association with the U.S. Navy. He accepted a position with the Naval Ordnance Test Station (now Naval Weapons Center) at China Lake, California, where he devoted himself to research and development in the field of ordnance. He became head of the Aviation Ordnance Department in 1950 and technical director of the laboratory in 1954.

With the reorganization of the Navy laboratories in 1966, a new laboratory, the Naval Undersea Warfare Center (now Naval Undersea Center) was established and ultimately located at San Diego, California. Dr. McLean served as technical director of the new laboratory until his retirement.

During his years at the Naval Ordnance Test Station, Dr. McLean is credited with the development of many important Navy weapons, such as Sidewinder, ASROC, and the torpedo Mark 46. Upon his transfer to the Naval Undersea Center he was able to more fully indulge his interest in undersea work vehicles. To the two submersibles, MORAY and DEEP JEEP, which resulted from work at China Lake, were added CURV, RUWS, MAKAKAI, DEEP VIEW, and NEMO. Here also, his interests in marine mammals led to a research program that successfully demonstrated the ability of such mammals to perform search and recovery tasks.

Dr. McLean's many awards for his achievements included the President's Award for Distinguished Federal Civilian Service (1958), the American Ordnance Association's Blandy Gold Medal (1960), the Rockefeller Public Service Award for Science, Technology and Engineering (1965), the California Institute of Technology Alumni Distinguished Service Award (1969), and the Institute of Electrical and Electronic Engineers' Harry Diamond Award (1972). He was further recognized by elections to membership in the National Academy of Engineering (1965) and the National Academy of Sciences (1973).

DR. ROBERT A. FROSCH

ROBOTS, PEOPLE, AND NAVIES

DR. ROBERT A. FROSCH
Vice President, General Motors Corporation
General Motors Research Laboratories

I approach my topic from the point of view of engineering design as a form of imaginative artistry. I do not aim at forecasting and predicting, or even insist on a particular direction, but rather will try to provide materials and ideas concerning the way a particular line of future technology may affect the way in which we build and operate future navies. It will be clear that navies are only a special case for the possibilities and problems I will raise. I will sketch ideas, give impressions, and provide materials for future elaboration. I intend this to be a stimulus to the thought of an audience of highly educated insiders.

Approaching the problem from that point of view, I have little concern with questions like "why is there a navy, which war must we prepare for, or what future scenarios will the systems fit into?" I am talking about technology that is available for the solution of future problems, and particularly about technological ideas that are intended to have sufficient flexibility to be of interest for a large class of future problems.

Over the years, I have been afflicted by budget types, and other forms of bureaucrats who seem to think that there is something wrong with a technology that is looking for problems to solve. As I read the history of human problems and their interaction with technology, I find this approach legitimate and more likely to occur than the supposedly proper situation: problems looking for technology. Problems are not completely defined separately from their solutions; the solutions and the problems are always redefined together. Frequently, it is the possibility of a solution that makes it possible to refine the meaning of a general problem such that the meaning of the solution itself can be defined.

I will discuss three related topics. The first (and longest) will cover what I will call for convenience "The Robot Revolution" and will describe some of the possibilities that I see in technology. The second will treat some ideas about the areas of application of these technological ideas to navies. Finally, I will discuss some problems to be faced in using the technologies for the purposes described.

THE ROBOT REVOLUTION

It is clear that one of the major technological changes in the past 15 years has been the explosive growth of capability in both computer hardware and computer software. The cost, size, and weight of unit computing power can best be described as having collapsed, the factors of decrease in the past few decades being orders of magnitude per decade. At the same time, our ability to handle the complexity produced by larger and larger aggregations of computer capability seems to have generally kept pace with the increase in complexity and capability. We are able to build computers that are physically smaller and smaller, but intellectually larger and larger. To date, we have managed to keep pace with the need for ever more complex problems, so that we use all of the capability that we succeed in producing.

This increase in complexity has forced us to attempt a deeper understanding of the modes by which we ourselves think and by which we control ourselves and our tools. In order to understand how to program a computer to carry out computational processes that are at least akin to thought, and apparently identical to thought, in a way that is efficient and adaptive, it has been necessary for us to analyze the processes by which we think and by which we adapt our thoughts and actions to incoming data to a depth that had not been achieved before we had thinking machinery to program. We have learned to break analytical and intentional thoughts into long chains of "if, then" statements with multiple connectivity, not necessarily linear thought trees, but networks with feedback statements in them and very complex connectivity, in order to produce programs that carry out our computational intentions.

A second important line of development has been the ability to marry these complex computational systems with sensors—measurements of temperature, stress, flow, etc.—and with machinery to act as effectors. We can now construct control systems, for example, for the operation of turbine engines and internal combustion engines, which, if not intelligent, can at least be called clever.

I try to eschew the word "intelligent," not because I am fearful of the discussion of whether machines are intelligent, but because the idea of intelligence even in people has become such a problem in the social sciences that it is better, I think, to avoid the term. I cannot always succeed in avoiding it, and I will return to the question of its definition for a moment later.

Our ability to build clever controllers has meant that we can give a machine a rather complex set of modes of sensation and thought that enable it to carry out a set of contingent instructions, that is, instructions that tell it that if it senses certain things, it should carry out further actions and computations, and may await other "sensations," which will signal it to carry on a further train of actions. This seeking and finding of sensations followed by other actions is all preprogrammed in a contingent way, but with the details of the order and occurrence of the contingencies not specified. Thus the machine adapts a general set of instructions to a specific set of circumstances.

When we use computers for computation, or for word processing or business purposes, we are very much aware of the machine and of our interface with it: the keyboard is under our fingers, the screen is before our eyes, the light pen may be in our hand. The brain and memory of the machine may be elsewhere, but our connection with it is clear to us. On the other hand, when a computer is used as an action controller for a machine, in most cases the computer is effectively transparent. That is, we drive the car or fly the aircraft as we used to before electronic computers, quite unaware that the chip is readjusting the engine operation continually, in fact doing the real control job; our driving or flying controls are a means of informing the computer what we want the craft to do. The computer operates the controls to carry out our intentions.

