Undersea Teleoperators and Intelligent Autonomous Vehicles

Edited by
Norman Doelling
Elizabeth T. Harding

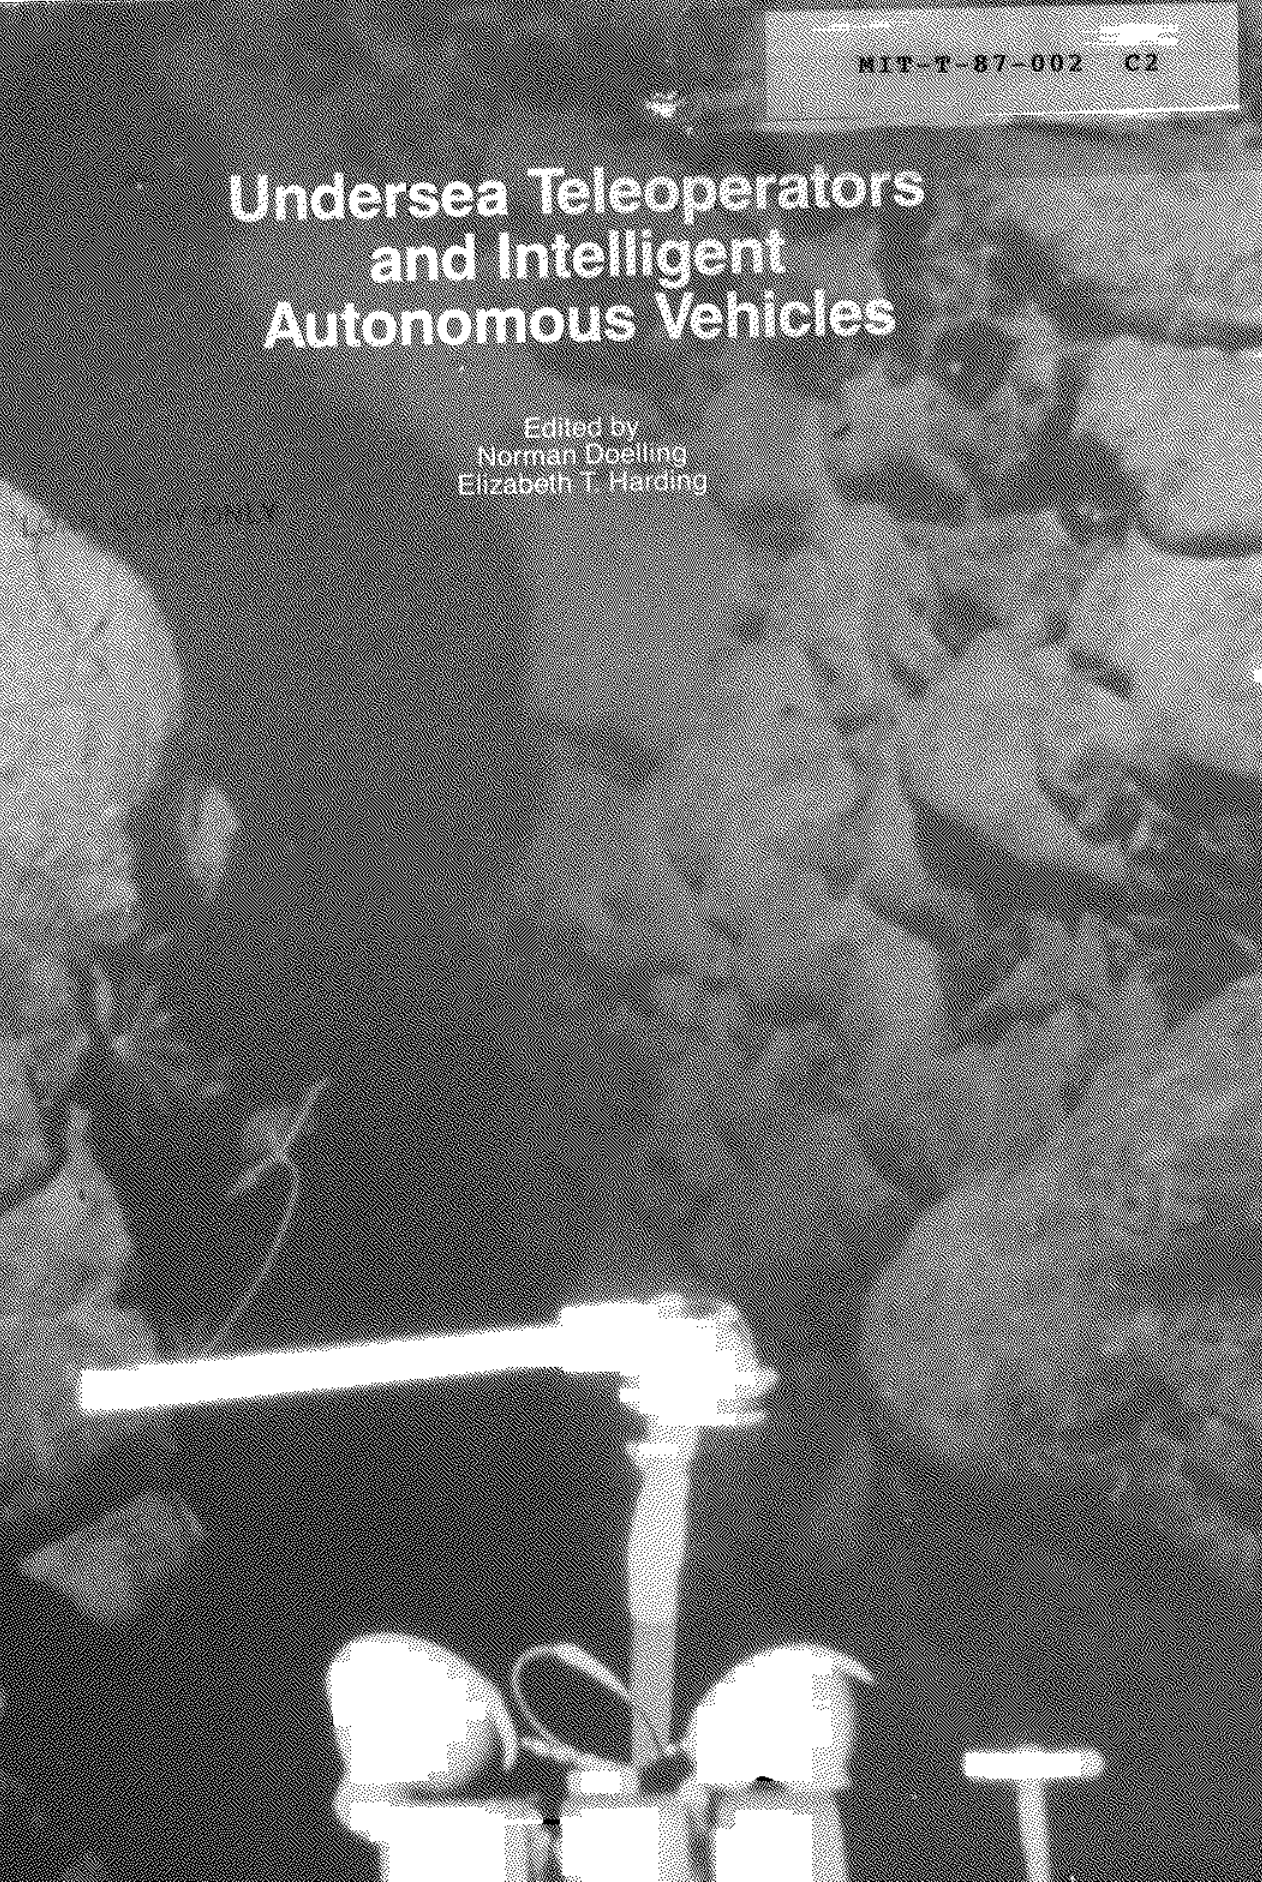

Errata

Undersea Teleoperators
and Intelligent
Autonomous Vehicles

Two important errors of omission have been brought to our attention.

The support of the Ocean Sciences Division of the National Science Foundation was inadvertently omitted from the Acknowledgements section. Our apologies to Larry Clark.

Also, Figure 2, originally to appear on page 73, was omitted. Figure 2 is reproduced below.

If any further errors are found, please call or write me.

Norman Doelling

April 21, 1987

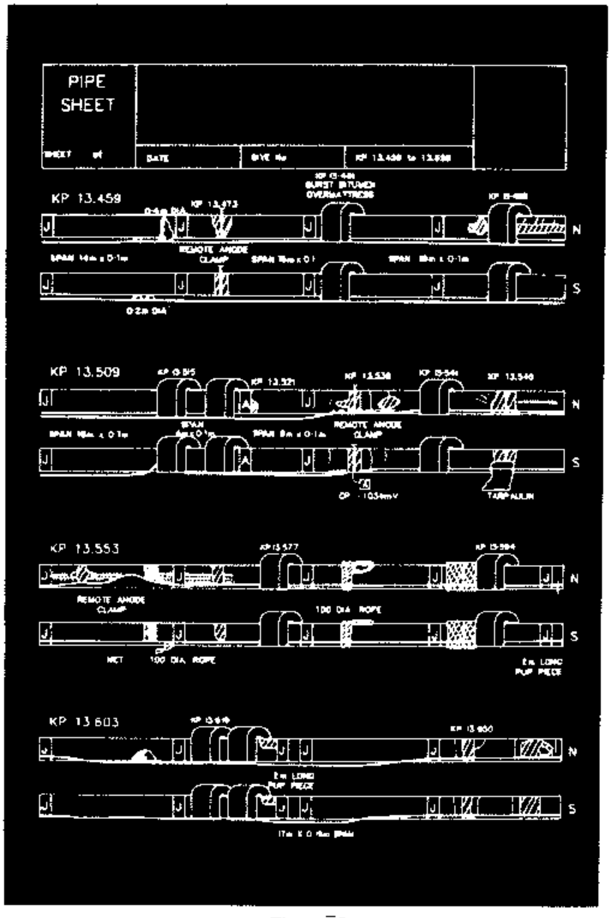

Figure 2
Pipe Detail

Undersea Teleoperators and Intelligent Autonomous Vehicles

MIT Sea Grant College Program
Cambridge, Massachusetts 02139

MITSG 87-1
January 1987

This volume, *Undersea Teleoperators and Intelligent Autonomous Vehicles*, is the proceedings for the 1986 MIT Department of Ocean Engineering Robert Bruce Wallace Lecture and the 1986 Annual MIT Sea Grant College Program Lecture and Seminar.

Copies are available through
MIT Sea Grant Information Center
Massachusetts Institute of Technology
292 Main Street, E38-306
Cambridge, Massachusetts 02139

Cataloging in Publication Data
Main Entry under title:
Undersea Teleoperators and Intelligent Autonomous Vehicles
1. Vehicles, remotely piloted—Congresses
2. Oceanographic Submersibles—Congresses
3. Manipulators (mechanism)—Congresses
4. Robotics—Congresses
I. Harding, Elizabeth T., editor
II. Doelling, Norman, editor
III. Series

Contents

Extending the Reach of Human Operators

Teleoperation for Marine Biology and Oceanography Research

The Sixth Robert Bruce Wallace Lecture

The Fourteenth MIT Sea Grant College Program Lecture

Mobile Robots in Unstructured Environments

Emerging Manipulator Technology

Foreword

Not many years ago, robot employees seemed a science fiction writer's dream of the future. However, quantum leaps in microprocessor technology and engineering have made robots a common sight in factories around the world. Theirs are the jobs which are repetitive and boring, requiring no knowledge, no judgment. But as the rapid evolution of microprocessor technology, control theory and artificial intelligence continue their prospects are changing.

In this volume we discuss the development of robots which can operate in an extraordinarily complex environment, the ocean. The research challenges are considerable. To review what has been accomplished and identify what remains to be done to create teleoperators and intelligent autonomous vehicles which can operate below the sea's surface, we have asked a group of AI experts, marine scientists, robotic engineers, and underwater vehicle designers to present their views and work. This book, based on papers presented at a conference held at MIT in October 1986, presents a multidisciplinary look at the evolving fields of teleoperation, robotics and artificial intelligence, as they apply to underwater systems.

In the first 150 pages, the authors present an historical perspective of the current state-of-the-art in teleoperation and the evolution of underwater remotely operated systems. Applications to marine biology, offshore oil and even space flight are highlighted.

The last section of the book explores the requirements of future systems. It would seem from the papers that the continuing revolution in computer technologies, control systems and artificial intelligence herald exciting futures for these undersea robots. There is no consensus on the form, shape or capabilities of future systems, but the possibilities are all interesting, exciting, thought provoking. In fact, one author—Hans Moravec—suggests in a very few decades, the complexity of computer hardware will imply potential intelligence approaching that of lower animals and perhaps even humans.

Both the conference and these proceedings were planned to encourage debate and to inspire new alliances between people and groups who are now taking different paths toward the same destination—the advancement of robot systems which can explore new frontiers or supplant human beings in work environments which are life threatening. The ocean is such a place—deep, dark, mysterious, and dangerous, but with enormous, unexplored potential. Robot vehicles can be a great tool in a continuing human quest to explore and understand the sea and its myriad resources.

