INDUSTRIAL ROBOT TECHNOLOGY AND PRODUCTIVITY IMPROVEMENT

James S. Albus
Industrial Systems Division
National Bureau of Standards

Robots have received a great deal of publicity recently. The movie "Star Wars" and several television series such as "The Six Million Dollar Man" and "The Bionic Woman" have raised the consciousness of the public to the subject of robots. The enormous influx of foreign cars manufactured in part by robots has aroused awareness of the press and many politicians to the fact that robots can have a profound effect on industrial productivity. Many people today believe that the robot revolution is well under way, that factories are full of armies of highly intelligent robots, and that human workers are being displaced in droves. The facts are quite different.

First of all, there are only about 3000 robots installed in the entire country, secondly, the great majority of these are quite primitive, with no capacity to see or feel or respond to their environment in any significant way.

Most people think of a robot as an android, which walks and talks, sees and feels, and looks much like C3PO, or at least R2D2. Real robots are much more primitive. In its simplest form a robot is nothing more than a mechanical device that can be programmed to perform some useful act of manipulation or locomotion under automatic control. An industrial robot is a device that can be programmed to move some gripper or tool through space so as to accomplish a useful industrial task.

These robots are typically programmed by recording each task as a series of points in space. This recording is then simply replayed whenever the task is to be performed.

This simple procedure is adequate to perform a surprising number of industrial tasks, from spot welding automobile bodies, tending die casting machines, loading and unloading machine tools and presses, spray painting, and performing a wide variety of materials handling tasks.

Even arc welding can be performed by a robot which can neither see nor feel, so long as the parts to be welded are positioned in exactly the right place, and the welding parameters are controlled by some automatic system.

However, the great majority of industrial tasks are beyond the capacities of present day robot technology. Most tasks are too complex and unstructured, or involve too many uncertainties, or require too much ability to see and feel and adapt to changing circumstances. Before robots can significantly impact productivity of the economy as a whole, they must be used in hundreds of thousands and even millions of applictions. This will not be possible before a large number of technical problems are solved.

TECHNICAL PROBLEM AREAS

One of the first problems is accuracy. Robot positioning accuracy needs to be improved. Although the repeatability of most robots is on the order of 0.050 inch over its working volume (and in some cases as good as .005 inch), the absolute positioning accuracy may be off as much as 0.250 inch, or even 0.500 inch in some regions of the reach envelope. Thus, it is not possible to program a robot to go to an arbitrary mathematically defined point in a coordinate space and have any assurance that the robot will come closer than a half of an inch. This creates major problems in programming a robot from a computer terminal, or in transferring programs from one robot to another. Each robot must be taught its program separately by leading it point by point through its job, a tedious and costly task.

Presumably, this accuracy problem could be solved through closer robot manufacturing tolerances, although not without cost. Alternatively, calibration procedures such as illustrated in Figure 1, might allow each robot to offset its off-line program points to compensate for its mechanical inaccuracies. However, no efficient methods of robot calibration have yet been developed, and robot control software is not presently designed to use calibration tables for improving absolute positioning accuracy. Until this absolute position accuracy problem is solved, robot assembly in the small batch environment will be uneconomical. Teaching a robot every point in the trajectory of a complex assembly task is a time consuming job which may take many times longer than would be required to perform the same task by hand. Thus, using a robot for small lot batch assembly cannot be economical until software can be efficiently produced by off-line programming (i.e., programming from a computer terminal).

Second, dynamic performance must be improved. Present day robots are too slow and clumsy to effectively compete with human labor in assembly. Two possible exceptions to this are in arc welding where speed is governed by the welding process itself, and spot welding where the task corresponds to moving a heavy welding gun through a simple string of points in space -- a procedure which the robot is particularly adept at executing. However, if robots are to perform other types of assembly and construction tasks, they must be able to execute much more complex routines with much greater grace, dexterity, and speed than they are now capable of. Control systems need to be alternately stiff and compliant along different axes in space (which do not generally coincide with joint coordinates). This requires much more sophisticated cross-coupled servo control computations than are presently employed.

Furthermore, robot structures are typically quite massive and unwieldly. Most robots can lift only about one tenth of their own weight. Many cannot even do that. New mechanical designs using light weight materials such as carbon filament epoxies and hollow tubular construction are needed. Advanced control systems that can take advantage of such light weight structures and high speeds will be a major research project.

Much also remains to be done in gripper design. Typically, robot hands consist of pinch-jaw grippers with only one degree of freedom -- open and shut. Contrast this with the human hand which has five fingers, each with four degrees of freedom. No robot has come close to duplicating the dexterity of the human hand, and it is not likely that one will in this century. Certainly, dexterous hands with jointed fingers for industrial robots are a long way in the future. The problem is not so much in building such a mechanical structure, but in controlling it. No one has any idea how to design control algorithms to make use of such complexity and very little research is being done in this area.

Third, sensors of many different types must be developed. Robots must become able to see, feel, and sense the position of objects in a number of different ways. Processing of visual data must become faster and be able to determine 3-dimensional shapes and relationships. Robot grippers must become able to feel the presence of objects and sense the forces developed on those objects. Proximity sensors are needed on robot fingertips to enable the robot to measure the final few millimeters before contacting objects. Longer range proximity sensors are needed on the robot arm to avoid colliding with unexpected obstacles. Force and touch sensors are needed to detect and measure contact forces. A variety of acoustic, electromagnetic, optical, x-ray, and particle detectors are needed to sense the presence of various materials such as metals, ferromagnetics, plastics, fluids, and limp goods, and to detect various types of flaws in parts and assemblies. Both the sensing devices and the software for analyzing sensory data represent research and development problems of enormous magnitude.

Robot sensors is an area where there is much research activity. Robot vision is by far the most popular research topic, and also probably the most difficult. A computer must treat a visual image as an array of brightness dots called picture elements, or pixels. A typical scene may consist of from 16 thousand to over a million pixels. Interpretation of such a large volume of data is an enormous task even for a high speed computer. It often takes many seconds to several minutes to analyse a single picture by computer. This is far too slow for the robot to respond in a timely fashion to what it sees. Various tricks are used to speed up this response time. One is to illuminate the scene so that the objects appear as black and white silhouettes. Another is to assure that no two objects of interest touch or overlap. However, even under such artificial circumstances robot vision is a very complex problem and subject to many difficulties. Such techniques obviously limit the use of robot vision to a few select applications.

Other robot sensory inputs such as touch and force appear to be simpler in principle, but much less work has been done in these areas.

Fourth, control systems are needed which can take advantage of sophisticated sensory data from a large number of different types of sensors simultaneously. Present control systems are severely limited in their ability to modify a robot's behavior in response to sensed conditions. Robot control systems need to be able to accept feedback data at a variety of levels of abstraction and have control loops with a variety of loop delays and predictive intervals. See for example, Figure 2. Sensory data used in tight servo loops for high speed or high precision motions must be processed and introduced into the control system with delays of no more than a few milliseconds. Sensory data used for detecting the position and orientation of objects to be approached must be available within hundreds of milliseconds. Sensory data needed for recognizing the identity of objects or the relationship between groups of objects can take seconds. Control systems that are properly organized in a hierarchical fashion so that they can accommodate a variety of sensory delays of this type are not available on any commercial robot.

Fifth, robot control systems need to have much more sophisticated internal models of the environment in which they work. Future robot control systems will have data bases similar to those generated by Computer-Aided-Design (CAD) systems, and used for computer graphics displays. These can describe the three dimensional relationships of both the workplace and the workpieces. Such data bases are needed to generate expectations as to what parts should look like to the vision system, or what they should feel like to the touch sensors, or where hidden or occluded features are located. Eventually, such internal models might be used in the automatic generation of robot software; for example, by describing how a finished assembly should look, or even how each stage of an assembly or construction task should appear in sequence.

Sixth, techniques for developing robot software must be vastly improved. Programming-by-teaching is impractical for small lot production, especially for complex tasks where sensory interaction is involved. Shop floor personnel unskilled in computers must be able to instruct robots in what to do and what to look for in making sensory decisions. Eventually, it will be necessary to have a whole range of programming languages and debugging tools at each level of the sensory-control hierarchy. The development of compilers and interpreters and other software development tools, as well as techniques for making use of knowledge of the environment derived from a number of different sensors and CAD data-bases are research topics that will require hundreds of person-years of highly skilled systems software talent.