This program of fly-by-wire has been used extensively in the space shuttle, and in some military aircraft, in which manipulation of the stick is interpreted by the computer in terms of a set of flight-control equations. The computer uses the pilot's control instructions together with a body of flight data from sensors in and on the craft, and, in some circumstances and modes, with data from the ground, to work out, through the control equations, what instructions to use to control reaction jets and/or aerodynamic control surfaces. The pilot does not directly control anything; he informs the computer what he wants the craft to do. The operations take place as operations responsive to the pilot, but the beauty of them is that the operator may be quite unaware of what is going on in detail to carry out his instructions. He does not need to know.

In an additional line of development, the most recent to be seen in use, obvious physical operations are undertaken by a machine under the control of a computer, without a human operator. The most common expression of this development is the articulated arm that is used to do arm and hand operations under programmed computer control. Because these computer-controlled arms are humanoid in appearance, we choose to call them by an old name, robots. In my view, the small computer that senses operations inside an engine and controls the various valves, injectors, and timings is as much a robot as the computer-controlled arms and at least as clever, though less obvious. The distinction is of no importance.

An important additional technological trend to insert into this discussion is the great increase in our ability not only to provide communication links, but also to use computers to control and route the information that goes through communication systems, to control the nodal connections and details of large and complicated networks, and even to provide format translations for different parts of the network.

The earliest effect of the communication and networking capability was the users' realization that the computer did not have to be displayed in the front office, where space is expensive. The computer can be put in the most convenient place suitable for its housekeeping, while the users and the rest of the network can be nearly anywhere. It is interesting that biology tends to put the computer (the brain) in what is in a sense the front office, closely connected with the principal sensors and communication outputs (the eyes, ears, and nose), but we have broken that close connection in our machine computing and communication capabilities.

Having introduced some of the technological elements that are now available, I would like to begin to bind them together into some concepts that are useful in thinking about the future. I will approach this problem by talking about biomedical engineering, since it has been the area in which some of the things I want to discuss have been most interesting, and because it introduces the possibilities in a way that involves nature.

We have built, experimented with, and begun to use what amount to implanted moderately intelligent substitute control systems for operations of the human body. The same technology that allows us to build a machine that can be sent to Saturn, where it can take data, care for itself, send information concerning itself and the data back to us, and receive instructions as to what to do next, is being used in the same way inside and around the human body. This symbiotic relationship between man and his clever machines is increasingly at the heart of the application of biomedical engineering. With an increasing understanding of the workings of the human body, coupled with an ability to sense and measure what is happening, small computers that can analyze these data, make decisions based upon previous instructions, and apply corrective control are making possible a new kind of replacement organ, supportive control system, and extension of human capabilities.

Implanted in the body, the heart pacemaker, in its simplest version, provides the timing and signal initiation for the electrical system that controls the contractions of the heart. A more complex version can be in a standby state, sensing the control waves and contractions of the heart and applying a pulse only when the rhythm of the heart falters or goes out of the heart rhythm that the pacemaker is programmed to accept. This activity requires an ability to sense the electrical or mechanical operation of the heart, and the capability to interpret these signals in terms of a preassigned program that decides whether they are within the range that it was told to leave alone. If outside the range, an appropriate signal may be applied. The device then decides whether the heart has responded suitably; it may then cease its activity or repeat its intervention.

Such a device may be carried one step further. In some types of heart disease and dysfunction, the rhythmic pulsation of the heart, which is initiated by regular electrical waves, gives way to an uncoordinated twitching and a set of electrical signals that are noise-like rather than regular; this is called fibrillation. One can build and implant an automatic defibrillator that can be combined with a pacemaker. The defibrillator senses when the motion and electrical current of the heart are nonrhythmic and noise-like. It then gives the heart a suitable electric shock through electrodes placed on the ends of the heart. Such a shock frequently reinitiates the rhythmic operation of the heart and is given in emergency medicine by the application of external chest electrodes. One formulation of such an implantable device will sense the required condition, produce such a shock, pause for a preassigned period of time to sense whether the heart has resumed its normal action, and then, if it has not, initiate a new shock. It will do this a specified number of times. In the program that I am aware of (Applied Physics Laboratory, Johns Hopkins University (APL/JHU) in conjunction with the Johns Hopkins Medical School), it will do this three times. It is almost as if a rather rudimentary machine paramedic that can sense a set of dysfunctions and take the appropriate prescribed action has been implanted in the chest of the patient.

A second kind of implanted machine is a drug-administration machine that medicates the patient in accord with a preprogrammed prescription. The instructions may operate in conjunction with data taken in the body of the patient. In one realization of this idea, the implant is in the abdomen. It consists of a reservoir of insulin, a computer, a power supply, and some communication and sensor devices. The insulin reservoir is connected via a catheter to the portal vein of the liver, which is an appropriate place for insertion of insulin. The computer is programmed by the physician to release insulin into the bloodstream of the liver in accordance with a prescription fitting the amounts and schedule that the physician feels should be administered. The program could also use some physical data from the patient, such as pulse rate. The computer can use these data as part of the input for its program. Given a good sensor for blood sugar, one could also incorporate that direct measurement into the program, and make it a feedback control system.

In the design of one such implant (APL/JHU), the patient has partial control over its operation. Using an external device, the patient can signal to the implant that he has made some decisions concerning future activities. In the case of a diabetic with the insulin administrator implanted, this might include such information as an intention to have a heavy meal within an hour or so, or an intention to exercise. Within the preprogrammed prescription, the machine will respond to this by beginning to alter the metering of insulin in anticipation of the expected event. Since this anticipation is restricted by the overall prescription, it will adjust the system to the patient's desires, but in such a way that it does not provide a danger should the anticipated events not occur. The insulin reservoir contains enough material for about three months of medication and can be refilled via an injection through the skin into a reservoir valve, which is immediately beneath the skin. This system is a fairly close equivalent to having a physician and some laboratory facilities in constant attendance on the patient.

Once implanted, these devices are not left completely alone. Via a telephone call, the implant can be made to report back to a computer under the control of the physician, and, via the same communications link, the physician may alter the programming, that is, the prescription, that is applied to the patient. In addition, in the case of the implanted insulin administrator, the machine records its own actions and any data it collects from the patient, and this too can be read out by the physician.