Norman Doelling
Elizabeth T. Harding
Editors

Acknowledgments

The papers in this volume were presented at a symposium, "Undersea Teleoperators and Intelligent Autonomous Vehicles," in October, 1986 at the Massachusetts Institute of Technology Bartos Theatre in the Jerome Wiesner Building. The Sixth Annual Robert Bruce Lecture and the Fourteenth Annual Sea Grant Lecture were presented during the symposium. The Lectureship has been made possible by a gift from Mr. and Mrs. A.H. Chatfield in memory of Mrs. Chatfield's father, Robert Bruce Wallace MIT '98. The Henry L. and Grace Doherty Foundation have generously provided a permanent endowment for the Sea Grant Lecture.

Additional support was provided by the MIT Department of Ocean Engineering, the MIT Department of Mechanical Engineering, the MIT Sea Grant College Program, Martin Marietta Aerospace, Baltimore, Hydro Products, A Division of Raytheon, National Bureau of Standards, Robot Systems Division, Naval Explosive Ordnance Disposal Technology Center, Naval Surface Weapons Center, and NOAA's Office of Undersea Research.

The organization of the conference and the production of these papers would not have been possible without the considerable help and efficiency of the MIT Sea Grant Program support staff, particularly Ms. Re Quinn. Francis Ogilvie, Head of the MIT Ocean Engineering Department and his administrative secretary, Mary Kreuz made important contributions to the success of the meeting.

Several others who served on the Program Advisory Committee should be acknowledged for their thoughtful suggestions and criticisms. They were, Richard Blidberg, Marine Systems Engineering Laboratory, University of New Hampshire, Eric Jackson, International Submarine Engineering, Ltd., and Michael Triantafyllou, MIT Department of Ocean Engineering.

Contributors

David L. Akin
Assistant Professor
Space Systems Laboratory
Department of Aeronautics and Astronautics
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

Rodney A. Brooks
Assistant Professor
Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

Andrew M. Clark
Head, Mechanical Engineering Department
Harbor Branch Oceanographic
Institution, Incorporated
Fort Pierce, Florida 33450

Robert W. Corell
Director
Sea Grant College Program
University of New Hampshire
Durham, New Hampshire 03824

F. Richard Frisbie
Senior Vice President
Ocean Systems Engineering
Houston, Texas 77084

W. Eric L. Grimson
Assistant Professor
Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

Neville Hogan
Associate Professor
Department of Mechanical Engineering
Laboratory for Manufacturing and Productivity
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

Stephen C. Jacobsen
Professor
Center for Engineering Design
University of Utah
Salt Lake City, Utah 84112

Samuel E. Landsberger
Research Assistant
Department of Mechanical Engineering
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

James R. McFarlane
President
International Submarine Engineering Ltd.
British Columbia, Canada V3H 1X1

Hans P. Moravec
Senior Research Scientist
The Robotics Institute
Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213

Tyler Schilling
Director
Schilling Development, Inc.
Davis, California 95616

Kenneth P. Sebens
Associate Professor
Marine Science Center
Northeastern University
Nahant, Massachusetts 01908

Thomas B. Sheridan
Professor
Man-Machine Systems Laboratory
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

Jean-Jacques E. Slotine
Assistant Professor
Department of Mechanical Engineering
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

Dana R. Yoerger
Assistant Scientist
Department of Ocean Engineering
Woods Hole Oceanographic Institution
Woods Hole, Massachusetts 02543

Extending the Reach of Human Operators

Teleoperators, Telepresence and Telerobots

Thomas B. Sheridan
Professor
Man-Machine Systems Laboratory
Massachusetts Institute of Technology

ABSTRACT

By now it is well recognized that remotely operated vehicles (ROVs) are replacing both free swimming human divers and manned submersibles in a variety of deep ocean inspection and work tasks. At the same time NASA is designing for their Space Station a new generation of teleoperators (devices for doing work at a distance from the human operator's eyes and hands e.g., ROVs which include at last one manipulator arm). This paper outlines some relevant history and concepts of teleoperation and discusses some current trends in the development of this technology, emphasizing in particular the relation between human operator and teleoperator. A number of examples are described of teleoperator research in the MIT Man-Machine Systems Laboratory over the last twenty years (with more extensive references to the latter than would constitute an even-handed literature review).

A TINY HISTORY OF TELEOPERATION

As with any new technology there is a richness of precedent. People have manipulated long sticks and ropes, hand tools and weapons to "extend their reach" to do work at a distance since long before recorded history. Over time these hand-extensions became specialized to mix, to sew, to grab, to hammer or to hit golf balls. They acquired separate power sources and became plows and earth moving machines and drills and printing presses and mechanical hands reaching into toxic or otherwise hazardous environments. Then they were augmented by sensor extensions - radio and television. Now the artificial sensors are being accompanied by computer brains, and the "hand tools" are becoming "subordinate intelligent and active entities". But that's getting ahead of the story......

Today the term "teleoperator" implies a relatively general purpose mechanical manipulator having at least three degrees-of-freedom and mounted on a base which itself may have some mobility, all remotely controlled by a human operator. There may or may not be some form of remote sensing such as video. Table 1 illustrates the progression in degrees-of-automation from the unencumbered human diver/operator using his own eyes and hands directly to the purely autonomous underwater inspection/work vehicle, with examples from undersea and space, and showing comparisons with respect to how observation, vehicle control and manipulation control are done.

Today's teleoperator technologists would probably judge Raymond Goertz, director of the AEC's Argonne (Illinois) Laboratory's Remote Handling Project during the 1940's and 1950's, as the person most responsible for the modern teleoperator (Johnson and Corliss, 1967). It may be testimony to Goertz's skill or to others' lack of imagination that in many respects telemanipulators have changed relatively little since Goertz's 1950 prototype E2 manipulator (which currently resides in our laboratory). But in many other respects teleoperators are changing now, very rapidly.

Another giant in the field was Jean Vertut, for many years director of the remote handling developments for the French CEA (Commission Energie Atomique) near Paris and also developer of some exotic undersea ROVs such as ERIC (Vertut and Coiffet, 1984). Perhaps the most exotic of the early teleoperators were Ralph Mosher's (General Electric's) 1954 HANDYMAN, having ten servoed joints on each of two arm-hand combinations, and his 1964 quadruped, a huge walking truck servo-controlled by harnesses attached to the four limbs of the operator (usually Mosher) who rode inside.

The first undersea ROVs (that term seems to have stuck for undersea teleoperators) emerged concurrently with the first space teleoperators. In 1963 Scripps Institution of Oceanography mounted a manipulator arm on an Army tank and drove it into the sea. In 1967 the SURVEYOR lunar landing

GENERAL CONCEPT OR TERM	UNDERSEA MODE	SPACE MODE	OBSERVE ENVIRONMENT	CONTROL VEHICLE	MANIPULATE ENVIRONMENT
DIRECT HUMAN CONTROL	diver	astronaut in EVA	direct	direct	direct
MANNED VEHICLE	manned submersible	manned spacecraft	direct plus video, force reflect	in-the-loop remote	in-the-loop remote
TELEOPERATOR (WITH ASPIRATIONS OF TELEPRESENCE)	conventional ROV	in-the-loop teleoperator	video, force reflect	in-the-loop	in-the-loop
SUPERVISORY CONTROLLED TELEROBOT	supervisory controlled ROV	space telerobot	computer mediated artificial sensors	human supervised computer	human supervised computer
AUTONOMOUS ROBOT	autonomous underwater vehicle	autonomous space vehicle	artificial sensing and intelligence	automatic computer control	automatic computer control

Table 1. Degrees of Automation of Remote Environment Work Systems

spacecraft had a simple scoop for shoveling material from the surface of the moon. Because of the three second time delay (round trip time to send a radio message from the earth to the moon and get confirmation information back) it was necessary to teleoperate SURVEYOR's manipulator in "move-and-wait" fashion. This had been shown to be necessary by earlier laboratory research (Ferrell and Sheridan, 1967: Black, 1970), meaning the operator could at best make a small "blind" movement of the manipulator (as though with his eyes closed), wait three seconds for feedback, then make another small "blind" move, wait for feedback, and so on. It is interesting to note that as undersea acoustic telemetry links get longer, similar time delays and similar control problems result for ROVs controlled through such links (one second round trip per 2500 feet depth).

In 1966 the Navy CURV teleoperator proved itself by successfully fishing an H-bomb off a precarious deepwater perch near Polomares, Spain. By 1979 commercial "roving eye-ball" (inspection) ROVs were being used commonly in off-shore oil operations, with a dozen or more companies participating (Busby, 1979). By this time the USSR was also well along with their ROV developments (Yastrebov and Stefanov, 1978). Soon afterwards very large special purpose teleoperators were laying cable and performing other exacting tasks, and a few years later, at the other end of the size spectrum, the tiny "Mini-rover" and its Japanese counterpart had made their debuts. Now the encumberances of the long, thick tethers are being replaced by optical fibers and even acoustic links, though these in turn present new technological challenges.

TELESENSORY FEEDBACK AND TELEPRESENCE

"Telepresence", a current buzzword in space technology, will surely come to be applied to marine teleoperator systems. It means visual, kinesthetic, tactile or other sensory feedback from the teleoperator to the human operator that is sufficient and properly displayed such that the human feels that he is present at the remote site, and that the teleoperator is an extension of his own body. In Table 1 "telepresence" is shown to apply equally well to extending the human operator's senses in the manned submersible, the conventional ROV, or the more robotized ROV (to be discussed further below). Though the term is new, the idea is traceable to the AEC Argonne Laboratory in the late 1940's (Johnsen and Corliss, 1967). Some salient problems of achieving (or moving closer to) the ideal of telepresence are:

(1) Television. Good progress has been made in developing low-light-level video cameras which do not require lots of energy-wasting artificial light (much of which is

bounced off of suspended particles and back into the camera). Also, it is now easy to digitize video images for real-time image enhancement, pattern recognition or superposition of computer-generated graphics on the screen. One application of the former is particularly appropriate when a necessarily low-bandwidth acoustic link is employed between the surface and an ROV. In this technique the video image is digitized and then adjusted according to operator specifications to have particular frames-per-second, pixels-resolution-per-frame and bits-of-gray-scale-per-pixel, where the product of these three is bits-per-second. It has been shown that when bits-per-second are severely constrained the operator would much prefer to adjust this tradeoff to suit his immediate task than to be stuck with arbitrary settings (Ranadive, 1979; Deghuee, 1980).

One particurly effective means of providing telepresence is to let the operator's side-to-side and up-down head motion drive the pan and tilt of the remote video camera in correspondence, and then mount a small video monitor directly to the operator's head. Then, wherever he moves his head he will see what he would see were he physically present at the site of the remote camera. Goertz, W.Bradley and others pioneered this technique twenty years ago (Johnson and Corliss, 1967) but its promise has hardly been exploited up to now.

(2) Teleproprioception. This term, coined from the Greek, means sensing where the telerobot's limbs are in space relative to its head (eyes) and body. Because the human operator naturally identifies the teleoperator's limbs with those of his own body, teleproprioception is particularly important. All human operators of teleoperators know that lack of stereopsis and visual accommodation cues for depth, absence of peripheral vision, unfamiliarity with particular manipulated objects and their shadow patterns, and inability to "move the head for a better view" all contribute real difficulty in keeping track of where the limbs of the teleoperator are relative to each other and to the vehicle and/or environmental objects in the remote task site. This difficulty is further aggravated by the fact that, while controlling in some degrees of freedom, relative positions may be drifting in other degrees of freedom. It has been shown that a regularly updated dynamic computer model can provide a view from any direction ("moving the head" is done easily in firmware) and aid teleproprioception greatly (Fyler, 1981).

(3) Force feedback and mechanical impedance at the human interface. Force feedback has proven its worth in nuclear hot laboratory remote handling for forty years (Corliss and Johnson, 1968), though it has no advantage for manipulator positioning when there are no force interactions with external objects. When significant time delay is present

continuous force feedback is mostly destabilizing (Ferrell, 1966). However there may be ways to provide short interval force feedback and still retain stable control, for example by cutting off forward control during brief force reflecton episodes, then automatically returning to master-slave coupling with no force feedback.

It is also important to consider the relation of human-arm-to-master-arm mechanical impedance with slave-arm-to-environmental-object impedance, for poor parameter choice or control technique here can produce undesirable transients. A computer simulation has been developed, based on one degree-of-freedom master-slave control, which includes adjustable impedances between human arm and master, slave arm and environmental object, feed-forward telemetry delay and lag, and the same for feedback telemetry. A one degree-of-freedom master-slave system has also been developed to explore the effects on performance of being able to adjust those mechanical impedances in-situ (Raju, 1986). I would assert that we are just beginning to understand the implications of master-slave impedance control.

(4) Teletouch. Just as force reflection and impedance adjustment are important in teleoperation, so too is touch. Touch, of course, is the sensing of spatial-temporal distribution of contact forces between hand and environment, not the sensing of the net force resultant. In the human body touch is sensed by receptors in the skin which are sensitive to changes of force with respect to space and time, but have poor absolute sense, while the force resultant is sensed by receptors in the muscles and tendons, which have better absolute sensing capability. So, teleoperator touch sensing requires special contact surface sensors, made up of many small electrical or optical transducers (Strickler, 1966; Schneiter and Sheridan, 1984).

A big problem is how to display touch information to the human operator. Displays on the skin of his hand where he is gripping a joystick or master arm have proven awkward at best. Computer graphic displays (i.e., to "touch with the eyes") appear easiest, especially with some computer processing already done to identify or suggest what object is being grasped (Schneiter, 1986).