Seventh, interfaces need to be defined in some standardized way, so that large numbers of robots, machine tools, sensors, and control computers can be connected together in integrated systems. Trends in the field of computer-aided-manufacturing are toward distributed computing systems wherein a large number of computers, robots, and machine tools all interact and cooperate as an integrated system. This creates enormous software problems. Particularly in the case where sensors are used to detect variations in the environment and to modify the control output to compensate for those variations, the software can become extremely difficult to write and virtually impossible to debug. In order for such systems to work at all, it is necessary to partition the control problem into modular components and then develop interface standards by which the various system components can communicate with each other. See Figures 2 and 3.

It is often felt that standards are an inhibiting influence on a newly developing field -- that they impede innovation and stifle competition. In fact, just the opposite is true. Well chosen interface standards promote market competition, technology development, and technology transfer. They make it possible for many different manufacturers to produce various components of modular systems. Standard interfaces assure that multivendor systems will fit together and operate correctly. Individual modules can be optimized and upgraded without making the entire system obsolete. Interface standards also make it possible for automation to be introduced incrementally -- one module at a time. Systems can be made upward compatible and automated piecewise. Thus, users can test the automation waters gradually, without a large initial capital barrier.

Eighth, many potential robot applications require robot mobility. Most robots today are bolted to the floor or to a tabletop. Small robots can reach only one or two feet while larger ones can grasp objects nine or ten feet away. But many applications need robots which can maneuver over much larger distances. For example, a robot used to load a machine tool typically spends most of its time waiting for the machine tool to finish its operations. Sometimes a single robot can be positioned between two or more machine tools so that it can be more fully utilized. However, this leads to severe crowding of the work environment and in many cases is simply not practical. There are a few applications in which robots have been mounted on rails so that they can shuttle between several machines. Unfortunately, to date this has proven too expensive and cumbersome for wide scale use.

In many applications, particularly in arc welding of large structures like ships or buildings it is not practical to bring the work to the robot; the robot must go to the work, sometimes over distances of many tens of feet. One example is in the construction of large machinery such as road building equipment. Another example is in the building of ships. A good ship building robot would be able to maneuver inside odd shaped compartments, climb over ribs and bulkheads, scale the side of the ship's hull, and weld seams several hundred feet in length. Similar mobility requirements exist in the construction of buildings. Construction robots will need to be able to manuever through the cluttered environment of a building site. In some cases they will need to climb stairs, and work from scaffolding.

Robots will also be used in undersea exploration, drilling, and mining. Robot vehicles will someday explore the moon and planets. These applications will require significant new developments in mobility mechanisms.

Robot mobility in the factory using rails, carts, or overhead conveyors is a relatively simple problem that undoubtedly will be solved in the decade of the 1980's. Robot mobility on the construction site, under the sea, and in outer space however, is another issue entirely. The sensor, data processing, and control problems associated with these aspects of robot mobility will require years of concentrated research.

For the most part, these eight problem areas encompass profound scientific issues and engineering problems which will require much more research and development. It may be possible to improve the mechanical accuracy of robots, and to improve servo performance with little more than careful engineering. But much more fundamental research and development will be required before the sensor, control, internal modeling, software generation, systems interface, and mobility problems are solved. Much remains to be done in sensor technology to improve the performance, reliability, and cost effectiveness of all types of sensory transducers. Even more remains to be done in improving the speed and sophistication of sensory processing algorithms and special purpose hardware for recognizing features and analyzing patterns both in space and time. The computing power that is required for high speed processing of visual and acoustic patterns will even require new types of computer architecture.

Sensory-interactive control systems that can respond to various kinds of sensory data at many different levels of abstraction are still very much in the research phase. Current commercial robot control systems do not even allow real-time six-axis incremental movements in response to sensory data. None have convenient interfaces by which sensory data of many different kinds can be introduced into the servo loops on a millisecond time scale for true real-time sensory interaction. None of the commercial robot control systems have anything approximating CAD data bases or computer graphics models of the environment and workpieces. Finally, current programming techniques are time consuming and not capable of dealing with internal knowledge or sophisticated sensory interactions.

These are very complex problems that will require many person-years of research effort. It is thus not surprising that the robot applications are still extremely limited.

WHAT LIES IN THE FUTURE?

All of the problems listed above are amenable to solution. It is only a matter of time and expenditure of resources before sensors and control systems are developed that can produce dexterous, graceful, skilled behavior in robots. Eventually, robots will be able to store and recall knowledge about the world that will enable them to behave intelligently and even to show a measure of insight regarding the spatial and temporal relationships inherent in the workplace. High order languages, computer-aided-instruction, and sophisticated control systems will eventually make it possible to instruct robots using much the same vocabulary and syntax that one might use in talking to a skilled worker.

There is no question that given enough time and resources robotics will eventually become a significant factor in increasing productivity in industrial production. The question is: How much time and how many resources will be required before this becomes a reality?

In my opinion more than a few tens of millions, and less than a few hundreds of millions of dollars for research and development will be required to make robots capable of performing a sufficient number of tasks to make significant productivity improvements in industrial manufacturing. More than a few hundred and less than a few thousand person-years of high level scientific and engineering talent will be needed before robot software of sufficient complexity can be generated economically for small lot batch production. In other words, a national research and development effort of at least one, and perhaps two, orders of magnitude greater than what has been done to date will be required to produce a significant impact on industrial productivity. And more than just total dollars spent is important. Robotics research is systems research. At least a few stable, consistantly well funded research centers of excellence will be required.

The questions then are:

"How fast are we progressing along the road to the solutions?"
and
"Who are the researchers that are leading the way?"

In the United States there are four types of research laboratories:

1. University
2. Non-profit
3. Private Industry
4. Government

UNIVERSITY RESEARCH

Among the principal university labs are:

Stanford University: The robotics effort at Stanford is of long standing. Tom Binford has been doing pioneering work in three-dimensional vision for over a decade. His students have developed one of the most advanced robot programming languages available today called AL, for Arm Language. The Stanford artificial intelligence lab has produced a long list of ground breaking research projects in manipulation, hand-eye coordination, and robot assembly. Stanford is presently working on robot vision, a three-fingered hand, force sensing, robot programming languages, and geometric modeling for vision and programming. They also have a cooperative program with Unimation for robot mobility.

Stanford received about $200K in FY81 from NSF. There are about 14 graduate students working on various projects.

MIT has had a major robotics effort at least as long as Stanford. At present, Danny Hillis and John Hollerbach are building robot skin made of thin sheets of rubber lined with tiny wires that detect pressure. These are being used to give robots a sense of touch. MIT also is active in robot vision and programming languages. Tom Sheridan of MIT is working on Supervisory Control of Teleoperators. This work is currently directed toward undersea work and is partially funded by Naval Ocean Systems Center in San Diego. Total MIT funding is around one million per year. Office of Naval Research provides approximately 700K of this amount.

Carnegie-Mellon University has recently formed a Robotics Institute directed by Raj Reddy with funding from Westinghouse, ONR, DARPA and other industrial sponsors. The Institute has programs in flexible assembly, machining, sensory systems, vision, mobility and intelligent systems. In its less than two years of existence the Institute has recorded significant achievements in the expansion of sensory capabilities of machines, the integration of several machines into cells carrying out complex tasks, the application of vision and optics to a wide range of industrial tasks, the development of new robot mechanisms, and the application of artificial intelligence to the management of evolving intelligent technologies. Total funding is over $3 million, making it one of the best funded major university projects. Office of Naval Research contributes approximately 500K per year to Carnegie-Mellon University.

Rhode Island University has an impressive effort directed by John Birk on general methods to enable robots with vision to acquire, orient, and transport workpieces. The Rhode Island robot was the first to pick parts out of a bin of randomly oriented parts. Rhode Island is also doing work on dexterous robot grippers and robot programming languages. Funding from NSF is $210K per year and from industrial affiliates, about $750K per year.

University of Florida under Del Tessar is doing work in teleoperators, force feedback, and robot kinematics and dynamics. Funding from the Department of Energy, NSF, and State of Florida amounts to about $1 million per year.

Purdue University is doing research in robot control systems, robot programming, languages, machine vision, and modeling of part flow through industrial plants. Total NSF funding to Purdue is about $400K over a four year period.