The implanted device can also monitor its own state of health, that is, examine its circuitry for normal indications. This includes the state of the power supply. Should it sense a difficulty with its own health, the device can signal the patient, by buzzing or by some similar signal, that it is not functioning properly and wishes the patient to consult with the physician. The battery may be recharged by a system of electromagnetic induction from a device on the skin outside the body to a coil under the skin.

Thus we have the beginnings of a medically programmable and controllable responsive implant system that is more than a mechanical replacement; it begins to partake of some of the internal control systems of the body itself and thus becomes nearly symbiotic, i.e., it becomes very close to a machine organism living with and performing functions in conjunction with the body itself. It seems clear that these examples are the beginning of a family of devices that can increasingly behave in more responsive ways as replacements and support systems in the body. An artificial heart will need to be more than a mere pump; it will need to be responsive to the requirements and needs of the body and therefore will need the kind of control system that I have described.

The analogy between the implant devices and the Voyager Spacecraft that recently flew by Jupiter and Saturn can be extended easily. The philosophy, design, and logic of some implants, the particular realization of the insulin administrator being one, are precisely the same as those involved in the design of a specialized spacecraft. The kinds of logic circuits, redundancy, communications systems, and control systems used are also the same. It is almost as if a miniature spacecraft were placed in the human body. Many of the system lessons from spacecraft design are now being applied in implant design.

I give these examples not because they are necessarily directly relevant to the discussion later, but rather because they introduce very graphically the idea of the use of involuntary or normally involuntary human activities as a signaling system for an internal robot.

Here a slight digression is in order. In most of our control of computers and robots to date, we have carefully delineated the interface between the person and the machine. As was noted earlier, the interface usually occurs in two places: at the tips of the fingers and at the eyeballs. That is, we operate a keyboard or a light pen or a joy stick with our hands, and we sense what the computer is doing by whatever it has been programmed to write out on a screen. This adaptation of traditional interfaces between people and machines to the human and computer interface has been convenient, but more intimate and less conventional interfaces may be important in the future.

External to the body, but also in response to the consequences of illness, we have begun to see the application of such technology to the operation of devices that can replace limbs or support inadequate muscle or nervous function. In the simplest sense, this has begun to be applied to the operation of the wheelchair for the paralyzed or otherwise extremely physically handicapped individual. It is clear that a buffering microcomputer can take simple signals in nearly any alphabet or signaling system and convert them into control signals for a wheelchair, a mechanical arm, or a similar device. For the purposes of such machinery, it does not matter whether the signals given by the person are puffs of breath, movements of the tongue or eyelids, fingers, toes, or muscles elsewhere in the body, or other controllable actions that produce changes in skin potential. In successful experiments, individuals have learned to control nerve signals directly as signals to a computer.

I do not discount the possibility of voice control, although at present, it is easier to use a coded system of sounds, whistles, or hums than it is to make devices directly responsive to ordinary language in any voice. Soon, however, we will be able to fit a detector to a trained voice and a specialized language, giving a signaling system that is very close to ordinary language for control purposes.

Given any of these signal systems and a receiver for them, a simple computer could interpret the control signals for, let us say, a wheelchair. This need not only to be control in the very simple sense of forward, back, left, or right; it could be more complex control, for example, an instruction to go to a describable place. This would be easiest if described in terms of distance and direction, but it would not be terribly difficult to make an instruction to go to a place that could be pointed out with a light beam, for example. We could build in obstacle avoidance, or stops for definable instructions or for special situations.

The same principles can be applied to a mechanical arm or leg, which could be driven by a separate power system and controlled via a computer, using the kinds of signals described above or body-motion signals. This is simply the application of a somewhat refined version of an industrial mechanical arm to direct control by human signals, with attachment to the body. By this means, it is possible to supply an individual with a very responsive and controllable replacement for a missing or totally damaged limb or to create a "dynamic brace" that would both support a weakened limb and, by being responsive to the remaining functions in that limb or to signals from some other part of the body, be a way of using the weakened arm or leg with a responsive external strengthening system that could make it function much more normally. We could construct a "walking brace" that would either amplify the remaining power and action in the weakened limb or walk in response to signals from the other leg or from other parts of the body. The individual would not need to decide how to walk (e.g., when to flex the knee), the program could contain that detail. The existing body is thus strengthened, and its motion made more possible and natural by the combination of sensory systems of the body, instructional controls, computational capability, and a power source.

There is no particular reason to restrict this technology to people who are handicapped or dysfunctional. Such a technique could be used on a perfectly healthy person in order to provide an extension of strength, dexterity, or reach. Indeed, such walking machines have been studied in the Department of Defense, but are more likely to find their most important realizations elsewhere.

This kind of capability is an extension in the direction of intimacy of control and naturalness of use from things that we already do. An automobile may be seen as a motorized wheelchair for healthy people, and the manipulators used in radiochemistry laboratories and the micromanipulators used in biological laboratories are rudimentary versions of what I am describing. The difference is that the control systems that we use for automobiles and micromanipulators are totally external to the body and are based on the use of hands and feet to maneuver controls that manipulate the machine. They also are based on separation of the control functions into a set of dimensions that originated in mechanical convenience. The new element is the ability to use other kinds of signals as control signals, and to combine them through computational capability in a way that bypasses many intermediate characteristics.

This entire technology can be described as an extension of human capability. I can combine my ability to use television cameras, communications means, controllable arms and legs, and the computational capability that puts this system together to place myself in the middle of the control loop of the system designed so that I may be in one place and yet see and hear what is happening in another place and manipulate a "body" in that place with a good sense of reality and presence. This technology, which has been dubbed "extensors," is an emerging capability that will have important uses, particularly in allowing humans to operate by remote control in places that are dangerous, distant, or otherwise difficult to deal with. Coal miners could mine coal from a control room outside the mine. A hand placed in a control glove can be moved to directly control the motions of a backhoe arm and jaw.

The general idea is not new: it appeared forty years ago in a story called "Waldo" by Robert A. Heinlein, and I am indebted to Bill McLean, and to Bill Bradley, formerly of IDA (Institute for Defense Analyses), for many interesting discussions of the possibilities.