SUPERVISORY CONTROL AND TELEROBOTICS

"Supervisory Control" and "Telerobotics" are assumed to be equivalent terms as applied to teleoperation. "Supervisory control" emphasizes that the human operator, acting from the ocean surface or from a comfortable non-hazardous location, supervises a lower level intelligence embodied in the teleoperator itself by intermittently monitoring and reprogramming as necessary for either routine or emergency situations (Sheridan, 1960; Ferrell and Sheridan, 1967; Sheridan and Hennessy, 1984). This corresponds to a corporation executive supervising a subordinate worker, possibly giving instructions and receiving back reports of work accomplished or difficulties encountered. "Telerobotics", a somewhat newer term currently fashionable in the space community, emphasizes that the teleoperator carries sufficient sensors, effectors and computer intelligence on-board to perform simple tasks automatically - with the programs being updated by the human supervisor over a telecommunications link.

The reasons for a supervisory controlled telerobot in preference to either an autonomous robot or a human diver (free swimming or inside a submersible) performing the necessary tasks are simple to explain. Autonomous robots at the present time are neither intelligent nor reliable enough to perform any but the simplest routine tasks. Human divers certainly have good manipulation and sensing capabilities, but the economic costs are very high and the risks for long duration and unrehearsed tasks are considerable. In time there is no question but that autonomous underwater vehicles (AUVs) will be viable for some tasks. For shallow water tasks human divers will remain hard to beat.

The reasons for supervisory control of a telerobot in preference to direct manual control of a teleoperator (e.g. by master-slave or joystick control) are less easy to explain, but the principal advantages of the former are: (1) quicker task completion in the case of telecommunication time delay due to transmission time or to low bandwidth; (2) quicker automatic reflex responses to unanticipated events; (3) greater accuracy in continuous control of position and force; and (4) ability to perform long-duration tasks without continuous attention by the human operator. Pressures to do future undersea teleoperations with relatively few human operators militate in the direction of telerobotics.

Salient problems and developments in telerobot control are:

(1) Computer-based planning. Both the telerobot being controlled and the components of the structures and devices being manipulated are likely to be known, i.e., drawings of them are likely to be available, so that they are amenable to rather precise geometric computer modeling. Insofar as initial conditions of positions and velocities can be established for various bodies relative to each other, their interactions can be simulated by computer, contingent upon various hypothetical control signals sent down. Assuming there are no outside disturbances or that these also can be modeled, it is possible to simulate the effects of a variety of alternative commands prior to committing to any actual command (Sheridan, 1984). There remain, of course,

constraints on computer speed and the extent to which fine-grain simulations can be run in any fixed time period. Some such test trials may be extensive, others simple.

An example of a simple model-based simulation run in almost real-time is a "predictor display" to the operator to compensate for time delay. In this technique real-time signals from the master arm or joystick generate a computer-graphic display of what the remote arm will do after one round-the-loop time delay. Superposition of such a stick-figure teleoperator hand onto the actual (delayed) video feedback provides anticipative or "lead" control capability and can reduce the time to perform a telemanipulation by forty percent (Noyes and Sheridan, 1984; Mar, 1985; Hashimoto and Sheridan, 1986).

Automatic optimization is tempting to try for, but has proven to be elusive because of the sheer complexity of manipulation tasks and the "curse of dimensionality" (Whitney, 1968). What appears more promising is to let a human supervisor assemble gross task elements, then have dynamic programming or other automatic search routines optimize how the well defined low-level task elements are done (Hardin, 1970).

(2) Command language. This includes software, interface hardware and operating procedures by which the human supervisor may communicate instructions and knowledge to the computer. A number of laboratory prototypes of command languages for supervisory control of manipulator systems have been built, starting in the late 1960's (Barber, 1967). It has been shown that even without time delay, some manipulation tasks can be accomplished more quickly by a human supervisor commanding and initializing each of a series of short subroutines, letting the computer then execute each in turn, than for a human operator to do the whole task in direct manual mode (Brooks, 1979). What appears to work best are typed symbolic statements for the gross commands, with accompanying analogic or "demonstration" commands interjected by means of master-slave or joystick (e.g. using a computer model to generate feedback) to in order to "point" (specify objects, points or trajectories). Symbolic commands often take the form of "Do A until B, then do C, unless D, in which case do E". Such statements can be arranged in a hierarchy, preferably with certain freedom of the supervisor to define and nest statements to his own taste (Yoerger, 1982).

(3) Resolved motion or rate. The human operator, whether controlling in the direct or the supervisory mode, prefers to think and give commands in terms of translations and rotations of the manipulator hand or end effector - not the motions of the intervening segments. Techniques have been developed to perform in real-time the required inverse Jacobean (given the end-point position and orientation,

specify the joint angles) in order to control any series-link manipulator arm (Whitney, 1969). Additional software has been developed to help the operator ensure that the joints are not moving into forbidden configurations or velocities (Minhas, 1986) or to automatically compensate for known relative motion between the reference frame of a manipulator and the reference frame of a body on which manipulations are being performed (Hirabayashi, 1981). The new parallel-link arm developed by Landsberger and Sheridan (1985), which is fully described in Landsberger's paper in this volume, does not require an inverse Jacobean in the forward loop.

(4) End effector dexterity and impedance. It is common practice these days to cry for greater teleoperator dexterity. But what is dexterity? It seems to be two things. First, it is multiple-degrees-of-freedom at the end effector, more than simple one-degree-of-freedom grasp. It requires more fingers with more adaptability in how the fingers can configure themselves to grasp, hold and ungrasp objects. Well-known multi-fingered hands have been developed at MIT and the University of Utah (see Jacobsen in this volume). Even simple one degree-of-freedom additions to conventional master-slave arms may add considerable function. "Flyable" computer graphic simulation for example, permits many degrees-of-freedom in each of several fingers to be controlled simultaneously (Yang, 1986), though to build such a device in hardware may be prohibitive.

Closely coupled with the ability to control several degrees of freedom at the end effector, dexterity is the ability to control positions in combination with interactive forces between end-effector and environmental object. Where the human neuromuscular system easily and naturally adapts to different task demands (e.g., tight position control for threading a needle, soft compliance for throwing and catching a baby in the air), this is difficult to do with present teleoperators. One would like to impose "target" impedance on the end effector - object interaction. The theoretical basis for this has recently been developed, and an experimental demonstration has been built (Kazerooni, 1985). Hogan discusses impedance control in another paper in this volume.

DEVELOPMENT AND EVALUATION OF TELEROBOTIC SYSTEMS

The above discussion points out several problems of telerobotic control and sensing which are deemed critical. In this final section the author opines on some "policy" aspects of telerobotics applied to undersea problems:

(1) Evolution vs revolution. When should free-swimming divers and/or manned submersibles be replaced by ROVs? When should directly controlled (man-in-the-loop) ROVs be replaced by telerobotic (supervisory controlled) ROVs? When should

conventional (power-plus-communications) tethers be replaced by optical fiber (communications only) tethers? When should tethered submersibles be replaced by untethered submersibles controlled over acoustic links? When should ROVs be replaced by autonomous underwater vehicles (AUVs)? Though the latter is the ultimate dream of the technologist it may be the user's nightmare unless care is taken not to claim too much or to try to move too rapidly through the stages of Table 1. Quite possibly all of the levels of Table 1 will continue to be in demand for particular applications. And probably it is better to let the developments evolve, letting each stage be a platform for launching the next stage.

(2) Computer simulation as substitute for prototype hardware development. Even though the claim is made that teleoperation by ROV is ultimately much preferred to directly manned operations, the construction and at-sea testing of hardware prototypes are expensive and time-consuming. When can computer simulation be an effective substitute? The development of aircraft and spacecraft provide close analogies because of the essential human-vehicle interaction. Here the relative advantages of simulator vs. hardware prototype are well understood. We now have a modicum of experience building teleoperator hardware prototypes for undersea applications, and have experimented with some computer graphic "human-flyable" teleoperator simulations (Royer, 1985; Kuentz, 1986). Experimental results suggest that such simulations are useful provided they are kept simple. Otherwise existing computers of either vector or raster-graphic variety may be too slow for real-time rendering of vehicle, manipulator and external objects and ocean floor. Software for touching and grasping environmental objects has proven particulaly computer-time consuming (Winey, 1981; Yang, 1986) and needs further effort.

For experimental purposes as well as training ROV operators computer-simulated arms, hands and objects can be reconfigured quickly and systematically as compared to actual hardware. At the same time we all recognize that it usually takes an actual hardware demonstration to convince a skeptical potential user or funder. Thus easily modifiable ROV hardware "test-beds" such as the MIT Sea Grant 1 will continue to be used in conjunction with computer simulations, no matter how realistic.

(3) Evaluation. The usual demand is that the teleoperator be able to do everything a human diver can do (were he able to "be there"). One would like to put numbers on manipulation performance, but alas, no generic theory of manipulation has yet emerged. Probably the best one can do for some time is to standardize task boards and refine time-accuracy-error scoring schemes to allow comparative performance evaluations of telerobots with each other and with human divers (Bertsche et al, 1978). These evaluation

methods necessarily start out being somewhat arbitrary, but ought to incude a variety of subjectively relevant features, much as does the standard IQ test of human capability. Of course, if one can afford it, for any particular applications one can always run the experiments to compare performance between a suitably trained and protected human diver and an appropriately selected and modified ROV.

REFERENCES

General References

Bertsche, W.R., Logan, K.P., Pesch, A.J and Wernli, R.L. (1978). Evaluation of the design and undersea work capability of the Work Systems Package, NOSC Tech. Rep. 214, San Diego.

Busby, F. (1979). Remotely operated vehicles, US Dept. of Commerce, Rep. No. 03-78-603-0136.

Corliss, W.R. and Johnson, E.G. (1968). Teleoperator Controls, NASA SP-5070.

Johnson, E.G. and Corliss, W.R. (1967). Teleoperators and Human Augmentation, NASA SP-5047.

Vertut, J. and Coiffet, P. (1984). Les Robots (Tome 3a) Teleoperation. Paris: Hermes.

Yastrebov, V.S. and Stefanov, G.A. (1978). Underwater robot/manipulator development, MTS Journal, 12,1.

References from the Man-Machine Systems Laboratory

Barber, D. (1967). MANTRAN: A Symbolic Language for Supervisory Control of an Intelligent Manipulator, SM Thesis, Cambridge, MA: MIT.

Black, J. H. (1970). Factorial Study of Remote Manipulation with Transmission Delay, SM Thesis, Cambridge, MA: MIT.

Brooks, T. L. (1979). SUPERMAN: A System for Supervisory Manipulation and the Study of Human Computer Interactions, SM Thesis and MIT Sea Grant Report. Cambridge, MA: MIT.

Deghuee, B.J. (1980). Operator-adjustable Frame-rate, Resolution and Gray-scale Tradeoff in Fixed-bandwidth Remote Manipulator Control, SM Thesis, Cambridge, MA: MIT.

Ferrell, W. R. (1966). Delayed force feedback. Human Factors. Oct., 449-455.

Ferrell, W. R. and Sheridan, T. B. (1967). Supervisory control of remote manipulation. IEEE Spectrum, 4-10, October, 81-88.

Flyer, D. (1981). Computer graphic representaton of remote environment using position tactile sensors, SM Thesis, Man-Machine Systems Laboratory Report, Cambridge, MA: MIT.

Hardin, P. A. (1970). "And Tree" Computer Data Structures for Supervisory Contro Control of Manipulations, Ph.D. Thesis, Cambridge, MA: MIT.

Hashimoto, T. and Sheridan, T.B. (1986). Paper to be published in the Japanese Journal of Ergonomics.

Hirabayashi, H. (1981). Supervisory Manipulation for Assembling Mechanical Parts while Compensating for Relative Motion. SM Thesis, Man-Machine Systems Laboratory, Cambridge, MA: MIT.

Kazerooni, H. (1985). A Robust Design Method for Impedance Control of Constrained Dynamic Systems, Sc.D Thesis and MIT Sea Grant Report, Cambridge, MA: MIT.

Kuentz, H. (1986). Simulation and Display of an Underwater Vehicle Manipulation, SM Thesis, Cambridge, MA:MIT.

Landsberger, S. and Sheridan, T. (1985). A New Design for Parallel Link Manipulators, IEEE Intl Conf. on Cybernetics and Society, Tucson, AZ, Nov. 12, 13.

Mar, L. E. (1985). Human Control Performance in Operation of a TLAYED Master-Slave Telemanipulator. SB Thesis, Cambridge, MA: MIT.

Minhas, F. (1985). Kinematics of RSI (7-F-SC-56) Manipulator. Man-Machine Systems Lab. Rep., Cambridge, MA: MIT.

Noyes, M. and Sheridan, T. B. (1984). A novel predictor for telemaniuplation through a time delay. Proc. 1984 Annual Conf. on Manual Control, Moffett Field, CA: NASA Ames Research Center.

Raju, G.J., PhD Thesis in progress. Cambridge, MA: MIT.

Ranadive, V. (1979). Video Resolution, Frame-rate and Gray-scale Tradeoffs under Limited Bandwidth for Undersea Teleoperation, SM Thesis, Cambridge, MA: MIT.

Royer, T. (1985). Human-interactive Simulation and Display of an Underwater Remotely Operated Vehicle, Cambridge, MA: MIT.

Schneiter, J. and Sheridan, T. B. (1984). An Optical Tactile Sensor for Manipulators, Robotics and Computer-Integrated Manufacturing, Vol. 1., No. 1, pp. 65-74.