A number of Universities have smaller robotics efforts, or efforts in related areas.

The University of Massachusetts is doing work in visual interpretation of natural scenes and design of parts for automatic assembly. ($125K per year) They have just received an NSF grant for $157K to study "Economic Applications of Assembly Robots".

University of Maryland Computer Vison lab under Azriel Rosenfeld is doing work on a number of image processing projects including robot vision and methods for using visual knowledge in interpreting images. (over $1 Million per year)

University of Rochester under Herb Voelcker is developing advanced methods of representing three dimensional shapes in a computer memory. The result of this work is a computer graphics language called PADL which is profoundly influencing the way future computer graphics systems are being designed. Much of this is being done with NSF funding. ($85,576 in FY81)

Rensselear Polytech Institute under Herb Freeman is also studying the generation of computer models for three-dimensional curved surface objects. ($98K)

University of Arizona is doing teleoperator work. ($113K)

University of Wisconsin is doing work in machine vision. ($60K)

Ohio State University under Robert McGhee is working on dynamics and control of industrial manipulators and legged locomotion systems. ($125K from NSF) DARPA has recently funded McGhee to build and test a man-carrying walking machine. This project is funded at $250K in FY81 and $630K in FY82. Battelle Labs are cooperating with Ohio State University in this effort.

University of Illinois, University of Pennsylvania, University of Washington, and the University of Texas all have small research projects in robotics, and robot related work.

Total National Science Foundation funding for university research in robotics and related fields is on the order of $5 million per year. Additional university funding from other sources such as industrial affiliates and internal university funding may run another $4 million per year. University research tends toward small projects of one or two professors and a few graduate students. The average NSF grant in robotics and related fields is around $150K per year.

Although support of university research by industry is on the rise, it is still small by European or Japanese standards. University efforts tend to be fragmented, progress is sporadic, and the issues addressed are often unrelated to the problems of industrial manufacturing.

NONPROFIT LABS

C. S. Draper Labs with Jim Nevins and Dan Whitney have been studying part-mating science and assembly system design for a number of years. They have performed a variety of assembly experiments, studied the use of force feedback, and developed a theory of the use of passive compliance in part-mating. Draper has also done economic modeling for designing industrial systems, and real-time simulation of the space shuttle remote manipulator system for NASA. NSF funding is about $200K per year. Draper also has a number of industrial clients for whom it performs design and construction of advanced assembly systems. Total funding is about $1 Million per year.

SRI International has an extensive robot research program that dates back to the SHAKEY Artificial Intelligence project that was funded by ARPA in the late 1960's. Presently SRI's program is headed by David Nitzan. Emphasis is on machine vision for inspection and recognition. Some very sophisticated robot vision research is being done on overlapping parts using structured light and a combination of binary and gray-scale vision. Work is also being done on printed-circuit board inspection, programmable assembly, vision-guided arc welding, and semiautomatic process planning. Funding from NSF is about $350K per year with about $350K per year from industrial affiliates. SRI was the first robotics lab to develop an industrial affiliates program. Office of Naval Research contributes approximately 250K for research in communication and negotiation between cooperating robots to distribute their workload. Additional $250K per year funding from NSF started in August 1981 for work on printed-circuit board inspection.

PRIVATE INDUSTRIAL RESEARCH LABS

General Motors has established a major robotics research effort at the G. M. Research Labs in Warren Michigan. They have concentrated on vision and have produced a new robot vision system called "CONSIGHT". This system has a unique method for obtaining silhouette images of parts on a conveyor belt that does not require back lighting and is not dependent on contrast between the part and the belt. General Motors is also interested in small parts assembly by robots and automatic inspection. Several years ago they contracted with Unimation to produce the PUMA robot; a small, accurate, computer controlled robot designed for assembly.

General Electric is becoming very active in robot research. G. E. has a substantial research effort in robot assembly, robot vision, robot controllers and new VLSI micro circuit technology. They have designed a very impressive laboratory robot which embodies a number of innovative concepts. G. E. also has a robot demonstration facility where they have one of almost every robot manufactured today. As a part of this facility they offer courses in robot programming and applications engineering. G. E. has also announced intentions of marketing the Italian PRAGMA robot in this country under the name of ALLEGRO, as well as the Hitachi Process Robot.

Westinghouse has established a productivity center in Pittsburgh with a robotics research lab containing 15 robots of all different kinds. This center supports Carnegie-Mellon University with $1 million per year grant for manufacturing research. Westinghouse also has a cost sharing project with NSF called APAS for Adaptable Programmable Assembly System. This research project will be complete in 1982. It has been funded by NSF at about $500K per year. Westinghouse also has a R&D center which is working with the University of Florida to assess what teleoperator technology is needed for nuclear power plants.

IBM has been involved in robotics research for a number of years. IBM has developed robot programming languages called AUTOPASS and EMILY and has studied the problem of robot assembly. IBM has also developed its own robot which it uses in its own manufacturing operations. All of the IBM robotics effort is internally funded and details of the projects are not available.

Texas Instruments also has developed a robot which they use for assembly and testing of hand calculators. No details of this effort are available.

Martin-Marietta has a robotics effort directed primarily toward NASA and DOD interests. They are working on automated diagnosis and checkout of avionics, cockpit simplification, and various autonomous devices. Martin is also studying the speed requirements for space shuttle manipulators, coordinate transformations, and two arm coordination. Funding is about $3 million per year.

Automatix is a small new company with a heavy emphasis on robotics research. Robot vision, microcomputer control systems, and applications engineering in arc welding systems are their main target areas.

Machine Intelligence Corporation is another small company, whose technical staff includes the principals who pioneered robot vision at SRI International. Machine Intelligence Corporation manufactures computer vision systems to be incorporated into turnkey inspection, material-handling and assembly systems. In cooperation with Unimation Corporation, they have developed the Univision system, the first commercially-available "seeing" robot, marrying an advanced vision system with the PUMA robot, programmable under a special language "VAL". They have an NSF Small Business Inovnovation grant for research on a method of person/robot communication, to permit programming a robot without need for a professional programmer.

ROBOT MANUFACTURERS

The major robot manufacturers, of course, also conduct a substantial amount of research. Unimation is working on advanced control systems, calibration techniques, mobility systems, and programming techniques.

Cincinnati Milicron has a research group working on new control system architectures, programming languages, and mechanical design.

Prab-Versatran, Autoplace, Advanced Robotics, Devilbiss, Mobot, Nordson, Thermwood, ASEA, KUKA, Tralfa, U. S. Robots, and perhaps ten other small new robot companies are all aggressively developing new and improved product lines.

The level of funding for research by the robot manufacturers is proprietary. However, based on the aggregate sales of about $150 million for the entire U. S. robot industry, it is probably around $15 million per year and scattered over about twenty companies. One or two of the largest manufacturers are spending around $5 Million per year on research. However, it is doubtful if more than three manufacturers are spending more than $1 million per year.

GOVERNMENT RESEARCH

The National Bureau of Standards is pursuing research related to interface standards, performance measures, and programming language standards for robot systems and integrated computer-aided-manufacturing systems. This work focuses on advanced concepts for sensory-interactive control systems, modular distributed systems, interfaces between modules, and sensor interfaces to the control systems of robots and machine tools. Funding from the Department of Commerce is about $1.5 million per year.

The Air Force Integrated Computer Aided Manufacturing (ICAM) project has funded several robot development and implementatiion projects. A contract with General Dynamics introduced robots into drilling and routing applications in aircraft manufacturing. A contract with McDonnell-Douglas resulted in a robot programming language based on the APT N/C tool language. A contract with Lockheed Georgia produced a study of potential future aerospace applications for robots. Total funding was about $1 million per year. This work is now completed. Technical Modernization, a related program is presently funding General Dynamics to design several aspects of an automated factory. Funding for this is about $4 million per year. Total ICAM funding is $17 million per year for computer based information, planning and control, and systems engineering methodologies for increased automation. Estimated future ICAM funding for robotics is $2 million per year.

NASA has a number of small robotics projects at several of its centers. JPL has a project in stereo vision, force feedback grippers, and the use of automatic planning programs for mission sequencing applications. Langley Research Center is doing research on robot servicing of spacecraft. Marshall Space Flight Center has developed a prototype robot arm for satellite refurbishing and is working on free-flying teleoperators. Johnson Space Center is managing the development of the space shuttle remote manipulator system. The total NASA reseach budget for automation is about $2 million.