Some very useful experiments have been done, including the remote driving of a truck from a remote control location. TV cameras on the truck were coupled through communication and control to TV receivers projected into a helmet in such a way that the cameras on the truck moved in synchronism with the motions of the helmet, e.g., when the operator turned his head, the TV camera eyes turned in the same way. The remote controls ran the controls on the truck. The remote operator "drove" the truck with a sense of realism and of being there.

At NOSC (Naval Ocean Systems Center), Kaneohe, Hawaii, some interesting experiments are currently being done with a helmet coupled to a TV-seeing robot, the head and torso of the robot moving with the motion (e.g., head turning) of the helmet wearer. The Argos/Jason project of Bob Ballard of Woods Hole Oceanographic Institution will exploit similar techniques and computer-controlled robots for remote, rapid exploitation of the seafloor.

There is, in fact, nothing to prevent a person from having several arms and hands under various kinds of control. For example, one could easily control an arm and hand with voice control while using one's own arms and hands, or extensions of them, and possibly even run multiple arms and hands with a combination of voice, foot, and hand control. It sounds rather like the realization of the image of a Hindu god.

There are other forms of connection between human and machine. A means of training robots in current use on production lines is as follows. The robot is set up at the job and a human hand guides it through the sequence of operations that it is later to follow. The appropriate sensory cues are pointed out to it while it is being guided through its operation. The brain of the machine is programmed to accept this information as a program-instruction series. In effect, the motion and operation of the machine are used as the instruction to its central computer for a later series of actions. One can instruct the robot in a whole variety of activities of different kinds, labeling each cluster of instructions with some cue. This cue may be an external sensory cue, either specifically provided or inherent in a task that is coming up. In effect, the robot recognizes what it is to do in the next sequence and carries out what it was taught. By this means, production line robots are taught to spot-weld a specific set of points on a car body or to spray paint a car in a particular sequence.

If you think it improbable or unlikely that you could learn to control an artificial arm, or an even more complicated sequence of mechanical actions, without thinking of the individual activities that are required to make the machine walk or reach, I remind you that you are quite familiar with learning sequences that do precisely that. After learning to drive or to fly, one does not usually think separately of the use of accelerator, pedals, steering wheel, stick, etc., but rather of a coordinated operation of driving or flying, which is carried out very much as one might walk, run, or play tennis. Another instructive example is the complicated set of emotional, intellectual, sensory, and motor activities that take place when one plays a musical instrument. The activity is not, except during the learning process, broken down consciously into the motions of fingers, muscles, and other detailed activity, but rather becomes a coordinated-control pattern, which, after a while, feels like making music, not like pressing keys.

Such a learning system is clearly applicable to the kinds of prostheses that I have described: not only prostheses for handicapped people, but also robot prostheses for healthy people. We do not know how many sets of intentional signals we could learn to operate in such a way, but I believe that we could operate many more than we have been able to so far and that the coordination problems could be overcome. The result would be a sort of "every man a Hindu god symbol" program, the idea being that one could with arms, or less obviously by using such control systems with other kinds of robots, exert considerably more coordinated human control than we are now capable of, or have even thought about using.

With our communication network capabilities, we can advance to a yet more complex robot operation system. An example of an extensive distributed communication network system is the ARPANET computer system, in which a number of computers around the country are linked in a communication network with a large number of users. A user can enter the net using a computer console and ask for data and for computational work to be done by the network. The network may then respond, and do the work, without the user's being aware what computer or computers in the network system are serving him, even to the extent that parts of the same computation may sometimes be done in different parts of the United States by different computers, with the results then assembled and presented to the user as a single result. The network looks like a single computer to the user.

With this in mind, it is clear that the robot arms, TV sensors, control computers, and programmed robots can be anywhere. One can control a system in which some parts of the system are doing things that they have been taught to do independently, by using preprogrammed or pretrained capabilities to recognize a task and carry it out, while other parts of the system are under direct control, being taught or being used directly as semi-intelligent hands, legs, arms, or other tools. As with the engine controller or the shuttle flight-control system, they may not look like arms or legs, but rather may operate as intricate machines controlled by computer instruction, or directly by a set of techniques learned by the user that are similar to those used in playing a musical instrument.

I am indebted to the writer Jerry Pournelle for the anecdote about the heavy-machine operator who grew tired of having to move from site to site as the jobs moved. He finally found a job in New York City. He came home and told his wife that he had a marvelous job, that he didn't have to move from site to site, and every day he went downtown and sat in an air-conditioned office in a fine office building. He sat in this office and ran a bulldozer on the moon.

I see no reason why, if the machines had a moderate degree of intelligence, he could not run a number of construction machines on the moon, dealing with them directly only when they have completed a task that they had been assigned and needed more instruction or when they sensed some difficulty. He would also have to maintain surveillance over their operations and be prepared to interrupt from time to time as necessary. Breakdowns and repair operations are another matter; at this point I am establishing a principle.

I would like to reiterate two points explicitly. First, it is important to recognize that the relationship between the machine or machines and the person running them is not necessarily an arm's length relationship, in the sense of being mediated by keyboard, that tells the machine step by step what to do in detail. It is far more likely that we will develop machine controls tailored to convenient learning processes for people; controls that are far more like playing a musical instrument, far more adaptive, both for the machine and for the person. I am really describing the road toward a kind of hybridization of people and their machines, in which one is not quite conscious of the distinction between self and machine but rather begins to use the sensory capability of the machine, and communication from the machine about what it is doing and what its reactions are, as part of one's own sensory and data-retrieving system. One might finally lose a sense of distinction between one's own motor activities and the motor activities of the machine, particularly if one introduces not only seeing and hearing but also force sensing translated to the operator as a pressure on a part of the body in proportion to the force sensed. The power outputs of portions of the machine, its muscle motors, for example, could also be transmitted to the operator, so that the operator can sense the effort being expended by the machine in a particular activity.