Schneiter, J. (1986). An objective tactile sensing strategy for object recognition and localization, Proc. 1986 IEEE Intl. Conf. on Robotics and Automation, Vol. 2. pp.1262-1273.

Sheridan, T. B. (1960). Human metacontrol. Paper presented at Conference on Manual Control, Wright Patterson L. Paper presented at Conference on Manual Control, Wright Patterson AFB, OH.

Sheridan, T.B. and Verplank, W.L. (1978). Human and Computer Simulation of Undersea Teleoperators, Man-Machine Lab. Rep., Cambridge, MA: MIT.

Sheridan, T. B. (1984). Supervisory Control of Remote Maniuplators, Vehicles and Dynamic Processes. In Rouse, W. B. (Ed.) Advance in Man-Machine Systems Research, Vol. 1. NY: JAI Press, 49-137.

Sheridan, T. B. and Hennessy, R. T., Eds. (1984). Research and Modeling of Supervisory Control Behavior: Report of a Workshop. Washington,D.C.:
National Academy Press.

Strickler, T. G. (1966). Design of an Optical Touch Sensing System for a Remote Manipulator, SM Thesis, Cambridge, MA: MIT.

Whitney, D. E. (1968). State Space Models of Remote Manipulation Tasks, PhD. Thesis, Cambridge, MA: MIT.

Whitney, D.E. (1969). Resolved-Motion Rate Control of Manipulators and Prostheses. Fifth Annual NASA-University Conf. on Manual Control, MIT, Cambridge, MA.

Winey, C. M. (1981). Computer Simulated Visual and Tactile Feedback as an Aid to Manipulator and Vehicle Control. SM Thesis, Man-Machine Systems Laboratory , Cambridge, MA: MIT.

Yang, J. (1986). Computer-graphic Simulation of Hands with Multi-motion Fingers, SM Thesis, Cambridge, MA: MIT.

Yoerger, D. (1982). Supervisory Control of Underwater Telemanipulators: Design and Experiment. Ph.D Thesis, Cambridge, MA: MIT

Task Resolved Robust Control of Vehicle/Manipulator Systems

Dana R. Yoerger
Assistant Scientist
Department of Ocean Engineering
Woods Hole Oceanographic Institution

Jean-Jacques E. Slotine
Assistant Professor
Department of Mechanical Engineering
Massachusetts Institute of Technology

INTRODUCTION

Significant advances must be achieved in several areas to create free-swimming underwater robots that can perform sophisticated manipulation tasks. Specific research topics include sensors and nonlinear state estimation, robustness and adaptation for nonlinear systems with strong bandwidth limitations, vehicle and manipulator design, and man-machine interfaces that free the human operator from the elements of the task which do not actually require his skills or decisions. A research program conducted jointly at Woods Hole and MIT has been focussed on these issues, with immediate application to the JASON system and its early prototypes.

This paper summarizes some of the work in teleoperated underwater vehicles and manipulators taking place in the Deep Submergence Laboratory of the Woods Hole Oceanographic Institution and the Nonlinear Systems Laboratory at MIT. This research will contribute to the development of JASON, an advanced teleoperator for scientific observation and manipulation tasks at depths to 6000 meters. A major goal of the JASON program will be to give the human operator high level control of the motions of the manipulator end effectors through finely controlled closed-loop movements of both the vehicle and its manipulators.

The first part of this paper summarizes recent progress in automating the movements of an underwater vehicle. Whether the vehicle is controlled jointly by a human operator and a computer system or by an autonomous system, the control theoretic problems remain. In each case, the state of the vehicle must follow a specified trajectory despite the uncertain dynamics and the disturbances. Any form of high level control must be built on a closed loop system that is stable and predictable. The vehicle ideally should be designed with automatic control in mind and must be equipped with the appropriate sensors. Automated vehicle movement control will improve performance in observation tasks, and will also permit the vehicle and the manipulator to work together to perform manipulative tasks.

The next part of the paper summarizes work in the supervisory control of teleoperated systems that allows the human operator and the computer system to share the trajectory generation task. This work builds directly on the previous trajectory control work. Given a model of the task to be performed, the degrees of freedom that are well defined can be assigned to the computer system, while the human operator retains direct control over the remaining degrees of freedom. Such techniques hold great promise for improved performance and reduced operator workload.

CONTROL TECHNIQUES FOR UNDERWATER VEHICLES

The control theoretic problems associated with underwater vehicles and manipulators are difficult for many reasons. The dynamics of such systems are significantly nonlinear and often poorly known. Unpredictable disturbances caused by currents or a tether are often present. In other cases, the navigation sensors employed present strong constraints.

In most instances the resulting control system design problems are poorly matched to existing linear design techniqes. Like manipulators, underwater vehicles lack obvious points about which to linearize their nonlinear dynamics. As a result, linear gain scheduling techniques yield large, complex design problems with little assurance of uniform performance or global stability. New nonlinear techniques allow a single controller to be designed based on a model that may be simplified, but does not need to be linearized.

Methodology

Recent advances in nonlinear control theory directly address many problems in underwater vehicle control. In particular, a methodology known as **sliding control** has great potential for making a substantial contribution in this area.

The method allows a controller to be designed directly for the multi-variable nonlinear system, with no need for linearization. This allows the control system designer to take better advantage of the sophisticated hydrodynamic modelling techniques that exist for underwater vehicles, as the resulting models are highly nonlinear. The resulting closed-loop systems can be shown to be stable in the Lyapunov sense [1].

Because uncertainty in the dynamic model can also be accomodated into the design, sliding control is a **robust** technique. Additionally, the trade-off between dynamic uncertainty and performance is quantitatively explicit. These robustness properties can be used to great advantage to simplify the modelling task in addition to dealing with the dynamic aspects that cannot be modelled. Often, the vehicle model can be dramatically simplified without significant performance

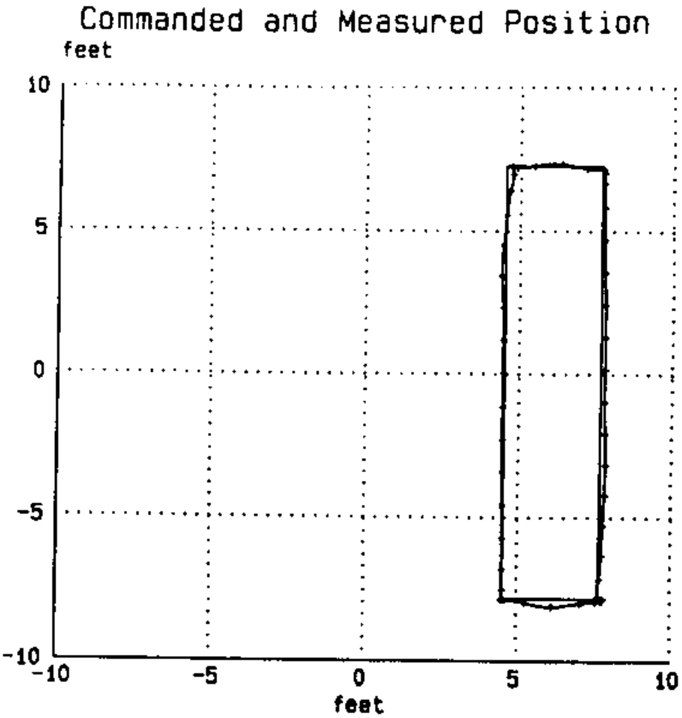

Figure 1. A test run illustrating performance of the closed-loop control system for vehicle translation. Finely controlled vehicle movements can permit a vehicle and manipulator to work together.

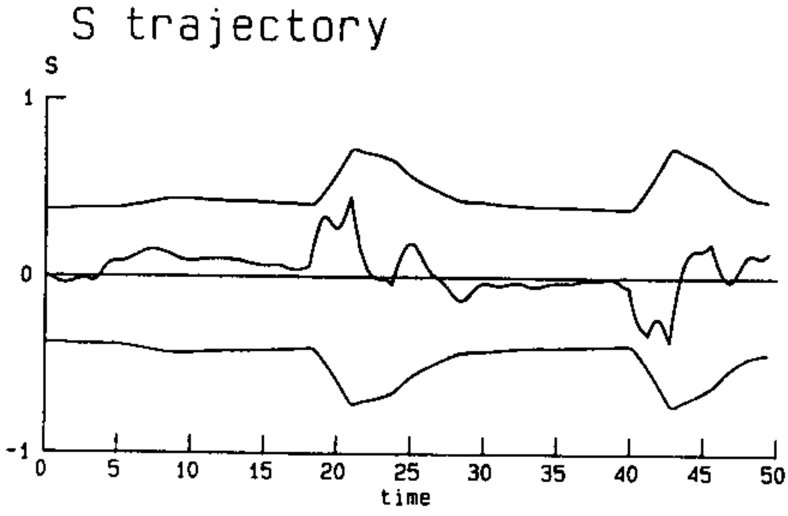

Figure 2. This plot shows the error metric 's' inside the time-varying envelope (called the 'boundary layer') of magnitude ϕ. This plot confirms that the magnitude of the errors is consistent with the projections of sliding control theory, which are based on the desired dynamics and estimates of uncertainty.

penalties, greatly reducing the modelling and parameter identification tasks. These properties can also be very revealing as to how vehicles can be designed that are better suited to automatic control. Demonstrations in simulation and in experiments using underwater vehicles confirm these points [2,3].

Experimental Results

A preliminary experiment was performed by members of the Deep Submergence Laboratory to investigate the effectiveness of sliding control to a vehicle tracking task. A test bed vehicle system (based on a Benthos RPV-430) and a prototype high precision tracking system (by Applied Sonics, Gloucester, VA) were used in the pool test.

The experiment demonstrated several of the important advantages of sliding control. The control system was based on a simplified model of the vehicle dynamics, where many details of the vehicle hydrodynamic behavior were treated as unknown disturbances. In addition, uninstrumented dynamic coupling between horizontal translation and attitude greatly limited the permissable closed-loop bandwidth of the system. Despite these limitations, excellent results were obtained. The level of performance was consistent with the predictions of sliding control theory.

Figure 1 shows a plot of the commanded and actual track of the vehicle in the pool. The commanded trajectory consisted of constant acceleration, constant velocity, and constant deceleration segments. The positional errors are quite low, even with substantial model uncertainty and the limited bandwidth.

Figure 2 confirms that the performance lies within expectations. The plot shows the error metric 's' inside a time-varying envelope or "boundary layer" $\phi(t)$. This plot summarizes performance for a single degree of freedom (forward translation). The magnitude of the boundary layer is directly predicted by sliding theory from estimates of model uncertainty and the bandwidth limit. Not only can sliding controllers deal with uncertainty in a nonlinear setting, but the trade-off between dynamic uncertainty and performance is explicit.

INTERACTIVE TRAJECTORY SPECIFICATION

There are a variety of philosophies for making teleoperators more effective [4,5]. In one approach, the teleoperator is made as anthropomorphic as possible, the goal being to allow the human to work as if he or she were actually at the remote site. Another approach, supervisory control, provides for a sharing of the task between the human operator and a computer control system.

The anthropomorphic approach attempts to project the human operator's effectors and sensors into the remote environment. Devices such as force-reflecting master-slave manipulators and head-coupled stereo television are employed to reproduce the experience of working at the remote site [6]. Such an arrangement is very versatile but often mechanically complex at both ends. Operator fatigue may be a problem.

Supervisory control may take two basic forms depending on the how the task is shared. In one form, the human operator executes part of the task, after which the computer control system executes part, and so on. Such shared control can be termed "serial" supervisory control. In another type of supervisory control, the human operator and the computer control system jointly execute the task, often termed "parallel" supervisory control.

In serial supervisory control, the operator executes part of the task, after which the computer control system completes the next portion while the human operator monitors, and so on. Often, the human operator's responsibilities are to teach positions or movements after which the computer system can execute the task based on a description of the task [7,8]. The task description is entered with an interactive programming interface that emphasizes the teaching of motions or interaction forces when the task is executed. The system can include real-time graphics for verification through simulation and for monitoring. This approach has great potential when communication to and from the remote system is limited in bandwidth or delayed.

In the parallel supervisory control approach, the human operator continuously executes the task but with substantial aid from the computer control system. Resolved Motion Rate Control (RMRC), for example, combines a kinematic model of the remote manipulator to provide the operator with cartesian motion inputs, rather than commands to the individual degrees of freedom [9]. In addition, models of the environment and task can be used in conjunction with sensors to provide the human operator with a higher level interface that can simplify the operator's job and can result in improved performance.

Task-resolved control is a form of parallel supervisory control where the operator is given movement command inputs that correspond geometrically to the task. The goals of task resolved movement control are to improve performance and reduce operator workload in teleoperated tasks.

This concept can be illustrated by several examples. If the teleoperated task is to inspect a vertical cylinder (such as a piling) with a remotely operated underwater vehicle, the operator can be given control of the vehicle in a cylindrical coordinate system centered on the piling while the heading of the vehicle is constrained to point at the axis of symmetry of the cylinder. If the task is to write on a remote blackboard with a piece of chalk, then the computer control system can keep the chalk oriented properly and can also maintain a desired force between the chalk and the blackboard. The operator is then given control over translation in the plane of the board.