The Naval Air Rework Facility in San Diego is funding the development of robots to remove rivets and fasteners from airplane wings, to strip and repaint aircraft, and to perform wire assembly. Total funding for these three projects is about $3 million per year.

The Naval Ocean Systems Center is currently exploring various military applications of robot and teleoperator systems. There are specific interests in teleoperated and robot submersibles, teleoperated and robot land vehicles, teleoperated lighter than air vehicles, underwater manipulators, stereo optic and acoustic vision, remote presence, autonomous robot knowledge representation and decision making and complex robot system specification and verification. These interests are distributed among six projects funded at a total of $650K per year.

The total government funding for robotics is about $10 million per year.

OVERSEAS RESEARCH

Overseas robotics efforts are considerably better funded.

Although exact figures are hard to obtain, most knowledgeable observers estimate that the Japanese are spending from three to ten times as much as the United States on robotics and related research. The Western Europeans are estimated to be spending from two to four times as much as the U. S. Certainly the corporate giants of Europe and Japan are heavily involved. Fiat, Renault, Olivetti, and Volkswagon have all developed their own robots, and many other European firms are marketing a wide variety of very sophisticated robots. In Japan, Kawasaki, Hitachi, Yasakawa, Fanuc, and Misubitshi all have major research laboratories and are aggressively marketing a wide variety of industrial robots. Fanuc has teamed up with Siemens of Germany to market a very competitive line of robots under the name General Numeric.

European and Japanese university efforts are heavily subsidized by the respective governments and university-industry collaboration is very close. Many university research laboratories are elaborately equipped with the most modern N/C machine tools and the best robots. Many of these machines are donated by private industry. Government support for salaries and overhead makes it possible for the universities in Europe and Japan to sustain large and coherent research programs. Even if the total U. S. effort were equivalent, the lack of U. S. centers of excellence supported on a consistent long term basis would put the U. S. at a serious disadvantage. The fact is, U. S. robotics research efforts are neither better funded nor better organized than those of our overseas trading partners. The Japanese have made the development of the automatic factory a high priority item of national policy. European research is heavily subsidized by the government funds. In both places robotics technology is treated as crucial to national economic development.

IMPLEMENTATION

In the United States at present, there are only about 3000 robots installed. That's less than the number of workers employed in a single factory in many companies. That's less than the graduating class of some high schools in this country. Today, there is a bigger market for toy robots than for real robots. So at least for the present, robots are having almost no effect one way or another on overall productivity in this country. Today, robots are being produced in the United States at the rate of about 1500 per year. Predictions are that this will probably grow to between 20,000 and 60,000 robots per year by the year 1990. In other words the production rate is growing at about a factor of 10 to 30 per decade. At that rate the U. S. will be lucky to have a million robots in operation before the year 2000. This means that unless there is some drastic change in the presently projected trends, there won't be enough robots in operation to have a significant impact on the overall productivity of the nation's economy before the turn of the century.

Of course, there will be some specific areas where the impact of robots will be large. In areas like automobile spot-welding, robots have already had some effect. By the mid 1980's there may be a significant effect on productivity in arc welding.

Arc welding is a hot, dirty, unpleasant job where the welder must wear heavy protective clothing and must work in the presence of a shower of hot sparks and choking smoke. Typically a human welder cannot keep his torch on the work more than 30% of the time. A robot welder, on the other hand can keep its torch on the work about 90% of the time. Thus, even though the robot cannot weld any faster than a human, it can turn out about 3 times as much work.

Unfortunately, present day robots cannot set up their own work. That requires a human assistant. So this reduces the productivity advantage. Also, the robot must be programmed to perform the welding task. Typically this takes much longer than would be required to actually perform a weld. Thus, unless the robot is used to perform many repetitions of the same welding task there is no productivity gain.

Of course, once robots become intelligent enough to assemble and set up their own work, productivity will improve. Once robots become clever enough to look at the job and figure out where to put the weld, productivity will improve even more. Eventually, welding robots will be sufficiently sophisticated to work from plans stored in computer memory and to correct errors which may occur during a job. Welding robots will then be able to work nights and weekends (four shifts per week) completely without human supervision. At that point productivity improvements over present methods of many hundreds of percent become possible. Unfortunately, we are a long way from that today. There are many difficult research and development problems that must be solved first. Unless the level of effort in software development is increased many fold, these improvements will not be realized for many years.

Let's look at another industry, the metal cutting industry, where robots are already being used to load and unload machine tools. This is a relatively simple task, so long as the parts are presented to the robot in a known position and orientation. During the 1980's, robot sensory and control capabilities will improve to the point where robots can find and load unoriented parts, or in some cases, even pick parts out of a bin filled with randomly oriented parts lying on top of each other. This may improve productivity by hundreds of percent because it will make it possible to install robots in many exisiting plants without major re-engineering of production methods. For example, in conventional N/C machine shops a single machinist could set up several machines which could then run for extended periods unattended. In some cases robot tended machines may run overnight and on weekends without human intervention.

By 1990 robots may begin to have a significant impact on mechanical assembly. There has been a great deal of research effort spent on robot assembly. Unfortunately, the results have not been spectacular--yet. On the one hand, robots cannot compete with classical so-called "hard automation" in assembly of mass produced parts. General purpose machines like robots are still too slow and too expensive to be economical for mass production assembly tasks. On the other hand, robots cannot yet compete with human assembly workers in small lot assembly. Humans are incredibly adaptable, dexterous, as well as fast, skilled, and relatively cheap compared to robots. A human has two hands and ten fingers with arms, and shoulders mounted on a mobile platform equipped with a total of 58 degrees of freedom. The human has a fantastically sophisticated vision system and can be programmed to perform a wide variety of tasks quite easily. Even in a relatively routine task such as the assembly of an automobile alternator (performed at the C.S. Draper Lab, Cambridge, MA), test results indicated that robot assembly would be only marginally effective economically even after every phase of the task had been optimized.

Nevertheless, progress is being made and will continue. Robot capabilities will gradually increase. Sensory systems will become more sophisticated and less expensive. The cost of computing hardware is dropping rapidly and steadily with no sign of bottoming out. Software costs are likely to be the major impediment to robot development for the foreseeable future, but even these are slowly yielding to the techniques of structured programming and high level languages.

Eventually, extremely fast, accurate, dexterous robots will be programmed using design graphics data bases which describe the shape of the parts to be made and the configuration of the assemblies to be constructed. Eventually, robots will be able to respond to a wide variety of sensory cues, to learn by experience and to acquire skills by self optimization. Such skills can then be transferred to other robots so that learning can be propagated rapidly throughout the robot labor force.

During the 1990's robots will probably enter the construction trades. Under the tutelage of a human mastercraftsman, apprentice robots will carry building materials, lift and position wall and floor panels, cut boards to size, and lay brick, block, and eventually stone. In the next century, labor intensive building techniques (using robot labor) may once again become practical. Homes, streets, bridges, gardens and fountains may be constructed of sculpted stone, quarried, cut, and assembled by robots. Eventually, robots will mine the seabed, and farm the surfaces of the oceans for food and fuel. And, of course, robots will play a major role in outer space, -- in the construction of large space structures, in space manufacturing, and in planetary exploration.

Sometime, perhaps around the turn of the century, robot technology will develop to the degree necessary to produce the totally automated factory. In such factories robots will perform most, if not all, of the operations that now require human skills. There will be totally automatic inventory and tool management, automatic machining, assembly, finishing, and inspection systems. Automatic factories will even be able to reproduce themselves. That is, automatic factories will make the components for other automatic factories.

Once this occurs, productivity improvements will propagate from generation to generation. Each generation of machines will produce machines less expensive and more sophisticated than themselves. This will bring about an exponential decline in the cost of robots and automatic factories which may equal the cost/performance record of the computer industry. For the past 30 years computing costs have spiraled downward by 20% per year. This, at least in part, is due to the fact that computers are used to design, construct, and test other computers. Once automatic factories begin to manufacture the components for automatic factories, the cost of manufacturing equipment will also fall exponentially. This, obviously, will reduce the cost of goods produced in the automatic factories. Eventually, products produced in automatic factories may cost only slightly more than the raw materials and energy from which they are made.