The second point to be reiterated is that the person who is working with, controlling, receiving sensory information from, or hybridized with the machine does not have to be close to the location of all or part of the machine. The machine can be anyplace where appropriate communication with it can be made available. In space, the only communication difficulty arises from the very large distances and the finite velocity of light. Communication between people and their machines must be tailored to how long it takes for that communication to arrive and the capability of the machine both to remember its instructions and to carry them out independently. In underwater use, the problem is the relatively small communication bandwidths available without cable (coax or glass fiber) connections and the mechanical design difficulty of such connections.

So far I have said little about the autonomous clever capabilities of machines. I have primarily described the ways in which machines, with a modicum of machine cleverness, can be used by people. It is clear, however, that we are beginning to construct and program machines that display "intelligent" characteristics. As I noted earlier, we can certainly make them adaptive, and we can teach them rather complicated algorithmic operations. For example, in the use of Landsat imagery of the earth, satellite imagery is composed of individual pixels (picture elements) corresponding to a few tens of meters on a side, each being represented by the satellite sensors as a breakdown of the reflectivity of an area on the surface of the earth in a number of electromagnetic spectral regions, some visible and some infrared. This gives a tremendous amount of information about areas on the surface of the earth, and this information can be manipulated to examine, for example, land use characteristics. The Census Bureau has made an interesting use of this capability with a clever computer program developed at the Goddard Space Flight Center. Using a particular synthetic display of a particular mathematical manipulation of the various spectral bands to represent the data, a human can look at the mosaic map of an area and characterize individual groups of picture elements as being characteristic of certain categories of land—farm or park land, single-dwelling residential areas, small-scale commercial areas, small-scale industrial areas, or large-scale industrial areas—depending on the spectral appearance of the regions in synthetic color. Goddard built an algorithm such that a human operator would identify a few regions on the picture display as having the characteristics of these particular kinds of land use. The computer system was then sufficiently clever to look at the entire picture in terms of the characteristics identified by the human and then classify all of the pixels in the display as falling into one of the land use categories. In effect, the computer program, the robot, identified and analyzed the characteristics the human was using for classification and applied them to the rest of the picture. Thus, with relatively little human intervention the computer could classify large tracts of U.S. territory with accuracies of above 95 percent. This is a particular human-machine combination in which the machine displays many important, clever characteristics, contingent upon the use of general program instructions.

We are using robots with vision (the eyes are separate from the arm) to select particular items from an assortment passing before them on a conveyer and to put them in an appropriate place or assemble them into a part on another conveyor. The robots cannot yet routinely pick specific parts from a jumble in a bin.

It is worthwhile to come back for a moment to the question of intelligence in machines. Since we have not yet been able to define human intelligence clearly or to agree on how to test for it, we are obviously in some difficulty. Turing, however, has proposed a test that does not depend on the definition of intelligence, but rather depends on the "humanness" of the machine. His proposition is that if you can communicate with a machine, have a dialogue with it, whatever the input-output mechanism, and at the end of the dialogue be unable to decide whether there is a concealed human being at the other end, then perhaps you are entitled to say that the computer, for the case in which there is no concealed human, is indeed intelligent.

There is now a class of programs called "expert programs" that do such things as medical diagnosis, given a set of clinical facts, or geological classification, given a set of geological findings. The programs are beginning to be able to mimic the operation of a human doing those tasks so well that one is entitled to be a little confused. Indeed, these programs are good enough to be useful as aids to human professionals. They are better than mediocre or poor professionals.

There are programs that behave as a certain kind of psychiatrist would and engage in a dialogue such that it is difficult to tell just from the dialogue, whether the conversation is with a machine or with that kind of psychiatrist, if not with any human being. With Turing's definition of machine intelligence in mind, this leads to the question of whether a psychiatrist operating in that particular therapeutic mode should be described as intelligent or merely programmed.

APPLICATIONS TO NAVIES

I now will discuss some applications of these concepts for possible use in naval systems and operations.

Let me begin with the simplest applications of sensor, computer, and control technology. Indeed, I will begin with a few small points that are only slightly related to robotics. In many areas the current design and operation of ships can best be described as antiquated. This has been true for at least a decade or two. The engines of a 747, which generate more power than those of most ships, can be run unattended for flight times of up to 13 hours, without more than cursory inspection between flights, for flight after flight for days or weeks, and so it seems ridiculous that we continue to operate engine rooms in naval ships as though modern control technology does not exist. To be fair, the 747 engine had the great advantage of being designed from scratch with new technology, while naval propulsion primarily consists of updated old technology, not well adapted to either the control systems or the maintenance and reliability philosophies that have developed in new systems. As I will point out later, however, it is entirely too expensive to save money by continuing with labor-intensive technological approaches.

This is not necessarily a plea for the replacement of steam turbines and/or diesels by gas turbines, or for the removal from ships of all maintenance people, who are the equivalent of the airline ground personnel who do routine servicing and maintenance; it is rather a plea for the replacement of old engineering with new engineering and technology. My own corporation, and indeed the entire auto industry, has not yet succeeded in carrying control and reliability technology as far as it should have to be a good example for the Navy. However, the auto industry too has been in a situation in which, to date, it has primarily been adapting new control technology to engines of older design. It is now in the business of designing the engines and the control systems together and should be able to produce the reliability necessary for the kind of operation I am suggesting for the Navy. In any case, automobile engines have been sufficiently reliable for quite a long time that opening the hood has not been a daily necessity, and we have certainly not had people in the engine compartment tending to them as they run.

Current human control activities for marine engine rooms, main engines, and all auxiliaries are unnecessary, unreliable, and poor design practice for naval ships. We need automated and roboticized control rooms, with the human controllers elsewhere in the ship, except for occasional maintenance and some extreme emergency operations. Design for continuous human tending is design for unreliability.

In a number of other areas, operations worthy of the trireme continue. An example is the operation of naval ships while one refuels the other. As far as I can tell, the control of such operations is still designed to be done by human estimation of distance and courses between the ships with no sensory or computer assistance. Control is through shouted orders to someone at a ship's wheel; the tension and inhaul of lines are adjusted by large crews of sailors also responding to shouted instructions. If this has recently changed, the fact has escaped my attention.

A rather simple sensor system of any one of a dozen kinds, working through a computer, could manage detailed ship's control with occasional human adjustment and do a better job of station keeping than a conning officer. The control of tension and winches also needs a good dose of modern technology.