Task resolved movement control requires models of the remote machine, environment, and task. Unlike systems that attempt to totally automate a task, simple models that are substantially incomplete can be used to great advantage. This discussion will assume that reliable servo control of the remote machine already exists and that the uncertain, nonlinear dynamics of the vehicle and manipulator are dealt with at that level.

The following discussion will be oriented toward the tasks for which current teleoperators are most commonly used: visual inspection, cleaning and structural inspection, and simple pick-and-place tasks. While tasks requiring impedance or force control are clearly important and are well suited to similar techniques, they will not be pursued here.

Model of the Machine

For these purposes, a simple kinematic model of the vehicle and manipulator will be sufficient. The vehicle and each manipulator link will be considered to be a rigid body. Each manipulator link will have a coordinate frame attached to its outward or distal end in the typical manipulator convention. For a vehicle with a single manipulator, a coordinate frame can be fixed at the manipulator base (considered to be the distal end of the vehicle "link"), although this may violate the vehicle modelling convention of placing the reference frame along the vehicle center line.

Model of the Environment

A wide variety of techniques can be used for modelling objects in the task environment. For the range of tasks considered here, the environment will be modelled as combinations of primitive surfaces. From these primitive surfaces, an artificial potential field may be constructed, as is done in Hogan's impedance control [10] or Khatib's operational space formulation [11]. To construct the artificial potential field, the only information required is the orthogonal distance from each surface. Khatib [11] describes the required computations for parallelepipeds, finite cylinders, and cones, and also suggests the use of n-ellipsoids and n-cylinders.

Construction of Task Space

For task-referenced teleoperation, the artificial potential fields will be used in a different fashion. Hogan and Khatib used the artificial potential field to generate a trajectory that reaches the goal while avoiding obstacles. For these purposes, the artificial potential field will be used to form a distorted space that conforms to the composite set of surfaces. As in impedance control or the operational space formulation, movements in the direction of the gradient of the artificial potential field will be directed toward the nearest objects. However, it will also be useful to move along an equipotential surface, for example to inspect a surface.

Khatib [11] suggests the following potential field formed by summing components from each surface:

$$P_i(x) = \begin{cases} 1/2\ \eta(1/\rho - 1/\rho_o)^2 & \text{if } \rho < \rho_o \\ 0, & \text{if } \rho \geq \rho_o \end{cases}$$

where $P_i(x)$ is the contribution to the total potential at the point x due to the i^{th} surface, ρ is the shortest distance from x to the i^{th} surface, η is a constant gain and ρ_o is a distance beyond which the i^{th} object exerts no influence. The gradient of the complete field is determined from the sum of the gradients of each component:

$$\nabla P_i(x) = \begin{cases} -\eta(1/\rho - 1/\rho_o)\partial\rho/\partial x & \text{if } \rho < \rho_o \\ 0, & \text{if } \rho \geq \rho_o \end{cases}$$

where $\partial\rho/\partial x$ is the partial derivative vector of the distance between the point x and the appropriate surface.

A unit vector in the direction of the gradient defines the direction for motions toward the complex surface:

$$v = \nabla P/|\nabla P|$$

This direction will correspond to the z axis of a coordinate frame.

A two dimensional example is illustrated in figure 3. Here the objects modelled could represent the walls of an underwater structure to be inspected by a remotely operated vehicle. Movements in the direction of the gradient move the vehicle toward the surface, movements along constant potential lines move the vehicle along the surface, and the heading of the vehicle remains pointed at the surface.

In three dimensional tasks, two other directions must be defined to establish a coordinate frame for which the x and y axes lie in the tangent plane to the equipotential surface. Like the direction defined by the gradient, these directions should vary smoothly between objects, and the likelihood of singularities or nonunique solutions should be low.

A solution to this problem can be obtained by assigning a reference direction to each primitive surface. This reference direction will be defined by a unit vector that points from the position x to a reference point. This reference direction will yield the y axis of the coordinate frame when projected onto the equipotential surface. In some cases, such a direction falls out naturally, such as the axis of symmetry of a cylinder. In other cases, the reference direction can be determined from task considerations. In the example of a plane, the reference direction may be any unit vector that lies in that plane.

The reference directions from each of the primitive surfaces can be combined to produce the reference direction in the plane tangent to the equipotential surface. As illustrated in figure 4, the magnitude of each reference vector is multiplied by the component of the potential associated with the appropriate surface. Each of these vectors can then be projected into the tangent plane and the reference direction considering all primitive surfaces can be computed by summing all the projections and then normalizing. The remaining direction u completes the righthanded coordinate frame. This technique will fail to produce a solution only when no reference directions have nonzero components in the tangent plane, or when all components sum to zero.

The cleaning and inspection of underwater welds is a representative task where this technique could be applied. Offshore oil production platforms are often built from tubular members, which are joined by welding. As the welding process may embrittle the metal, such joints are subject to cracking in the high fatigue environment. Key joints are routinely inspected using nondestructive testing techniques that require good manipulator control such as closeup stereo photography, acoustic crack detection, or magnetic particle inspection. All existing inspection techniques require that the entire heat effected zone (HAZ) be first cleaned of all marine growth, oxidation, etc. Such cleaning is typically done by grinding, bushing, or water jetting.

This type of work is now done primarily by human divers, as the vehicles and manipulators now available are not productive enough to be competitive. Using master-slave or joint rate controls, progress can be quite slow. Using task-resolved techniques, the operator can be given controls that allow the tool to be manuevered along the weld, across the weld, or normal to the weld.

Figure 5 illustrates how an appropriate task space can be constructed by superimposing simple models. The structure is a node consisting of a large cylinder (the primary member) with a smaller

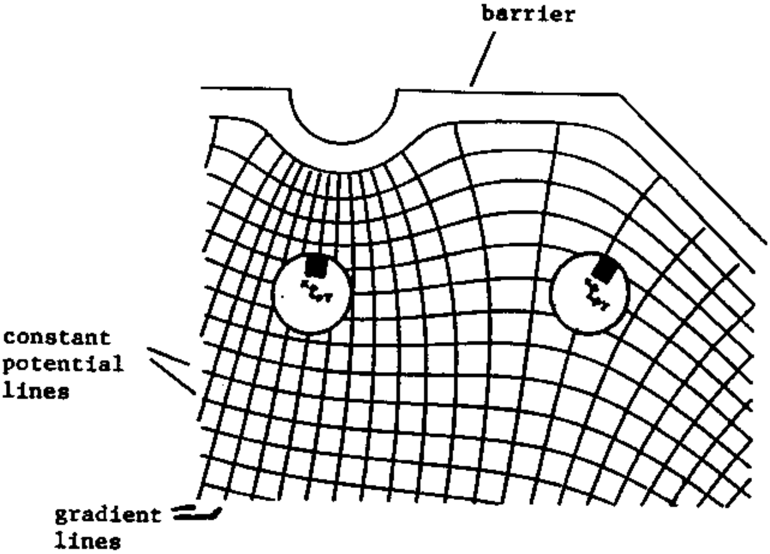

Figure 3. An example of a task space for a two dimensional environment. Artificial potential fields are assigned to objects, which are then summed to produce the total field. The gradient and constant potential lines then form natural directions for the vehicle motions to be commanded by the operator.

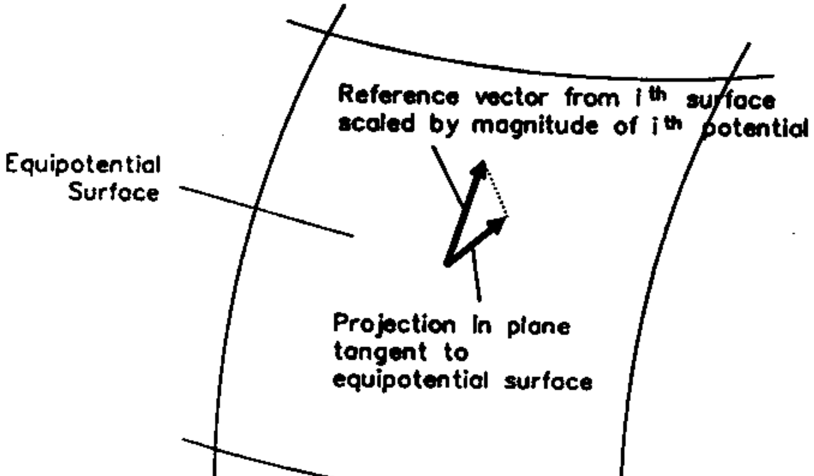

Figure 4. The y axis for the task frame may be constructed by considering reference directions associated with each surface. Each reference vector is scaled by the potential for that surface. The resulting vectors are projected into the plane tangent to the equipotential surface and summed. The resulting vector is then normalized to produce the y axis for the task frame.

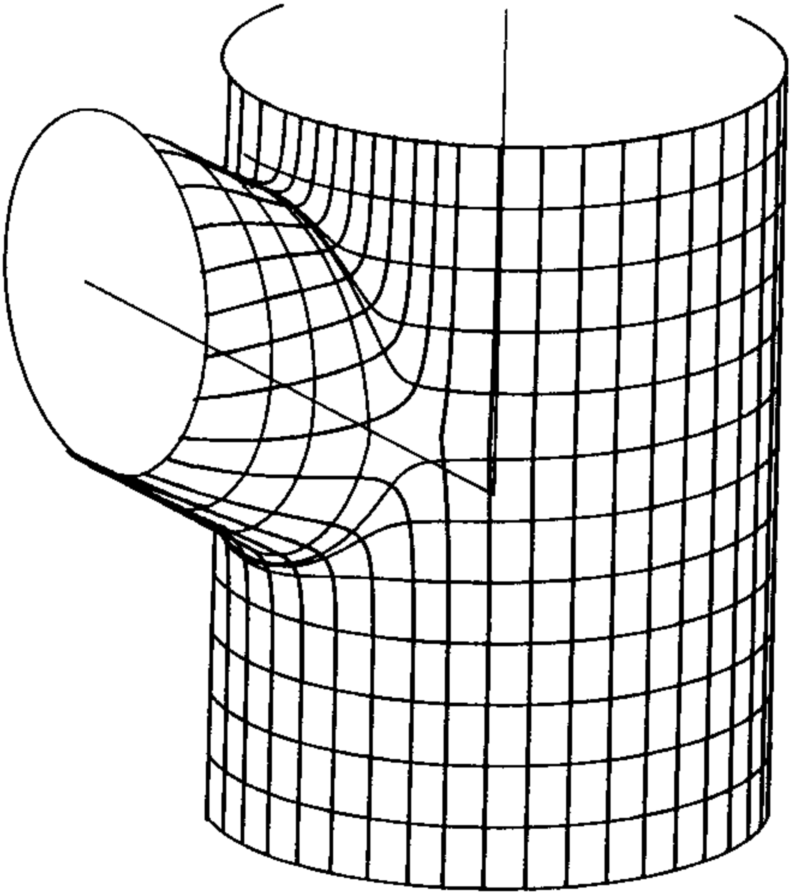

Figure 5. An equipotential surface for the inspection of the node of an underwater structure. For the large cylinder, the reference vector is the axis of symmetry. For the smaller cylinder, the reference vector is normal to the large cylinder and lies in the plane formed by the axes of symmetry of the two cylinders.

cylinder (the secondary member) butted to it. The reference point for the primary is the point where the secondary cylinder axis of symmetry intersects the primary, and the reference direction for the secondary is normal to the primary in the plane formed by both axes of symmetry. Movements in the w direction move the tool toward the node. Movements in the u direction move the tool approximately parallel to the intersection of primary and secondary (the weld). Movements in the v direction are normal to the weld.

CONCLUSION

Remotely operated underwater vehicles and manipulators can be made more productive through improved control systems. This paper reviews progress in two areas. The first involves techniques for automated movement control that directly address problems instrinsic in underwater vehicles: bandwidth constraints, nonlinear dynamics, and dynamic uncertainty. The second area illustrates techniques for producing high level movement control interfaces that allow the human operator and computer to jointly generate complex trajectories.

REFERENCES

1. Slotine, J.-J.E., Sastry, S.S., Tracking Control of Nonlinear Systems Using Sliding Surfaces, with Application to Robot Manipulators, Int. J. Control, 38, 1983.
2. Yoerger, D.R., Slotine, J.-J.E., Robust Control of Underwater Vehicles, IEEE J. Oceanic Eng., Vol OE-11 no. 4., Nov. 1985.
3. Yoerger, D.R., Newman, J.B., Slotine, J.-J. E., Supervisory Control System for the JASON ROV, IEEE J. Oceanic Eng., Vol OE-11 #3, July 1986.
4. Sheridan, T.B., Man-Machine Systems, MIT Press, 1977.
5. Bejczy, A.K., Data-Driven Automation in Remote Applications of Robots, in Robotics Research, Brady and Paul, editors, MIT Press, 1984.
6. Vykukal, H.C, King, R.F., Vallotton, W.C., An Anthropomorphic Master-Slave Manipulator System, in Remotely Manned Systems, Herr, editor, Cal. Inst. of Tech., 1972.
7. Brooks, T.L., SUPERMAN: A System for Supervisory Manipulation and the Study of Human/Computer Interaction, SM Thesis, MIT Dept. of Mech. Eng., May 1979.
8. Yoerger, D.R., Supervisory Control of Underwater Telemanipulators: Design and Experiment, PhD Thesis, MIT Dept. of Mech. Eng., Aug 1982.
9. Whitney, D.E., The Mathematics of Coordinated Control of Prosthetic Arms and Manipulators, ASME J. Dynamic Systems, Measurment, and Control, 1972, pp 303-309.
10. Hogan, N., Impedance Control; An Approach to Manipulation: Part III - Applications, ASME J. Dynamic Systems, Measurement, and Control, March, 1985.
11. Khatib, O., Real-Time Obstacle Avoidance for Manipulators and Mobile Robots, 1985 IEEE Int. Conf. on Robotics and Automation, St. Louis, March 1985.