The long range potential of totally automated manufacturing is literally beyond our capacity to predict. It may change every aspect of industrial society. Automatic factories that can operate without human labor, and reproduce themselves, could lead to an entirely new era in the history of civilization.

Now, in the light of the unprecedented economic potential of robots, I suppose I should comment on why the implementation of this technology is proceeding so slowly.

First, at least in the U.S., funding for robotics R&D has been very modest. Every indication is that in the future, support will grow, but not dramatically. Certainly, there is nothing to suggest that a crash development program on the scale of the Manhattan Project or the Apollo Moon Program is imminent. Certainly, there are no plans for the federal government to launch such an effort and private investment funds are not likely to be committed on a massive scale because of the long time to pay back. Robotics is still a long term research topic. We are a long, long way from a sophisticated sensory interactive, intelligent, highly skilled, dexterous, economically feasible, and commercially manufacturable robot. Research in this area is long term, time consuming, and risky. Also, there is no certainty that inventions can be kept proprietary. There is therefore, no guarantee that the firms which make the investments can capture enough of the benefits to make the risk worthwhile.

Secondly, even after the research and development problems are solved, several decades and many hundreds of billions of dollars will be required to convert the present industrial base to robot technology. This enormous investment will severely tax available sources of capital. The transformation of the entire industrial plant of a country simply cannot be achieved except over an extended time period.

Thirdly, and perhaps most importantly, many voters question the desirability of rapid, massive deployment of robot technology. Despite the obvious benefits from productivity improvement, there would be serious social and economic adjustments necessary as a result of such a rapid productivity growth. Productivity improvement by its very nature reduces the amount of human labor needed to produce a given product. Thus, an obvious, but I believe incorrect conclusion is that a rapid increase in productivity would lead to unemployment. There is a wide spread perception that robots pose a threat to jobs. The fear is that if robots were introduced at the rate that is technologically possible, unemployment would become a serious problem.

However, widespread unemployment is not the inevitable result of rapid productivity growth. There is not a fixed amount of work! More work can always be created. All that is needed is a way to meet the payroll. Markets are not saturated. The purchasing power of consumers can always be increased at the same rate that more products flow out of the robot factories. At present, there is plenty of demand. The mere fact of inflation is prima facie evidence that consumer demand exceeds the ability of present production techniques and facilities to supply goods and services at constant prices. Work is easy to create. So is demand. What is hard to produce is goods and services that can be sold for a profit, at, (or below) the current market price.

Nevertheless, the average citizen is unconvinced that advanced automation would necessarily put increased spending power into his or her pocketbook. The question is -- If the robots have most of the jobs, how will average people get their income? In order for most people to be convinced that robots are going to bring more benefits than problems it will be necessary to demonstrate that a variety of alternative income producing occupations will be created to fill the void left by those jobs which are taken over by robots. Fortunately, this is not difficlt to do.

Perhaps, the most obvious source of new jobs is in the industries which must be created in order to convert to a robot based economy. Certainly if robots are to be manufactured in large enough quantities to make a significant impact on the existing industrial system, entirely new robot manufacturing, sales, and service industries will emerge and millions of exciting new jobs will be created. A typical industrial robot costs from $30,000 to $80,000 and sometimes more by the time it is installed and operating. This means that every robot installed creates from 2 to 4 person-years of work somewhere in the economy. The robot market is presently growing at about 35% per year, which means it doubles about every 3 years. As long as this growth rate continues, robot production will add jobs to the economy about as fast as robot installation takes them away.

It will be many years, perhaps many decades, before robots can design, manufacture, market, install, program, and repair themselves with little or no human intervention. In the meantime, the manufacture and servicing of robots will produce an enormous demand for mechanical engineers, technicians, computer programmers, electronic designers, robot installation and repair persons. New robot companies will require secretaries, sales persons, accountants, and business managers. It seems likely that the robot industry will eventually employ at least as many people as the computer and automobile industries do today.

Converting the world's existing industrial plants from manual to robot labor will require many decades and will cost as much as the total existing stock of industrial wealth. This is a Herculean task which will provide employment to millions of workers for several generations. For a country like the United States which has a strong technological base, the world market in robots could easily create twice as many jobs in robot production as were lost to robot labor. Needless to say, the export of robot systems (as well as products made by them) could have a strong positive effect on the balance of trade and the strength of the dollar on the international market.

In general, industries that use the most efficient production techniques grow and prosper, and hire more workers. Markets for their products expand and they diversify into new product lines. Workers displaced by automation are simply transferred into new growth areas or retrained for different occupations. It is in the industries that fall behind in productivity that job layoffs are prevalent. Inefficient industries lose market-share to competitors, shrink, and eventually die. Thus, the biggest threat to jobs is not in industries that adopt the latest robot technology, but in those which do not.

For example, there are almost one-half million jobless workers today in the American automobile industry. This is not because of a couple thousand robots. It is because of the energy crisis and because of foreign competition. U. S. auto workers are suffering unemployment more because of robots in Japan than because of robots in Detroit. If America continues the present low rate of productivity growth, we cannot help but have even greater unemployment. Foreign trading partners are modernizing at a rapid rate. If we do not innovate, our products cannot compete, and our workers will find their jobs being taken away by foreign competition.

Improving productivity is not easy. It requires research, development, education, capital investment, and incentives to do better. The new technology of advanced automation is not a quick fix. It is a long range solution. Robots have much promise but a long way to go. We are only beginning to understand some of the technical problems. We are many years, perhaps several decades from making truly intelligent, highly skilled robots. But technical solutions will come. It is only a matter of time, money, and intellectual resources. The real question is whether we can evolve a society in which robots will complement, not compete with, humans for their livelihood. If this problem can be solved, then the prospects for the future may be very bright indeed. Robots and automatic factories have the potential to increase productivity virtually without limit. This potential, if brought to reality, could create a material abundance and standard of living which far exceeds the horizon of today's expectations. Over the next two centuries the technology of robotics and advanced automation could make everyone rich. Robots someday could provide the economic foundation for an "everypersons' aristocracy." However, this will require that we find a way to make them work for us, and not in competition with us. To protect the human worker's livelihood in the coming decades there are several steps which can and should be taken.

First, we must provide retraining for workers displaced by robots for new and better occupations.

Second, (after a decade or so when robots begin to make a significant impact on productivity) we can decrease the workweek. It is nowhere written in stone that humans must work 40 hours per week. As robots take over more and more work, humans can improve their work environment and decrease their work periods to 30, 20, or even 10 hours per week. Education and leisure activities can be increased virtually without limit. Eventually, all "work" could be voluntary.

However, in order to achieve this we will need to explore a wide variety of mechanisms for broadening our ownership of robots and automatic factories. Employee stock ownership plans, individual robot-owner entrepreneurs, and even semi-public mutual fund ownership plans might be developed in the future. If everyone could own the equivalent of one or two robots, everyone would be financially independent, regardless of whether they were employed or not.

Finally, in the next few years and decades, we must recognize that it is premature to worry about insufficient work to go around. There is virtually an unlimited amount of work that needs to be done in eliminating poverty, hunger, and disease, not only in America, but throughout the world. We need to develop renewable energy resources, clean up the environment, rebuild our cities, exploit the oceans, explore the planets, and colonize outer space. The new age of robotics will open many new possibilities. What we humans can do in the future is limited only by our imagination to see the opportunities and our courage to act out our beliefs.