The idea that a cold, bored 19-year-old talking into an inadequate telephone is a suitable substitute for a radar and an optical and infrared detector system, properly designed, working with a properly programmed computer, with the correct displays and alarms in a reasonable place, is ridiculous. It is a sample of antiquity sanctioned only by the sillier boundaries of admiralty law. If it is necessary to have lookouts, let us at least put them in a comfortable place with some sensory equipment to help them. I am familiar with all the tales about how radar is not a good substitute for visual scanning, and I agree entirely, unless we really begin to think about the problem of sensing the horizon and what is on it as something worth doing and build a proper system for the purpose. I have not insisted on taking people out of the loop, but rather on putting them intimately into the loop in a way that is useful.

It seems clear that the kind of machinery and possibilities that I described in the section on robots and the hybridization of people and machines describes a new set of approaches toward the operation of aircraft and ships, whether for sensing or weapons launch, and for the weapons and weapons systems that they use and control. What I am describing is a system in which rather complex machinery for sensing and controlling weapons, aircraft, ships, and devices is mixed with people in a new relationship. Some of the intelligence and capability, and many of the primary sensors, will be in the machines, with the people placed as intervenors in the operations of the machines. The intervention will be based either on separate information that the machines do not have (intelligence in the military intelligence sense, for example) or on conclusions from information coming from the machines. These conclusions may differ from the conclusions the machines alone would have reached given their programs and programmed level of cleverness. We have begun to approach some of this new man-machine relationship in electronic warfare systems in which a great many responses are made by the computer, but the operators control tactics and strategy and may interrupt and reinstruct the machines.

This approach also seems applicable to collections of weapons, in which the weapon may be intelligent not only in the sense of using sensory information to find things and decide whether or not they are targets and then arming itself toward them (this has begun with "smart" weapons), but also in the sense that the weapon may also keep a human operator informed of what it is doing and take further instructions from the operator wherever he is. This is an adaptation of my Hindu god symbol operation in which a person may have a variety of sensors and arms that he controls all at once or intermittently. The operator can do so because the sensors and arms are somewhat clever themselves, and therefore continuous attention to all of the machine partners is not necessary.

I should also remark that there is no particular reason why the launch source of the machines or weapons that are controlled by a particular human should be in any particular place. That is, a weapons controller need not be a weapons launcher and need not be a sensory systems carrier. We have begun that process of dissociation with E-2C's, AWACS, and similar devices, but it can be carried much further with profit.

It is clear that without ever referring to it specifically, I have introduced the concept of the remotely piloted vehicle. I would prefer not to refer to it as a remotely piloted vehicle or aircraft, but rather as a semicontrolled flying robot, so that it may be fitted more neatly into my overall hybridization scheme. It is clear that our control capabilities are such that the operations of landing and takeoff of aircraft can now be made essentially completely automatic. Even the restrictions that apply to the automatic carrier landing system are now no longer necessary, as has been demonstrated by the microwave landing systems developed by NASA for the FAA.

It is clear that a human operator who can operate heavy machinery, and even collections of heavy machinery, remotely can operate a collection of aircraft remotely, whether they carry weapons or are weapons. The operator may see out through their eyes and may manipulate or instruct their controls. Thus a person can have his control power greatly magnified through the use of a number of robots, with their cleverness and sensors, and the new control and computer technology.

The problem is not whether this is possible, the problem is to determine what tactical dimensions and uses of people and machinery this introduces that will be of interest. It certainly allows one to use machinery as a people magnifier for all sorts of military operations.

Having hinted at some of the applications of the robot technological revolution that I described in the first part of my talk, I would like to leave the rest of the application portion as an exercise for the student.

PROBLEMS

I will turn now to some problems, fundamental and otherwise, that will be encountered in trying to move in the directions I have described.

The first problem is that of reliability. Clearly, the systems that are implied by my discussion are complex, and the question of whether they will work and continue to work becomes crucial to the entire discussion. Reliability analyses, design for reliability, and great care in execution will be important in trying to use the kind of complexity that is inherent in these systems. Fortunately, we are entering an era in which it is possible to carry out the idea of superreliability for solid-state sensor systems and computer circuitry. The reliability of chip computers is already far greater than the reliability of most mechanical components. Due care must be paid to design, execution, and packaging for this to be true, but it is certainly possible that the era of "welded-up electronics" is upon us. The clever uses of redundancy in parts, in system functioning, and in network diversity are techniques that are beginning to be understood. They can lead both to high system reliability and to preplanned and largely controlled "graceful degradation" of the system if parts of individual machines, or of the network, do begin to fail.

It may be particularly important to change some of the traditional ideas about the economics of reliability in design and to pay the necessary costs for getting extremely high reliability in those areas of the system where it is possible, so that due attention to maintenance can be paid in those parts where superreliability is not possible.

A particular area of reliability that is only partly understood, and will be important in this kind of operation, is the reliability of software against the possibility of subtle problems in very complex software packages. In this connection, an early experience with the Voyager spacecraft, which went to Jupiter and Saturn and is now on its way to Uranus, may be instructive. Shortly after the first Voyager was launched, within the first day, it began to behave in a most peculiar manner. It was losing its navigational lock on sun and stars, it was reporting that it was healthy, and it was not accepting any instructions. The operators were in a frustrated state because they could not interpret what it was doing, nor could they manage to instruct it. It turned out that this problem was the result of carefully worked out software programs.

The underlying problem was that the naviagation system depended upon locking onto certain stars, and to the edge of the sun's disc, within a certain angular accuracy. The erection of the masts and other appendages after launch had initiated some vibrational disturbances that were greater than expected, so that the angular gates for finding and locking on the stars and sun were narrower than the excursions being forced by the vibrations. The robot was losing its lock upon stars and sun, expending gas to regain the lock, and then losing lock again due to the vibrations, which took quite a while to damp out. All attempts to communicate with the bird, to tell it to stop expending its gas supplies in reacquiring stars and sun, failed. There was concern that its gas supply would run low while the operators attempted a diagnosis.