Impedance Control of Contact Tasks Using Force Feedback

Neville Hogan
Associate Professor
Department of Mechanical Engineering and
Laboratory for Manufacturing and Productivity
Massachusetts Institute of Technology

INTRODUCTION: CONTACT TASKS

Applying robots to contact tasks is one of the challenges of developing robot technology. Contact tasks are those in which the robot grasps or pushes or works on its workpiece. They make up a large proportion of the tasks to which robots could be applied. Yet, ironically, to date most of the successful applications of robots are to tasks in which the robot stays clear of the workpiece. Spray-painting is one example; the robot brings the spray gun through a complex path in space, but (if it's working correctly) never touches the workpiece. Seam welding is another example; the robot end-effector tracks the weld seam, but never touches it. Even in those applications where the robot does contact its workpiece, as in materials handling tasks, it generally does so without precise control of the way it interacts with the object.

CONTACT INSTABILITY

Some of the reasons for this state of affairs can be understood by considering the mechanics of manipulation. When an object is grasped, and the manipulator pushes on it, the object pushes back; it interacts dynamically with the manipulator. In effect, the grasped object becomes a part of the physical hardware of the control system and changes how the robot responds to its actuators. These changes in the way the physical hardware responds to commands from the controller can have a profound effect on the performance of the robot. In fact they can easily destabilise the control system, causing catastrophically pathological behavior.

It might seem that this problem would only become serious when the robot attempted some dynamically demanding task, such as wielding a power tool, say for deburring or grinding. However, even the apparently simple contact task of wiping a surface has proven to be surprisingly difficult. Robot control systems which perform stably and well during free movements can break into unstable oscillation upon contact with a surface. Despite the substantial research effort which has been directed at controlling the force exerted by a robot, attempts to use sensory feedback (e.g. from a wrist force sensor) for this purpose have been plagued by this problem of contact instability. Indeed, a recent review by Whitney (Whitney, 1985) identified this as one of the major challenges of developing robot technology.

In fact, stable control of the force exerted on a surface was reported in the literature about a decade ago, (Whitney ,1977) but to achieve stability, either the surface pushed on had to be compliant, or the force sensor had to be. Theoretical analysis showed that stability could be achieved if the effective stiffness encountered by the robot (the series combination of the sensor stiffness and the workpiece stiffness) was substantially less than the effective stiffness of the robot itself (which is due to drive systems, servo controllers, etc.). More recent experimental work (Roberts et al., 1985) confirmed this limitation.

Restricting contact tasks to compliant surfaces is clearly unacceptable; most workpieces are quite rigid. Using a compliant force sensor seems more reasonable but has the drawback that the position of anything mounted on the force sensor becomes uncertain due to the compliance of the sensor. A compliant wrist force sensor degrades the positioning accuracy of the end-effector. Compliant force sensors at the tips of the fingers don't eliminate this problem, they simply relocate it; the accuracy with which a held object can be positioned is degraded by compliant finger pads.

IMPEDANCE CONTROL OF CONTACT TASKS

This paper presents some recent experimental results using a control algorithm which eliminates this problem. It will be shown to achieve stable control of the force exerted on a rigid surface without the need for a compliant force sensor. The algorithm is an implementation of impedance control, (Hogan, 1985) an approach to robot control which was formulated to deal with contact tasks. Impedance control also provides a unified approach to all aspects of manipulation. In this paper it will be shown that both free motions and contact tasks can be controlled successfully using a single control algorithm; it is not necessary to switch between control modes as task conditions change. Furthermore, the algorithm completely eliminates the need to perform the notoriously difficult inverse kinematic computations.

PRIOR WORK ON ROBOT CONTROL

Prior work on robot control has been dominated by position or motion considerations, at both the planning and execution levels. It is assumed that the robot's task may be planned and defined in terms of a series of desired or target motions. The execution or implementation problem is then to perform these motions with a minimum of error, usually as fast as possible given the limitations of the hardware. This approach is eminently reasonable and quite successful for non-contact tasks (e.g. arc-welding or spray-painting) in which the only aspect of the robot's behavior which needs to be controlled is its motion.

Prior approaches to contact tasks have also been built on this motion-control substrate, but they have been less successful. In general, an ideal kinematic constraint divides the workspace of the end effector into two mutually exclusive subspaces, one in which no motion is possible but force may be exerted, and one in which no force may be exerted but motion is possible[1]. In this way (Mason, 1981) a contact task may be translated into two dual geometry problems and the forces in the constrained directions treated in a manner exactly analogous to the motions in the unconstrained directions.

For example, when a robot end-effector is in contact with an ideal, frictionless rigid surface no force may be exerted along the tangent to the surface, but motion is possible, so it would seem reasonable to plan motion independent of force in those directions. On the other hand, motion along the normal into the surface is impossible; but a force may be exerted in that direction and it would seem reasonable to plan force independent of motion in that direction in much the same way that motion is planned.

Motion control is also at the heart of prior approaches to the execution or implementation of contact tasks (see Whitney, 1985 for a review). Typically, the robot is designed with a fast (frequently analog) motion control servomechanism for each joint. A force control feedback loop (usually digital and slower) is then closed around this motion controller to generate corrections to the commanded motions so as to regulate the force. Unfortunately, this architecture results in problems with contact instability.

To understand this, consider the following example. Suppose a robot (with a position controller on each joint) is wiping a horizontal surface and it is (correctly) generating a downward force in the vertical direction. For some reason (e.g. an irregularity in the surface) an error is generated: the actual position of the end-effector differs from the position command given to the position controller; say the commanded position is now inside the surface. The position controller asks for more actuator effort to eliminate this error. A change in the applied force is generated, a force error. A sensor measures the force error and generates a correction to the commanded motion in the upward direction to reduce the force error.

If the robot response and the controller action were instantaneous, this might work; but they aren't. Response speed is necessarily limited, and errors will go uncorrected for a time. In the typical situation, the position loop is fast but the force loop is less so. If the force sensor is stiff, then small position errors rapidly generate large force errors. These generate large motion corrections which could easily be enough to generate a bigger force error in the opposite direction. To recover the nominal force a correction to the commanded position is applied in the downward direction. Now the robot tries to get to a position further inside the surface than the original commanded position, in the process generates an even bigger force error, and this gives rise to unstable behavior.

This intuitive description of contact instability explains why a soft sensor (or a soft surface, which amounts to the same thing) helps. The same position errors will generate smaller forces, so for the same robot and the same control algorithm, smaller corrections to the commanded motions will be generated, reducing the tendency to instability. Of course, as outlined above, a soft sensor generates a new set of problems.

THE MECHANICS OF INTERACTION

The problem with this motion-based approach is that it unrealistically ignores the fundamental mechanics of interaction in a contact task. The control system senses force and in response tries to dictate the motion of the robot and the grasped workpiece. It embodies an implicit assumption that motions may be imposed on the workpiece and that it will determine how hard to push back in response.

---

[1] An ideal kinematic constraint is assumed to be frictionless.

However, if the workpiece is kinematically constrained, the nature of a kinematic constraint is such that arbitrary imposed motions may not be possible. Arbitrary forces may be applied; but the kinematic constraint determines whether the workpiece will move at all and if so, in what way and by how much. The mechanics of this kind of object make it a mechanical admittance.

By this reasoning it might seem that an appropriate strategy would be to control force, not motion (Ishida, 1977; Mason, 1981; Shimano, 1977). But to focus on controlling force alone unrealistically ignores the motions due to dynamic interactions. One must consider the dynamic relation between force and motion.

If the workpiece produces an output motion for an input force, then, because robot and workpiece are connected, the robot control system must have the opposite kind of behavior. It must accept an input motion and produce an output force, much the way a spring or an automobile shock-absorber does; it must behave as a mechanical impedance. The issue here is not merely one of semantics; it is a matter of what may be treated as an input and what may be treated as an output; a matter of what the robot physically can and cannot do.

The distinction being drawn here is that the motion-based approaches assume the task may be defined in terms of motions. The controller is then designed accordingly as though the interactions with the workpiece were a source of disturbances producing motion errors. The dual approach, force control, assumes the task may be described in terms of forces. The controller is then designed accordingly as though the interactions with the workpiece were a source of disturbances producing force errors.

AN ALTERNATIVE: IMPEDANCE CONTROL

An alternative is to recognise that the dynamic interactions are not a source of disturbances to be rejected, but an integral part of the task. This is the idea behind impedance control. It is assumed that the robot's task must fundamentally be described, not in terms of motions, nor in terms of forces, but in terms of the relations between them. In the case of a robot performing contact tasks, that relation should be an impedance. Accordingly, in implementation, the controller is designed to modulate and regulate, not the robot's motion, nor the force it exerts, but its output impedance.

In fact, when the interaction between a robot and its environment is considered, *that's what any feedback controller really does.* The control algorithm implements a relation between sensed quantities and actuator efforts. Combined with the hardware, this produces a change in the robot's total dynamic behavior. At the robot end-effector, that change shows up as a modified output impedance. Therefore it makes sense to design the controller to do what it naturally does — modify the robot's output impedance.

How does impedance control compare to some of the alternative approaches? A hybrid combination of motion control and force control in orthogonal directions has been explored (Raibert and Craig, 1981) but it simply combines the motion control and force control approaches and does not circumvent the problems inherent in each. Another approach is stiffness control (Salisbury, 1980). Because the objective of that approach is to implement a relation between force and motion, it is closely related to impedance control. However, stiffness is merely the static component of a robot's output impedance. Impedance control goes further and attempts to modulate the dynamics of the robot's interactive behavior.

Describing the action of the controller in terms of the *dynamic* impedance changes it produces is a considerable aid to — indeed, is the key to — understanding what the controller does. Force feedback does not merely regulate force; it changes impedance. For a typical robot, that impedance change is not a modified stiffness; as we will see next, the principal effect of force feedback is to change the robot's apparent inertia.

WHAT FORCE FEEDBACK DOES

A common model of a robot assumes that it may be described as an inertial mechanism driven by actuators which exert a controllable torque. An extremely simplified one-dimensional model of the basic mechanics is a mass (representing the inertia of the mechanism) acted on by two (opposing) forces, one due to the actuators and one due to the environment. The mechanics of this system are simple:

$$M_{actual}\, d^2x/dt^2 = F_{actuator} - F_{external} \qquad (1)$$

In the absence of a feedback controller, (i.e. $F_{actuator} = 0$) an observer in the environment would perceive the system as an inertia.

$$M_{actual}\, d^2x/dt^2 = - F_{external} \qquad (2)$$

Now assume the following simple force-feedback control algorithm is implemented.

$$F_{actuator} = G\,(F_{reference} - F_{external}) \qquad (3)$$

This controller generates an actuator force proportional to the deviation of the measured external force from a commanded reference force. For simplicity, assume for the present that the reference force is zero. Combining the controller with the physical system:

$$M_{actual}\, d^2x/dt^2 = - (G + 1)\, F_{external} \qquad (4)$$

Dividing by G + 1 and defining $M \equiv M_{actual}/(G + 1)$

$$M\, d^2x/dt^2 = - F_{external} \qquad (5)$$

This has the same form as equation 2. From the viewpoint of an observer in the environment, the controller changes the apparent dynamic behavior of the system. In the equations, M plays the role of mass, so for this simple system, negative force feedback reduces apparent inertia. This makes physical sense: whenever an external force is applied, the controller detects it and make the actuators assist the external force so that the mass accelerates more rapidly than it would in the absence of feedback. Thus the force required to produce a given acceleration — the apparent inertia — is reduced. Note that we reach this conclusion without having to make any assumptions about the dynamics of the environment.

A force feedback controller such as the above is usually considered to regulate force. The foregoing analysis does not contradict this, but augments it. To understand this we will make some simple assumptions about the environment, assuming that the workpiece is rigid and the force sensor is modelled as a linear spring of some stiffness, K, with some internal linear damping with viscous coefficient B. These elements generate the external force.

$$F_{external} = B\, dx/dt + K\, x \tag{6}$$

Rearranging and using the Laplace variable, s:

$$x = \frac{F_{external}}{(Bs + K)} \tag{7}$$

Combining equations for the controller and the system we obtain the transfer function:

$$\frac{F_{external}}{F_{reference}} = \frac{(Bs + K)G/(G + 1)}{M\, s^2 + B\, s + K} \tag{8}$$

This is the transfer function of a simple type-zero controller which regulates the external force to follow the reference force. For example, at steady state,

$$F_{external} = F_{reference}\, G/(G + 1) \tag{9}$$

However, note that to reach this conclusion we had to make assumptions about the dynamics of the environment.

This more conventional description of the controller action is correct, but incomplete. It does not specify how the system responds to its environment. Yet the response of the system to its environment is a major part of what is going on during a contact task. The point is that a feedback controller does not merely regulate a variable such as force or position, it changes the dynamic behavior of the system.