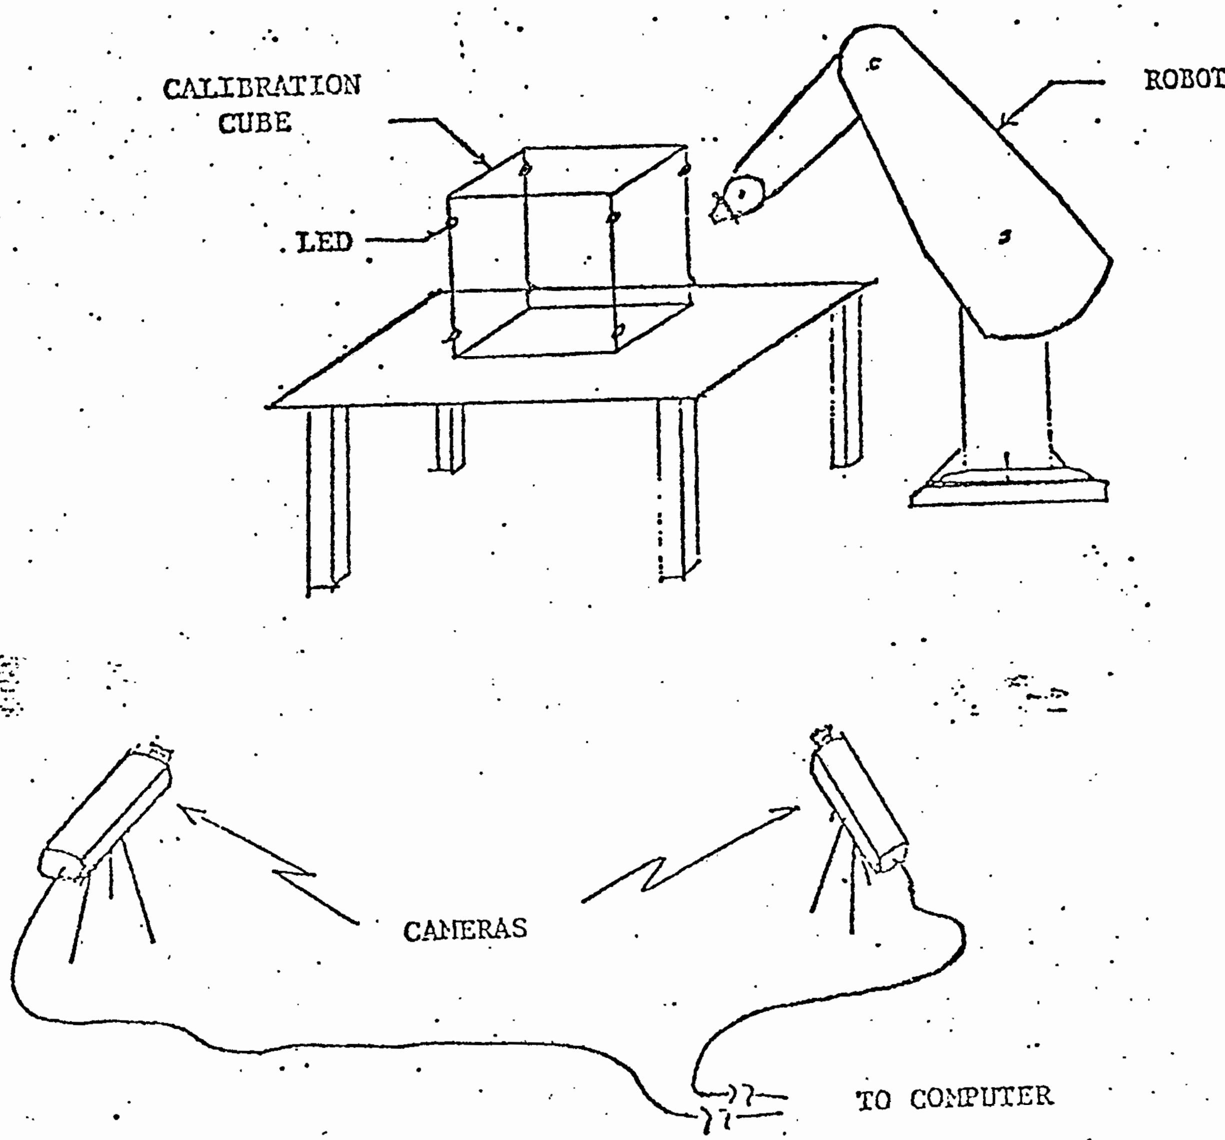

Figure 1. Remote, in situ robot trajectory calibration system. Each of the two cameras can measure the x and y position of light-emitting-diodes (LED's). Initially, a calibration cube with a set of LEDs at known points is used to compute the positions and viewing angles of the two cameras. Then the two cameras can track a LED on the robot so as to determine the 3-dimensional position accuracy of the robot over its working volume.

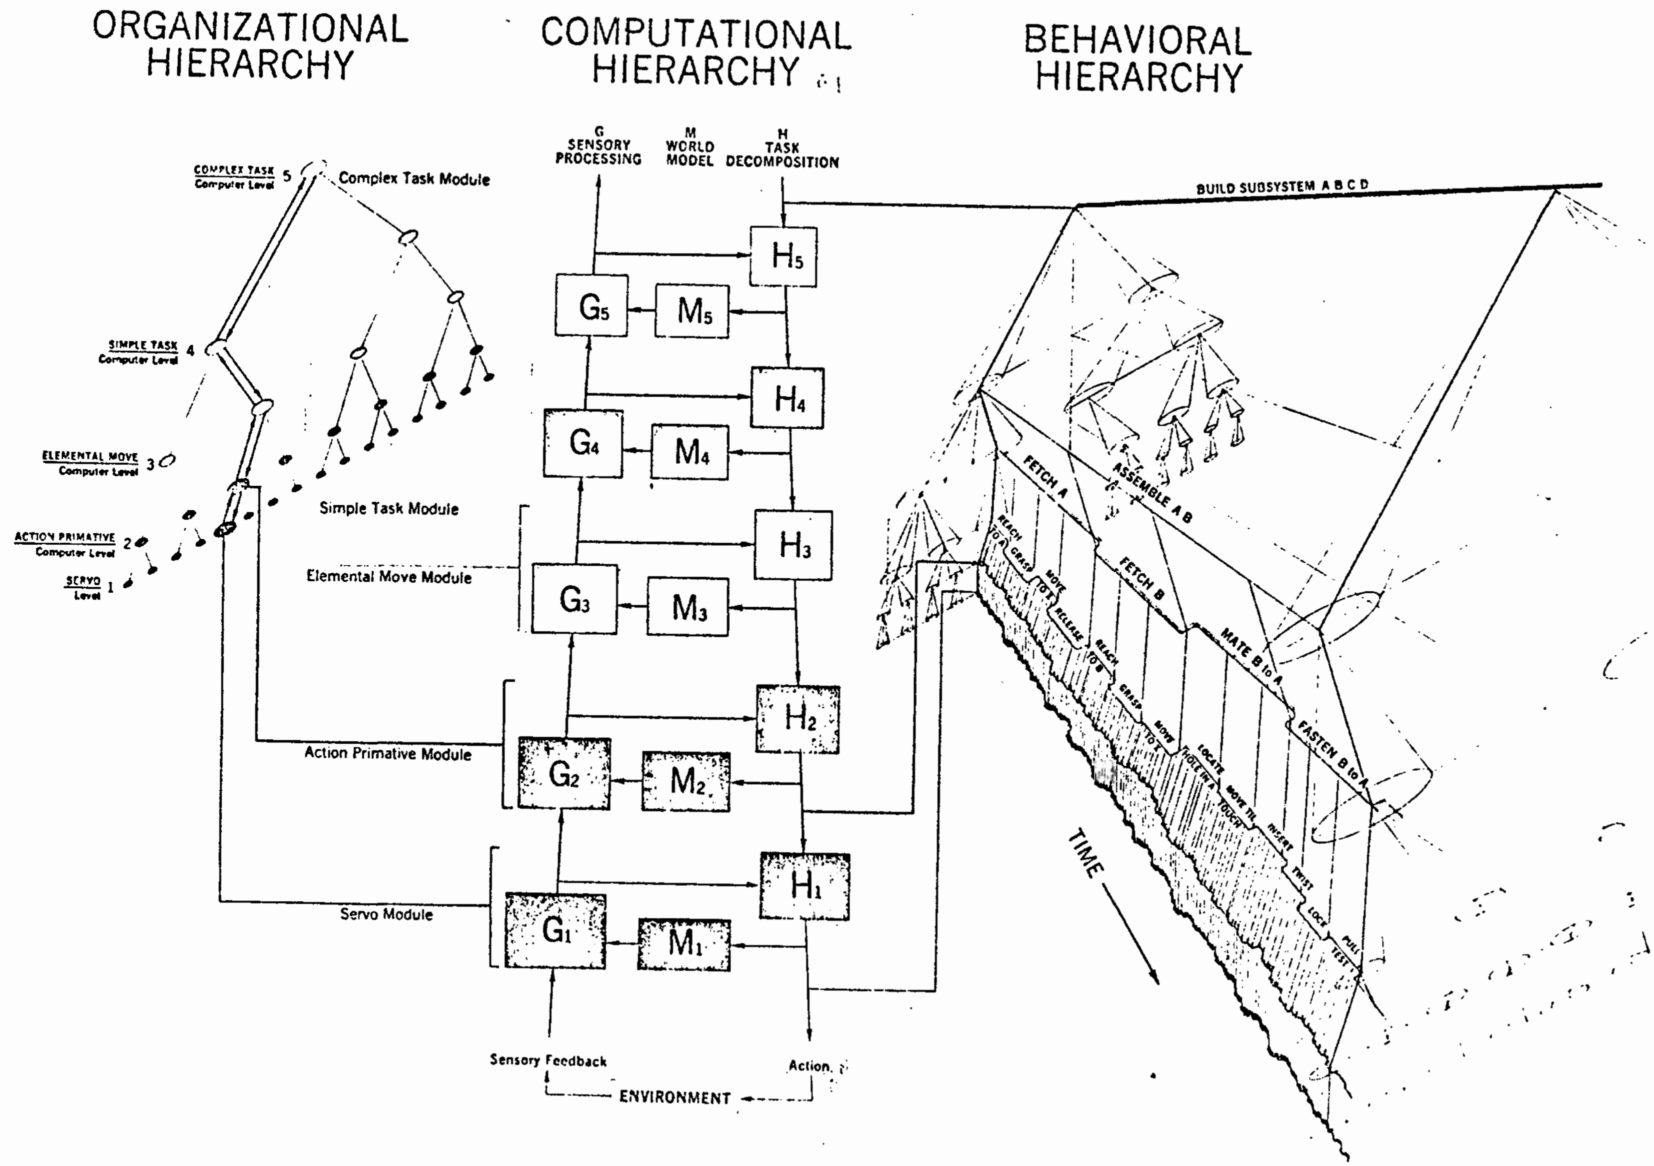
ORGANIZATIONAL HIERARCHY
COMPUTATIONAL HIERARCHY
BEHAVIORAL HIERARCHY
COMPLEX TASK Computer Level 5
Complex Task Module
SIMPLE TASK Computer Level 4
ELEMENTAL MOVE Computer Level 3
ACTION PRIMATIVE Computer Level 2
SERVO Level 1
Simple Task Module
Elemental Move Module
Action Primative Module
Servo Module
G SENSORY PROCESSING
M WORLD MODEL
H TASK DECOMPOSITION
H5
H4
H3
H2
H1
M5
M4
M3
M2
M1
G5
G4
G3
G2
G1
Sensory Feedback
Action
ENVIRONMENT
BUILD SUBSYSTEM A B C D
ASSEMBLE A B
FETCH A
FETCH B
MATE B to A
FASTEN B to A
TIME