Fortunately, the programmers found their error in time. The sequence of programmed instructions was as follows. The first instruction was, "when you lose lock on navigation more often than a certain frequency, you are sick." This a perfectly reasonable conclusion. The next instruction was, "when you are defined as sick, you will examine your state-of-health to find out what the problem is, and report." So far, so good. The system would lose lock, gain lock, and lose lock. After it happened often enough, it would then declare itself sick and proceed to examine itself for illness. Vibration beyond a certain degree was not by itself defined as illness, so the machine would always conclude that it was healthy, report, and go back to locking on, losing lock, etc. Unfortunately, there was one further instruction: "While you are examining your state-of-health you are not to accept any other instructions."

Once this unfortunate sequence of program instructions was understood, the programmers were able to examine the detailed sequence and find a very short time in between the various operations when the robot would accept instructions. At this point they inserted a "stop and listen" command and then corrected the various difficulties by reinstructing the machine as to new angles that it should accept for acquisition and lock-on of sun and stars. These angles, of course, did not need to be the same as the somewhat finer measurement angles.

This is a simple case of an unexpected difficulty in software, but in large programs such problems can turn up. This is an area in which great vigilance is required, because the software is crucial to what it is the machines really do in response to the instructions of their human partners.

I would like to say something about the question of economics and costs. In spite of the rapidly decreasing cost of computers, what I have described will be seen, probably correctly, as very expensive machinery and very expensive sensors and computers. Unfortunately, the military programs are costed in terms of the costs of objects and subsystems, not of total systems and navies. The difficulty is inherent in the procedures that are used; what is costed is widgets, not navies. That is to say, the comparison is always between the hundred-dollar bomb and the million-dollar smart weapon. I exaggerate for emphasis. Nobody ever costs the consequences of the facts that the bombs seldom hit anything and the missiles may or may not be reliable, much less the consequences of the fact that the bombs must be dropped by a pilot in an expensive airplane that goes nearly over, if not entirely over, the target, while the missile may be fired by somebody else from far away. The costs of reliability are generally exaggerated without being examined while the savings are seldom examined or adopted.

The consequences of this, as recently as the Vietnam War, were disastrous. I spent a good deal of effort trying to convince people that smart, long-distance weapons, though expensive, would be a lot cheaper than the cost of airplanes and pilots lost while dropping bombs that seldom hit the precise target. It would certainly have been worth some very expensive weapons if we could have attacked bridges, strong points, and obvious targets from fifty miles out at sea, with the pilot looking out through the nose of the weapon or otherwise sensing where it was he wanted the weapon to go. The cost arguments for not doing that were totally specious, because they did not take into account the cost of aircraft, the cost of MIA's and KIA's, in any cost sense whether dollar, or human, or political, nor did they consider the cost of the logistics and of providing for the people and the quantities of weapons and aircraft that were required.

When one replaces a person by a machine (which there is a possibility of doing if we play the reliability game right), then one can free a whole logistics tail, ship space, support of people, and operations support for use in a different way. I have been in some discussions on this subject, but they were never very sophisticated, nor were they carried to conclusion. The budget arrangements, the structure of the naval and budget system, and the construction of congressional committees, and indeed of the whole political system, are such that questions relating to this area cannot conveniently be asked; we cannot make the tradeoffs that are necessary to move into a new era of mixtures between people and machines.

Thus it will be important to configure an examination of the effect of this kind of system approach on logistics, on system economics, on manpower systems, on procurement, and on the political arrangements for deciding what systems to build and how to build them. This is the part that will be, from the point of view of the technologist, at least, the most difficult and least exciting part in the whole operation. It is, however, essential if anything useful is to come of this line of thought.

In order to carry this out, a better understanding of the kinds of technology that are possible for these purposes will have to be produced in the public at large and in the political leadership particularly. Only naval officers and technologists who can understand what this technology is and what it can do can be the agents of educational change for others so that they may understand that the social and economic computations must also be done correctly.

Finally, some of the human and ethical problems that can arise from this kind of technology are as important as the question of physical and engineering difficulties and cost problems.

It is not clear what kind of problems the habits and feelings of people who learned to be tightly coupled to machines would present. It is clear from my discussion that it is increasingly possible for a person to be intimately involved with a machine in such a way that the person's capabilities, whether damaged or whole, are extended beyond what they would be without the machine. Parts of this extension may be remote from the body of the person, operating through communication links, computers, sensors, and action systems that are elsewhere.

The remoteness of part of the system, and the nature of the programmed feedback control raise two issues that have been present in medical practice for a long time but have now become more pressing. First, the addition of the external devices, particularly if they are extended and partly elsewhere in space, raises the question "where are my own boundaries?" I might ask, "what part of this system is inherently me, and what part somehow attaches me to or blurs my connection with, the rest of things?" For example, if part of a medical therapeutic or support system is internal to me, but for resetting and some control purposes is in frequent or constant communication with a computer elsewhere, or a system monitored and controlled in part elsewhere, then I will have a particularly strong sense of connection to other people and other machines and other places. This gives me an intense feeling of dependency on a remote place, and a support that I cannot see and feel continually, which may give a new sense of insecurity and may furthermore blur my idea of where I am and perhaps even my idea of what it is that constitutes me. If I take the extensor control system, whether I am a healthy person or not, to the degree that Jerry Pournelle has suggested—that is, that my job in New York City is running a bulldozer on the moon—I may indeed develop a curious sense of difficulty about where I am. This raises some new psychological issues that have old philosophical connections but that will probably take some new learning to understand.

Second, there is the issue of control. Given the capability to use means that make people better than they were in certain physical directions, how do we distinguish between the repair of illness or handicap and the control or changing of a person? Referring back to the therapeutic implants I discussed earlier, is it always clear where the balance point on insulin administration is? Given that the patient and the physician have some control over this balance, in what sense is it the patient himself who controls his own body and in what sense is he Trilby to an absent Svengali, operating through a machine? While the issue of control is easy to consider in the case of insulin administration, when one comes to the possibility of the administration of brain-related chemicals in a program intended to correct a chemically based psychiatric disorder, fundamental questions arise about the ability of the person to be in control of himself while a prescription is operating in his body under the design of someone else. While this has been a problem in psychiatric practice for some time, it seems to me that the possibility of putting the prescription in an implanted program, or a program run by computer control, raises a deeper set of issues, since the patient now does not have the choice of taking the medication or not—that choice has been made.