IMPLEMENTING IMPEDANCE CONTROL

Impedance control is based on the recognition that the controller changes the dynamic behavior of the system. To implement impedance control, the first step is to specify the desired behavior of the robot, the target impedance. To arrive at a reasonable choice, consider the basic mechanics of the robot hardware. A common model of a robot assumes that it may be described as an inertial mechanism driven by actuators which exert a controllable torque.

$$I(\theta)\ddot{\theta} + C(\dot{\theta},\theta)\dot{\theta} = \tau_{actuator} - J^t(\theta)\, F_{external} \tag{10}$$

This is a reasonable model of the mechanism (described below) on which the impedance control algorithm was implemented. The model can be generalised to include terms accounting for gravity and friction in the mechanism without preventing derivation of the control algorithm, (Hogan, 1985) but this was unnecessary for the work reported here. The dominant behavior of the hardware along each degree of freedom is that of a second order system. Consequently a reasonable target impedance is also second order (but simpler) in each degree of freedom.

$$- F_{external} = M\ddot{x} + B\dot{x} + K(x - x_0) \tag{11}$$

This target impedance is specified by the programmer or a higher-level supervisory system. The quantity x_0 is the nominal equilibrium position of the end effector at steady state in the absence of any external forces. Because it is specified in software, it may go to positions beyond the reachable workspace of the robot and therefore it is referred to as a *virtual position*. It is used in the impedance controller to specify how the robot moves. Note, however, that the target impedance specifies much more than just the robot's motion.

A control system to implement this target impedance can be derived as follows. Purely for the sake of simplifying the derivation, assume that there is a fictitious actuator at the tip of the robot which can generate controllable forces. These controllable forces are related to the controllable torques of the actuators through the Jacobian of the mechanism.

$$\tau_{actuator} = J^t(\theta)\ F_{actuator} \tag{12}$$

The robot dynamic equations now become

$$I(\theta)\ddot{\theta} + C(\dot{\theta},\theta)\dot{\theta} = J^t(\theta)\ (F_{actuator} - F_{external}) \tag{13}$$

The acceleration of the robot joints is

$$\ddot{\theta} = I^{-1}(\theta)[J^t(\theta)\ (F_{actuator} - F_{external}) - C(\dot{\theta},\theta)\dot{\theta}] \tag{14}$$

The corresponding acceleration of the end-effector is

$$\ddot{x} = J(\theta)\ddot{\theta} + \dot{J}(\theta)\dot{\theta} \tag{15}$$

Substituting

$$\ddot{x} = J\ I^{-1}J^t\ (F_{actuator} - F_{external}) - J\ I^{-1}C\dot{\theta} + \dot{J}\ \dot{\theta} \tag{16}$$

For clarity and notational convenience, the dependence of I, J and C on the robot configuration has not been written explicitly.

The quantity $J\ I^{-1}J^t$ has an important physical meaning. It is the mobility tensor (Hogan, 1985) of the robot end-point. We will denote it by W.

$$W \equiv J\ I^{-1}J^t \tag{17}$$

Its inverse is the actual inertia of the robot end-point.

$$W^{-1} = M_{actual} \tag{18}$$

Note that M_{actual} depends on the robot configuration θ.

To obtain the control law, we first solve for the fictitious actuator force.

$$F_{actuator} = W^{-1}[\ddot{x} + J\ I^{-1}C\dot{\theta} - \dot{J}\ \dot{\theta}] + F_{external} \tag{19}$$

But the desired acceleration of the end-effector is

$$\ddot{x} = M^{-1}\{K(x_0 - x) - B\dot{x} - F_{external}\} \qquad (20)$$

Substituting

$$F_{actuator} = W^{-1}M^{-1}\{K(x_0 - x) - B\dot{x}\}$$

$$+ W^{-1}[J\ I^{-1}C\dot{\theta} - \dot{J}\ \dot{\theta}] + [1 - W^{-1}M^{-1}]F_{external} \qquad (21)$$

Now substitute the controllable actuator torques for the fictitious forces.

$$\tau_{actuator} = J^tW^{-1}M^{-1}\{K(x_0 - x) - B\dot{x}\}$$

$$+ J^tW^{-1}[J\ I^{-1}C\dot{\theta} - \dot{J}\ \dot{\theta}] + J^t[1 - W^{-1}M^{-1}]F_{external} \qquad (22)$$

Finally, use the robot kinematic equations to express the end-point position and velocity in terms of the joint position and velocity.

$$\tau_{actuator} = J^tW^{-1}M^{-1}\{K\ (x_0 - L(\theta)) - BJ\dot{\theta}\}$$

$$+ J^tW^{-1}[J\ I^{-1}C\dot{\theta} - \dot{J}\ \dot{\theta}] + J^t[1 - W^{-1}M^{-1}]F_{external} \qquad (23)$$

This is a nonlinear impedance control algorithm. Given measurements of joint motions and external forces it specifies the actuator efforts required to make the robot end-effector exhibit the specified target impedance.

This algorithm has several interesting features. It requires measurements of joint positions, joint angular velocities and external forces, each of which can be measured with reasonable fidelity in practice. It requires the specification of the target behavior in terms of x_0, M, B and K. This information must be supplied by the programmer or a higher-level supervisory system. The impedance controller permits parameter adaptation; each of these quantities may be changed with time to suit a particular task. Note, however, that the components of the target impedance need not be linear. For example, the position-dependent component $K(x_0 - x)$ can be nonlinear, and for some applications this can be quite useful. A nonlinear position-dependent component has been used in an impedance control algorithm which successfully protected a robot from collisions with moving obstacles. (Andrews and Hogan, 1983)

ELIMINATING INVERSE KINEMATICS

The typical motion-based approach to robot control consists of specifying the robot's task in terms of appropriate motions of its end-point. Because the control system which executes this planned task acts on the joints, not the end-point, it is necessary to translate the end-point motion specifications into a corresponding set of joint motion specifications. This requires the inversion of the kinematic equations of the robot mechanism. Unfortunately, this is frequently an extremely complex and difficult computational problem. Its solution is multiple valued and may require iterative procedures (which can be disastrous if the computation is to be performed in a real time control system). It has been described as the most difficult problem in robot control (Paul, 1981) and has received much attention in the literature.

The impedance control algorithm derived above completely eliminates the need to solve the inverse kinematics problem. When the virtual position x_0 is moved around the workspace the algorithm will make the actual position of the end-point follow it. How closely it follows is a function of the speed of movement and the choice of the impedance parameters. Given an appropriate choice of M, B and K, the deviation between the two can be kept small, and the algorithm successfully controls motion. But this is achieved without any inverse kinematic computations.

THE EFFECT OF FORCE FEEDBACK

The terms which multiply the measured external force determine the effective force feedback gain. To understand the action of the force feedback, consider the application of this algorithm to a one degree of freedom system. The term J^t becomes a scaling constant; for simplicity assume it is unity. It can then be seen that the term multiplying the force measurement depends on the ratio of the actual inertia of the end-point (which varies with position) to the target inertia. But (comparing to equation 3) this term is the negative of the force feedback gain, so

$$-G = [1 - M_{actual}M^{-1}] \qquad (24)$$

Rewriting, we obtain

$$M = M_{actual} /(G + 1) \qquad (25)$$

M is one of the parameters of the target impedance, the apparent inertia imposed by the action of the force feedback. This is exactly the same relation between force feedback and apparent mass obtained earlier. Again, the true action of a feedback controller is to change the dynamic behavior of a system; in this case force feedback changes apparent inertia.

EXPERIMENTAL HARDWARE

To establish the practicality of this nonlinear control algorithm, it was implemented using the apparatus shown in figure 1. Complete details of the implementation are presented in Wlassich (1986). The hardware consists of a mechanism with two links which are free to move in the horizontal plane. The popular SCARA robot design (Makino et al., 1980) employs a kinematically similar mechanism. The links are driven by disc-armature DC permanent-magnet motors, both of which are mounted on the supporting base. The inner link is directly mounted on the shaft of one motor. The outer link is driven through a parallel-link mechanism which is mounted directly on the shaft of the second motor.

The motors are driven by transconductance amplifiers with high-gain internal current feedback. Coupled with the low inductance of the iron-less armature, a voltage input to the amplifiers results in an accurately proportional torque between the armature and the stator. Thus, aside from any frictional losses, these actuators produce a controllable torque.

Low-friction ball bearings were used at all joints, so the drive system has an extremely low output mechanical impedance, similar to a Direct-Drive robot (Asada and Youcef-Toumi, 1984). The links and the motor armatures have low inertia and as a result the maximum acceleration and speed of this mechanism are considerably greater than that of a typical commercial robot. In terms of end-point motion, representative values are 100 in/second and 3.3 g respectively (Wlassich, 1986).

A low-noise analog position potentiometer was mounted at the pivot point of each link. A two-axis analog force joystick was mounted at the tip of the outer link to provide measurement of the horizontal interface forces between the mechanism and its environment. Each of these sensor voltages was digitised by a 12-bit A/D converter. An estimate of joint angular velocity was obtained by digital differentiation of the digitised position measurement.

The controlling computer was a DEC LSI 11/23. The nonlinear algorithm defined symbolically above was implemented in software and the commanded torques transmitted to the hardware through D/A converters. The time required for one computational cycle was 9 milliseconds, resulting in an effective sampling frequency of 111Hz. This means that the effective bandwidth of the controller was approximately 1/20th of this, a little over 5Hz. Translating this performance into movement times, if the robot were peforming repetitive motions between two points, each point-to-point move would last about 1/10th of a second.

Because this mechanism was designed as a laboroatory-scale bench-top test apparatus, and was therefore not subject to the constraints imposed on a commercial robot, its performance far exceeds that of a typical robot. The merit of this approach was that the mechanism provided an exacting and informative test of the impedance control algorithm. By restricting the apparatus to two degrees of freedom the significant kinematic nonlinearities of a robot were retained but our understanding of the behavior of the system was not occluded by the mechanical and kinematic complexity of six or more degrees of freedom. The relatively high speed and acceleration capability of the apparatus meant that the behavior of the algorithm, especially any tendency to instability, was not obscured by sluggish hardware performance.

TEST TASK

To test the effectiveness of the impedance controller, a simple task was chosen: the virtual position x_0 was brought at constant speed through a circle in the workspace of the apparatus. Then a rigid barrier was placed in the workspace so that its flat front surface intersected that circle. This task provided a comprehensive test of the impedance controller's ability to control the path of the end-point during free motions, to control the interaction forces during constrained motions, and to deal with the transitions between these two parts of the task.

If the algorithm successfully achieves the target dynamic behavior, then the time histories of force and motion should be as shown in figure 2. During free motion the force will be zero and the path of the end-point will describe an arc of a circle. Upon contacting the barrier the end-point should slide smoothly along a chord of the circle defined by the surface of the barrier. At the moment of contact with the barrier there should be a brief impulsive force due to the collision of the end-point inertia with the rigid surface, but immediately afterwards the force should drop to zero. Then, as the virtual point x_0 moves inside the surface but the end-point x remains on the surface, the difference between these two should result in an interface force determined by the target impedance. At low speeds, the dominant term will be due to the static component, $K(x_0 - x)$. If the target impedance has linear components, the time profile of the interface force should describe a circle (or a section of an ellipse, depending on the choice of force and time scales).

EXPERIMENTAL RESULTS

The actual performance of the impedance-controlled apparatus is shown in figure 3. The system remains stable throughout all phases of the task, both constrained and unconstrained. The combination of the environment (the barrier) and the force sensor is quite stiff. The barrier is a concrete block with essentially infinite stiffness; the force sensor stiffness was measured to be 700 lbf/ft. In this instance the effective stiffness of the robot (which is entirely due to the action of the controller) is quite modest, 4.5 lbf/ft. Similar results were obtained for all values of stiffness and viscosity throughout the range which could be achieved with this apparatus. Stable interaction with a rigid constraint was unaffected by the value of the stiffness and/or the viscosity of the robot. Clearly, the restriction that a compliant sensor be used to achieve stable interaction does not apply to this impedance control algorithm.

During the constrained portion of the task the interface force is as predicted from the ideal target behavior. For clarity, this data has not been smoothed or filtered, as can be seen from the high frequency fluctuations (which are due to the force sensor) superimposed on the ideal performance. Thus, within the limitations of this apparatus, the impedance controller achieves good control of interface force during a contact task.

Throughout the entire task, the path of the end-point is as predicted from the ideal target behavior. Again, for clarity, the data are presented without filtering. The apparently random fluctuations are due to noise in the system. For example, as the end-point "wipes" along the barrier, its true path is constrained to be a perfectly straight line. Within the limitations of this apparatus, the impedance controller achieves good control of end-point position. But remember, this is achieved without any inverse kinematic computations.

WHY USE FORCE FEEDBACK?

It could be argued that given this idealised laboratory-scale apparatus with its torque-controlled actuators and the low friction in the drive-train, good control of the interface force is to be expected whatever control algorithm is used. How much of the observed perfomance can be credited to the impedance control algorithm and how much is due to the apparatus? To examine this, the same test task was performed using the same computational algorithm, but with the force feedback sensors disconnected. The results are shown in figure 4. During the free motion phase of the task, the path control is essentially as before (as it should be). When the end-point collides with the barrier, the system still remains stable and the end-point slides along the constraining surface.

This is an interesting result in its own right; it demonstrates that, given proper design of the hardware components, force feedback is not necessary to achieve stable accomodation of a constraining surface. It confirms an earlier result in which a constrained motion task was successfully performed by an impedance controlled robot without force feedback (Andrews, 1983).