Figure 2

Explanation of Figure 2.

The command and control structure for successful organizations of great complexity is invariably hierarchical, wherein goals, or tasks, selected at the highest level are decomposed into sequences of subtasks which are passed to one or more operational units at the next lower level in the hierarchy. Each of these lower level units decomposes its input command in the context of feedback information obtained from other units at the same or lower levels, or from the external environment, and issues sequences of sub-subtasks to a set of subordinates at the next lower level. This same procedure is repeated at each successive hierarchical level until at the bottom of the hierarchy there is generated a set of sequences of primitive actions which drive individual actuators such as motors, servo valves, hydraulic pistons, or individual muscles. This basic scheme can be seen in the organizational hierarchy on the left of Figure 2.

A single chain of command through the organizational hierarchy on the left is shown as the computational hierarchy in the center of Figure 2. This computational hierarchy consists of three parallel hierarchies: a task decomposition hierarchy, a sensory processing hierarchy, and a world model hierarchy. The sensory processing hierarchy consists of a series of computational units, each of which extract the particular features and information patterns needed by the task decomposition unit at that level. Feedback from the sensory processing hierarchy enters each level of the task decomposition hierarchy. This feedback information comes from the same or lower levels of the hierarchy or from the external environment. It is used by the modules in the task decomposition hierarchy to sequence their outputs and to modify their decomposition function so as to accomplish the higher level goal in spite of perturbations and unexpected events in the environment.

The world model hierarchy consists of a set of knowledge bases that generate expectations against which the sensory processing modules can compare the observed sensory data stream. Expectations are based on stored information which is accessed by the task being executed at any particular time. The sensory processing units can use this information to select the particular processing algorithms that are appropriate to the expected sensory data and can inform the task decomposition units of whatever differences, or errors, exist between the observed and expected data. The task decomposition unit can then respond, either by altering the action so as to bring the observed sensory data into correspondence with the expectation, or by altering the input to the world model so as to bring the expectation into correspondence with the observation.

Each computational unit in the task decomposition, sensory processing, and world modeling hierarchies can be represented as a finite-state machine. At each time increment, each unit reads its input and based on its present internal state computes an output with a very short time delay.

If the output of each unit in the task decomposition hierarchy is described as a vector, and plotted versus time in a vector space, a behavioral hierarchy such as is shown on the right side of Figure 2 results. In this illustration a high level goal, or task, (BUILD SUBASSEMBLY ABCD) is input to the highest level in a robot control hierarchy. The H5 task decomposition unit breaks this task down into a series of subtasks, of which (ASSEMBLE AB) is the first. This "complex" subtask command is then sent to the H4 task decomposition unit. H4 decomposes this "complex" subtask into a sequence of "simple" subtasks (FETCH A), (FETCH B), (MATE B to A), FASTEN B to A). The H3 unit, subsequently decomposes each of the "simple" subtasks into a string of "elemental moves" of the form (REACH TO A), (GRASP), (MOVE to X), (REALEASE), etc. The H2 decomposition unit then computes a string of trajectory segments in a coordinate system fixed in the work space, or in the robot hand, or in the work piece itself. These trajectory segments may include acceleration, velocity, and deceleration profiles for the robot motion. In H1, each of these trajectory segments are transformed into joint angle movements and the joint actuators are servoed to execute the commanded motions.

At each level, the G units select the appropriate feedback information needed by the H modules in the task decomposition hierarchy. The M units generate predictions, or expected values, of the sensory data based on the stored knowledge about the environment in the context of the task being executed.