In the use of extensions of the normal body, the ethical questions, particularly if the extension is accompanied by changes in body chemistry, become even more acute.

Some of these problems have, as we know, been treated in fiction, in Michael Crichton's *The Terminal Man,* and elsewhere. Now they are becoming as real as they were imaginative.

Having introduced some examples of where current technology may be taking us and having mentioned some of the problems that arise in using this technology, I would like to end by asking a rather extreme question, going to this limit in order to illuminate some possibilities and problems.

I will phrase the question as follows: "Given an increasing capability of computers to store and analyze material and given that they are programmed appropriately, in what sense might it be possible to pour myself into a machine?" Clearly, there are a number of biological functions that would be irretrievably altered, but would it be possible for me in some sense to put my ideas, something about my reasoning processes, and my views into a computer, and what would that be like?

Let us go to the limit by steps. We are on the verge of the replaceable heart. We know something about the directions in which to go to provide implantable dialysis, i.e., the implantable kidney. Many muscular operations are relatively straightforward, and we can begin to use computers more and more directly as containers of a memory and information to use with current brain function. If I work day in and day out with a very capable machine, storing more and more of my desired information, my computations, the things I think about, my associations, and descriptions of my emotions (which may have a special meaning to the computer because it may be attached to the implants that monitor physiological functions responsive to emotion), will I approach some end state in which the distinction between the part of the machine that is machine and the part that is me becomes blurred, and perhaps finally proceed to a state in which I am in some sense more "in" the machine than "in" myself? Is it only the limitation of time needed for programming that prevents me from "becoming" a machine that operates in the same way that I myself would, given the same input sensations and information? Different schools of psychology give different answers to this question. When I started to contemplate it, I was certain the answer was in some sense yes, but the more I think about the problem, the less easy it is to come to that conclusion.

Man has been described as the tool-making and tool-using animal. Our tools are becoming more sophisticated and are beginning to have properties that enable us to fit them to ourselves, and perhaps ourselves to them, in a way that is far more intimate than the mere molding of the shape of a stone to the hand and fingers. While no different in kind, the intimacy now gives us not only a capability to repair and support certain ills that, at least as yet, we do not know how to deal with biologically, but also an opportunity to go far beyond our normal capabilities, and to do so in a way that will feel increasingly natural, although at the same time, as I have suggested, it may change the nature of what we mean by feeling natural.

While these final comments about man-machine relationships may seem remote from the naval questions, if these directions are seriously pursued, these comments will need to be considered. A naval system in which people and their robot and machine partners are as completely hybridized as I have implied they could be would be a new kind of experience for the people, and the real problems could be those implied in my philosophical coda, rather than the simpler ones to be encountered at the beginning of development in this direction.

I do not know what directions we will choose for the use of the technology of computers, communications, and robots for naval purposes, but I hope that this exercise in laying out some possibilities will be stimulating and hence useful.

CURRICULUM VITAE
DR. ROBERT A. FROSCH

Dr. Robert A. Frosch was elected vice president of General Motors in charge of the Research Laboratories on March 1, 1982. In that post, he heads a science team engaged in applied research and development, long-range research, and specialized service work for other GM units.

Born May 22, 1928, in New York City, Dr. Frosch attended Columbia University, earning his B.A. in 1947, M.A. in 1949, and Ph.D. in theoretical physics in 1952. While completing the studies for his doctorate, he joined Columbia's Hudson Laboratories in 1951 as a research scientist working on naval research projects. He rose to director of Hudson Laboratories in 1956 and held that position until 1963.

In 1963 Dr. Frosch became director of Nuclear Test Detection for the Advanced Research Projects Agency (ARPA) of the Department of Defense. In 1965 he was appointed deputy director of ARPA.

Appointed Assistant Secretary of the Navy for Research and Development in 1966, Dr. Frosch served in this capacity until 1973, when he was selected as assistant executive director of the United Nations Environment Program, with the rank of assistant secretary general of the United Nations.

In 1975 Dr. Frosch became Associate Director for Applied Oceanography at Woods Hole Oceanographic Institution and served in this position until joining NASA. He served from 1977 to 1981 as the administrator of NASA.

In January 1981 Dr. Frosch was appointed president of the American Association of Engineering Societies, Inc.

Dr. Frosch has received the Arthur S. Flemming Award (1966), the Navy Distinguished Public Service Award (1969), the Defense Meritorious Civilian Service Medal (1973), the Neptune Award of the American Oceanic Organization (1973), and the NASA Distinguished Service Medal (1981). He is a member of Phi Beta Kappa and Sigma Xi honorary societies.

His memberships in professional societies include the following: National Academy of Engineering, American Physical Society, Seismological Society of America, Marine Technology Society, Society of Naval Architects and Marine Engineers, Society of Exploration Geophysicists, and the American Geophysical Union. He is a fellow of the American Association for the Advancement of Science, the Acoustical Society of America, the Institute of Electrical and Electronics Engineers, the American Institute of Aeronautics and Astronautics, and the American Astronautical Society.

Dr. Frosch is married to the former Jessica Rachael Denerstein of Brooklyn, New York. They have two daughters, Elizabeth Ann, 21, and Margery Ellen, 19.

SEARCH COMMITTEE FOR THE CHARLES H. DAVIS LECTURE SERIES

DR. Harrison Shull (CHAIRMAN)
Vice President Academic Affairs
& Provost
Rensselaer Polytechnic Institute

CDR Ronald Atkins, USN
Director of Support Services
Naval War College

DR. George F. Carrier
Division of Applied Sciences
Harvard University

VADM John T. Hayward, USN (RET.)
3 Barclay Square
Newport, RI 02840

DR. Sidney G. Reed, Jr.
11104 Waycroft Road
Rockville, MD 20882

DR. Joseph M. Reynolds
Vice President, Instruction
and Research
Louisiana State University

DR. William M. Tolles
Dean of Research
Naval Postgraduate School

MR. Lee M. Hunt (STAFF OFFICER)
Executive Director
Naval Studies Board