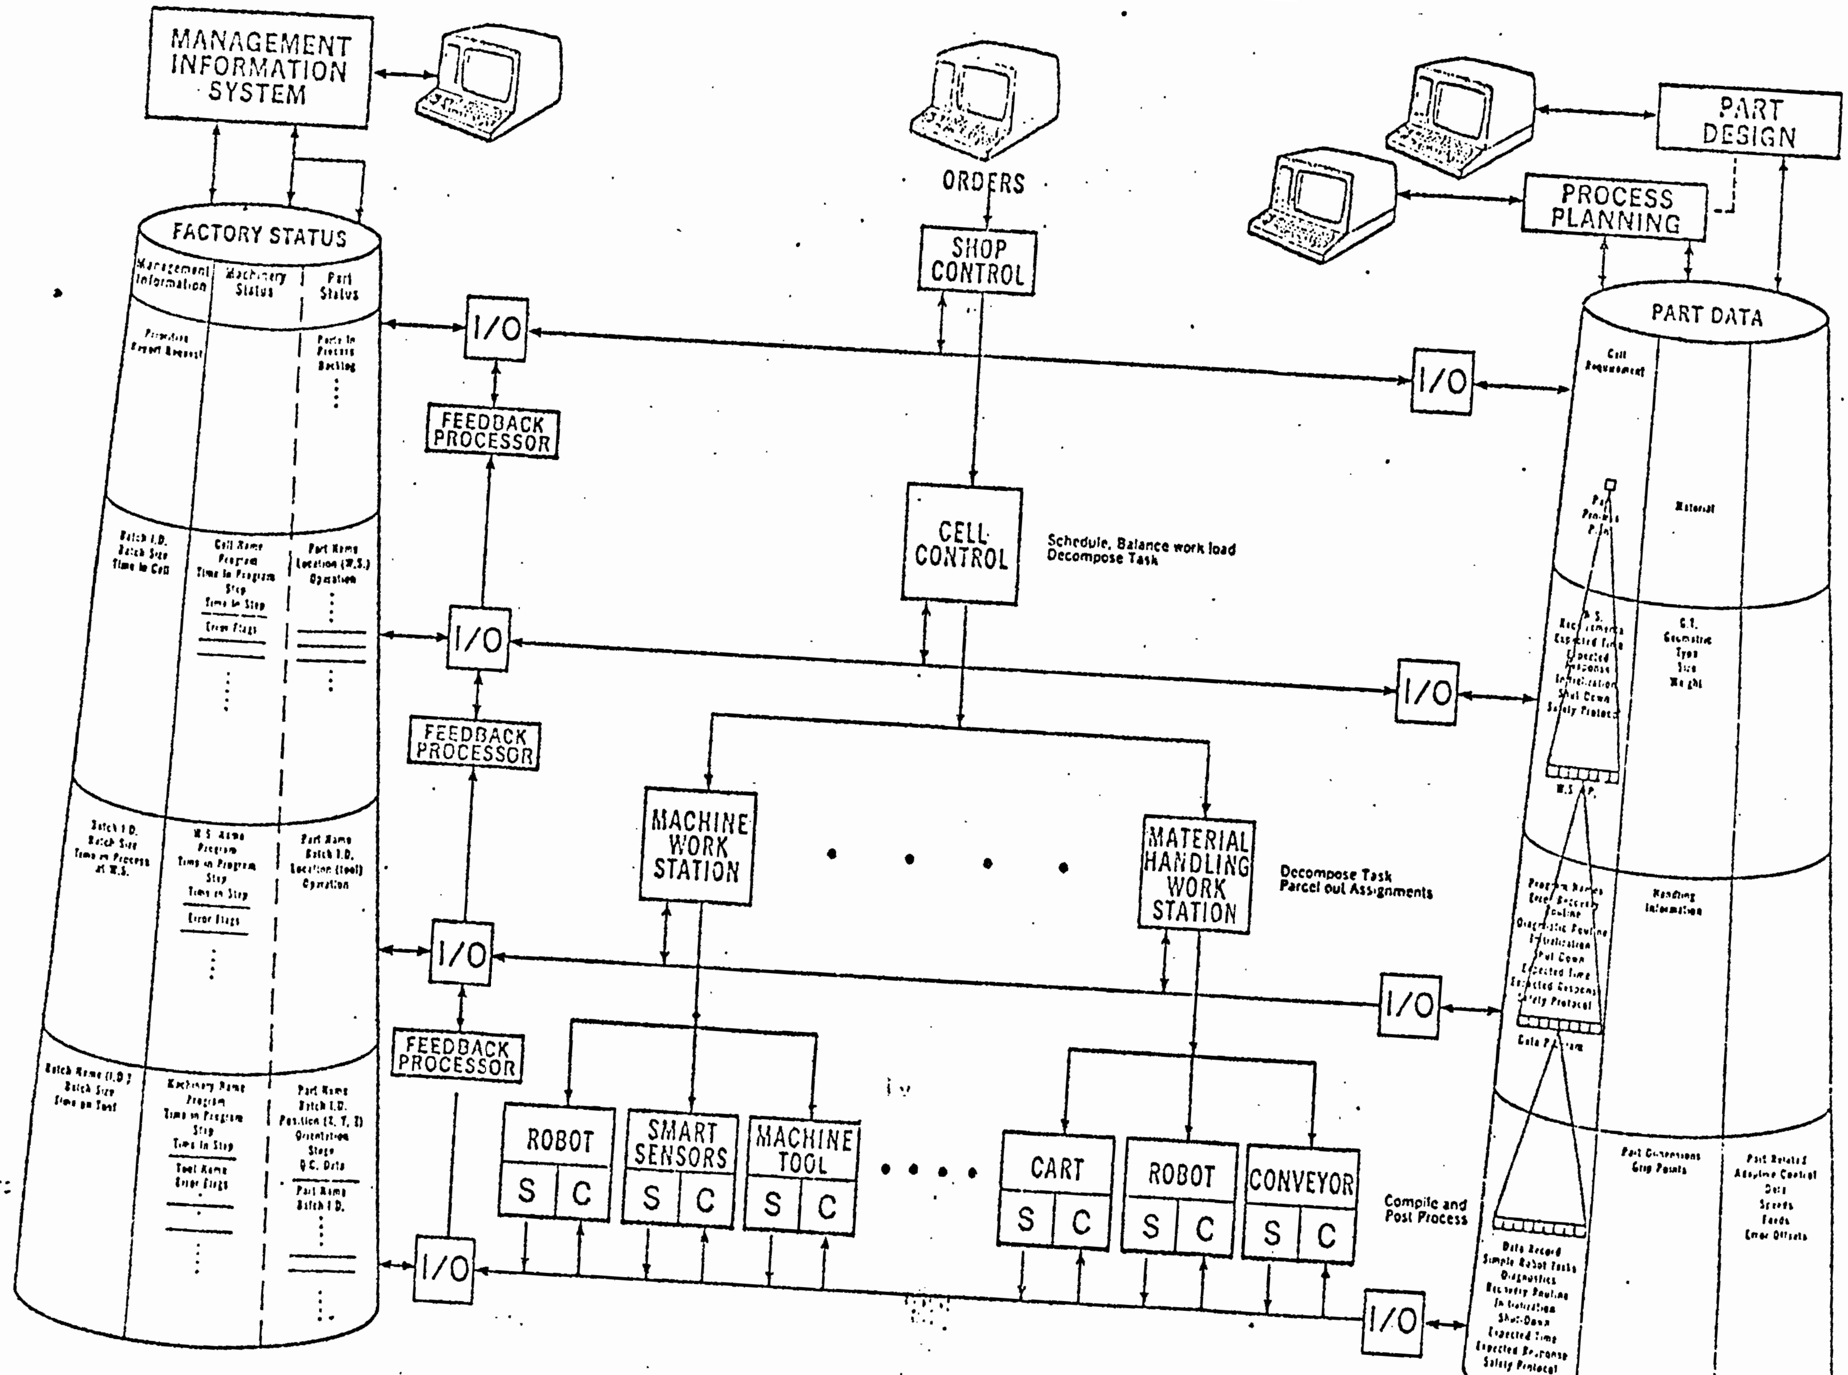
MANAGEMENT INFORMATION SYSTEM
FACTORY STATUS
Management Information
Machinery Status
Part Status
ORDERS
SHOP CONTROL
CELL CONTROL
Schedule, Balance work load
Decompose Task
PART DESIGN
PROCESS PLANNING
PART DATA
I/O
FEEDBACK PROCESSOR
MACHINE WORK STATION
MATERIAL HANDLING WORK STATION
Decompose Task
Parcel out Assignments
ROBOT
SMART SENSORS
MACHINE TOOL
CART
CONVEYOR
S
C
Compile and Post Process

Figure 3

Explanation of Figure 3.

The computing architecture shown in Figure 3 is intended as a generic system that can be applied to a wide variety of automatic manufacturing facilities and can be extended to much larger applications. The basic structure is hierarchical, with the computational load distributed evenly over the various computational units at the various different levels of the hierarchy. At the lowest level in this hierarchy are the individual robots, N/C machining centers, smart sensors, robot carts, conveyors, and automatic storage systems, each of which may have its own internal hierarchical control system. These individual machines are organized into work stations under the control of a work station control unit. Several work station control units are organized under, and receive input commands from a cell control unit. Several cell control units may be organized under and receive input commands from a shop control unit, etc. This hierarchical structure can be extended to as many levels with as many modules per level as are necessary, depending on the complexity of the factory.

On the right side of Figure 3 is shown a data base which contains the part programs for the machine tools, the part handling programs for the robots, the materials requirements, dimensions, and tolerances derived from the part design data base, and the algorithms and process plans required for routing, scheduling, tooling, and fixturing. This data is generated by a Computer-Aided-Design (CAD) system and a Computer-Aided-Process-Planning (CAPP) system. This data base is hierarchically structured so that the information required at the different hierarchical levels is readily available when needed.

On the left is a second data base which contains the current status of the factory. Each part in process in the factory has a file in this data base which contains information as to what is the position and orientation of that part, its stage of completion, the batch of parts that it is with, and quality control information. This data base is also hierarchically structured. At the lowest level, the position of each part is referenced to a particular tray or table top. At the next higher level, the work station, the position of each part refers to which tray the part is in. At the cell level, position refers to which work station the part is in. The feedback processors on the left scan each level of the data base and extract the information of interest to the next higher level. A management information system makes it possible to query this data base at any level and determine the status of any part or job in the shop. It can also set or alter priorities on various jobs.