However, it is important to clearly distinguish between stability and performance. Although the system is still stable, the performance is considerably poorer. Upon colliding with the barrier, the end-point bounces several times, repeatedly impacting the surface. As a result, the time history of the interface force departs significantly from the ideal. Comparing with the previous figure it is quite clear that the impedance control algorithm accomplishes a significant improvement in performance by making appropriate use of force feedback.

CONCLUSION

These experimental results clearly demonstrate that impedance control is a practical strategy for performing contact tasks. To achieve stable behavior while wiping a rigid surface it was not necessary to use a sensor more compliant than the robot, as has been reported in prior approaches to this type of task. In fact, changing the effective stiffness of the robot did not influence the stability of the apparatus during constrained motion. This should not be surprising as the analysis presented here showed that for this kind of mechanism, the true action of force feedback is related to the apparent inertia of the system, rather than its stiffness.

The experimental results also show that given an appropriately designed mechanism, stable interaction may be achieved without using force feedback. However, the results also demonstrate that a control algorithm which can make proper use of force feedback (such as the impedance controller presented here) offers some benefits to offset its cost: it yields a substantial improvement in performance.

In practice, contact tasks also involves phases when the robot is not in contact. The implementation presented here demonstrates that impedance control can provide effective control of unconstrained motions. Furthermore, it does so without the need to perform the complex and difficult inverse kinematic calculations.

A key point is that the same controller successfully deals with all parts of the task: the free motion, the constrained motion and the transitions between them. An alternative might be to restructure the controller (e.g. change from path control to force control) when the robot contacts objects in its environment. But this would require the controller to identify the moment of contact and rapidly switch between different modes of control. Using impedance control this is not necessary. Nor is it necessary for the system to use "guarded moves" (Mason, 1981) to make the transition between free and constrained motion. The impact due to collision with the environment does not destabilise the system. Impedance control is a unified approach to robot manipulation. The results presented here demonstrate one of the practical benefits of a unified approach; the same control algorithm is competent in all parts of the task. In this regard, impedance control is simpler than some of the alternative strategies.

Manipulation is fundamentally a process of mechanical interaction. During contact tasks dynamic interaction between a manipulator and its workpiece becomes the dominant factor determining stability and performance. The key idea behind impedance control is the recognition that the true action of a controller on physical hardware is not merely to regulate motion, or force, or stiffness, but to change its apparent dynamic behavior. Consequently, the controller is designed with the objective of changing the dynamic behavior to an appropriately chosen target impedance. The analytical and experimental results presented here demonstrate the practicality of this approach.

ACKNOWLEDGEMENTS

This work was supported in part by National Science Foundation Grant No. ECS-8307461 and in part by Daewoo Heavy Industries, Ltd.

REFERENCES

Andrews, J. R. (1983) Impedance Control as a Framework for Implementing Obstacle Avoidance in a Manipulator. S.M. Thesis, Department of Mechanical Engineering, M.I.T.

Andrews, J. R. and Hogan, N. (1983) Impedance Control as a Framework for Implementing Obstacle Avoidance in a Manipulator. *Control of Manufacturing Processes and Robotic Systems*, pp.243-251, ASME, New York.

Asada, H. and Youcef-Toumi, K. (1984) Analysis and Design of a Direct-Drive Arm With a Five-Bar-Link Mechanism. *Journal of Dynamic Systems, Measurement, and Control*, pp. 225-230.

Hogan, N. (1985) Impedance Control: An Approach to Manipulation: Part I - Theory, Part II - Implementation, Part III - Applications. *Journal of Dynamic Systems, Measurement, and Control*, pp. 1-24.

Ishida, T. (1977) Force Control in Coordination of Manipulator Fine Motions. *Fifth International Conference on Artificial Intelligence.*

Makino, H., Furuya, N., Soma, K. and Chin, E. (1980) Research and Development of the SCARA Robot. *Proceedings of the Fourth Annual International Conference on Production Engineering*, Tokyo.

Mason, M. (1981) Compliance and Force Control for Computer Controlled Manipulators. *IEEE Transactions on Systems, Man and Cybernetics*, pp. 418-432.

Paul, R. (1981) *Robot Manipulators: Mathematics, Programming, and Control.* M.I.T. Press, Cambridge, MA.

Raibert, M. and Craig, J, (1981) Hybrid Position/Force Control of Manipulators. *Journal of Dynamic Systems, Measurement, and Control*, pp. 126-133.

Roberts, R., Paul, R. and Hillberry, B. (1985) The Effect of Wrist Force Sensor Stiffness on the Control of Robot Manipulators. *Proceedings of the IEEE Conference on Robotics and Automation.*

Salisbury, J. K. (1980) Active Stiffness Control of a Manipulator in Cartesian Coordinates. *IEEE Proceedings, Conference on Decision and Control.*

Shimano, B. (1977) Force Control *Stanford University Artificial Intelligence Laboratory Memo 285.4.*

Whitney, D. (1977) Force Feedback Control of Manipulator Fine Motions. *Journal of Dynamic Systems, Measurement and Control*, pp. 91-97.

Whitney, D. (1985) Historical Perspective and State of the Art in Robot Force Control. *Proceedings of the IEEE Conference on Robotics and Automation.*

Wlassich, J. J. (1986) Nonlinear Force Feedback Impedance Control. S.M. Thesis, Department of Mechanical Engineering, M.I.T.

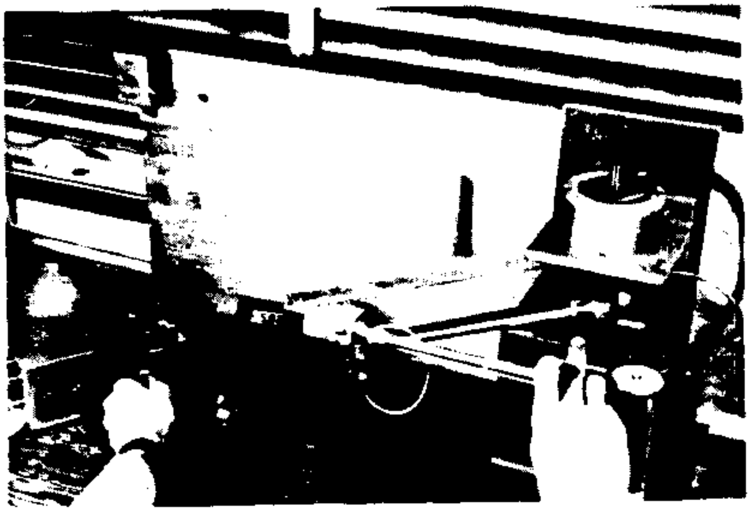

Figure 1. A photograph of the mechanical apparatus used to implement the impedance controller (from Wlassich, 1986).

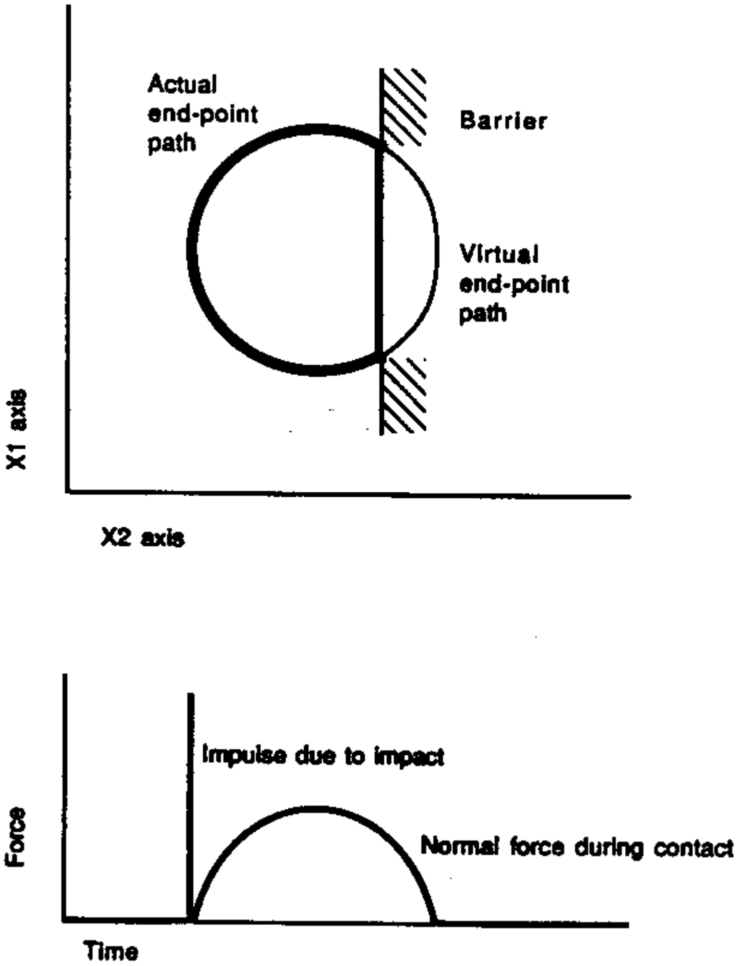

Interface Force vs. Time

Figure 2. The ideal performance of the impedance controller on the test task. A plan view of the actual path (heavy line) and the virtual path (light line) are shown on the top. The time history of the interface force is shown on the bottom.

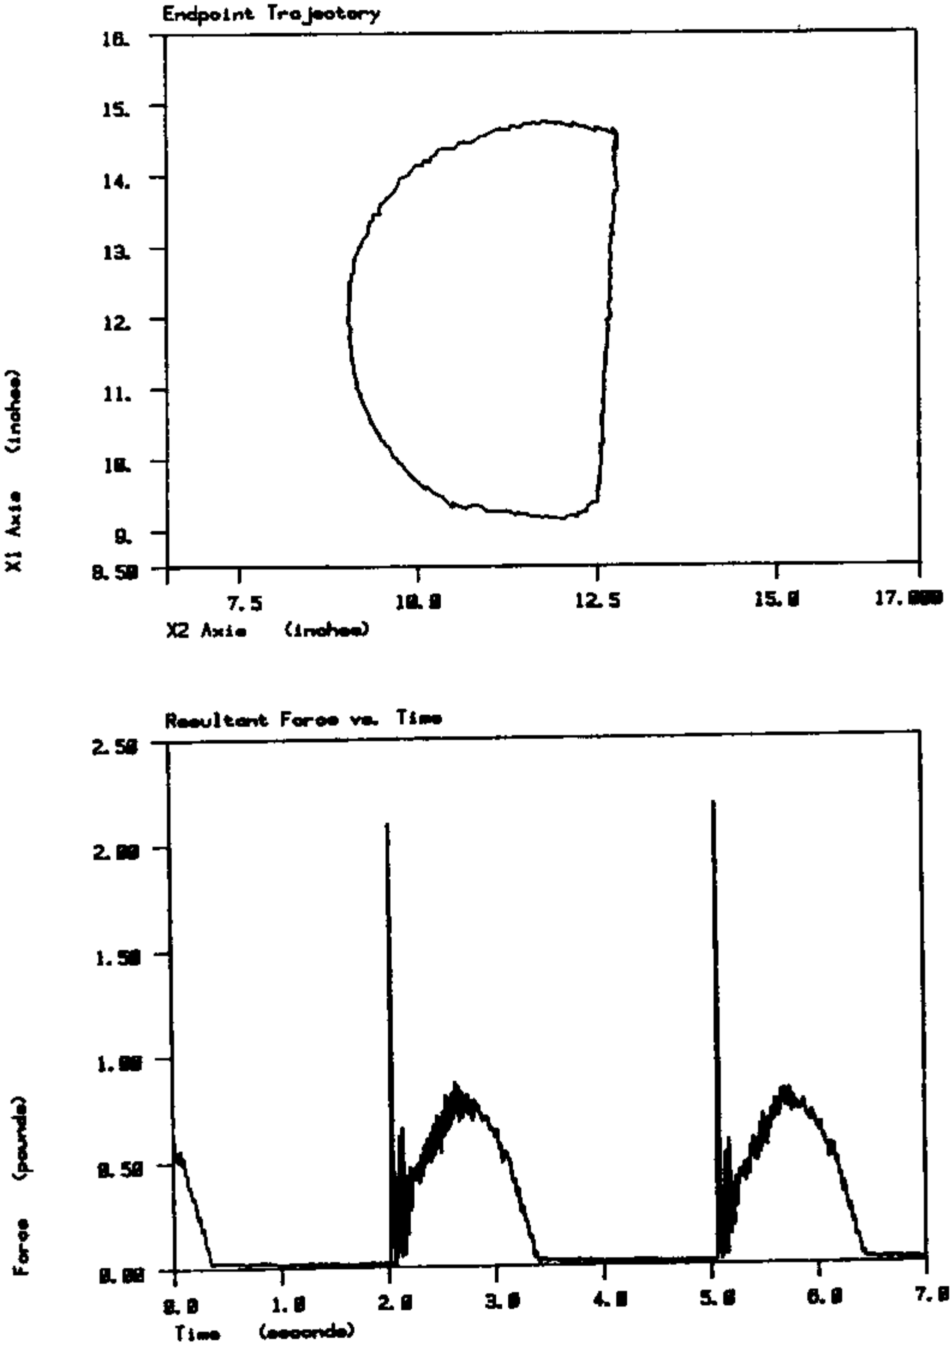

Figure 3. Actual performance of the experimental apparatus on the test task. The path of the end-point is shown on the top. The time history of interface force is shown on the bottom.

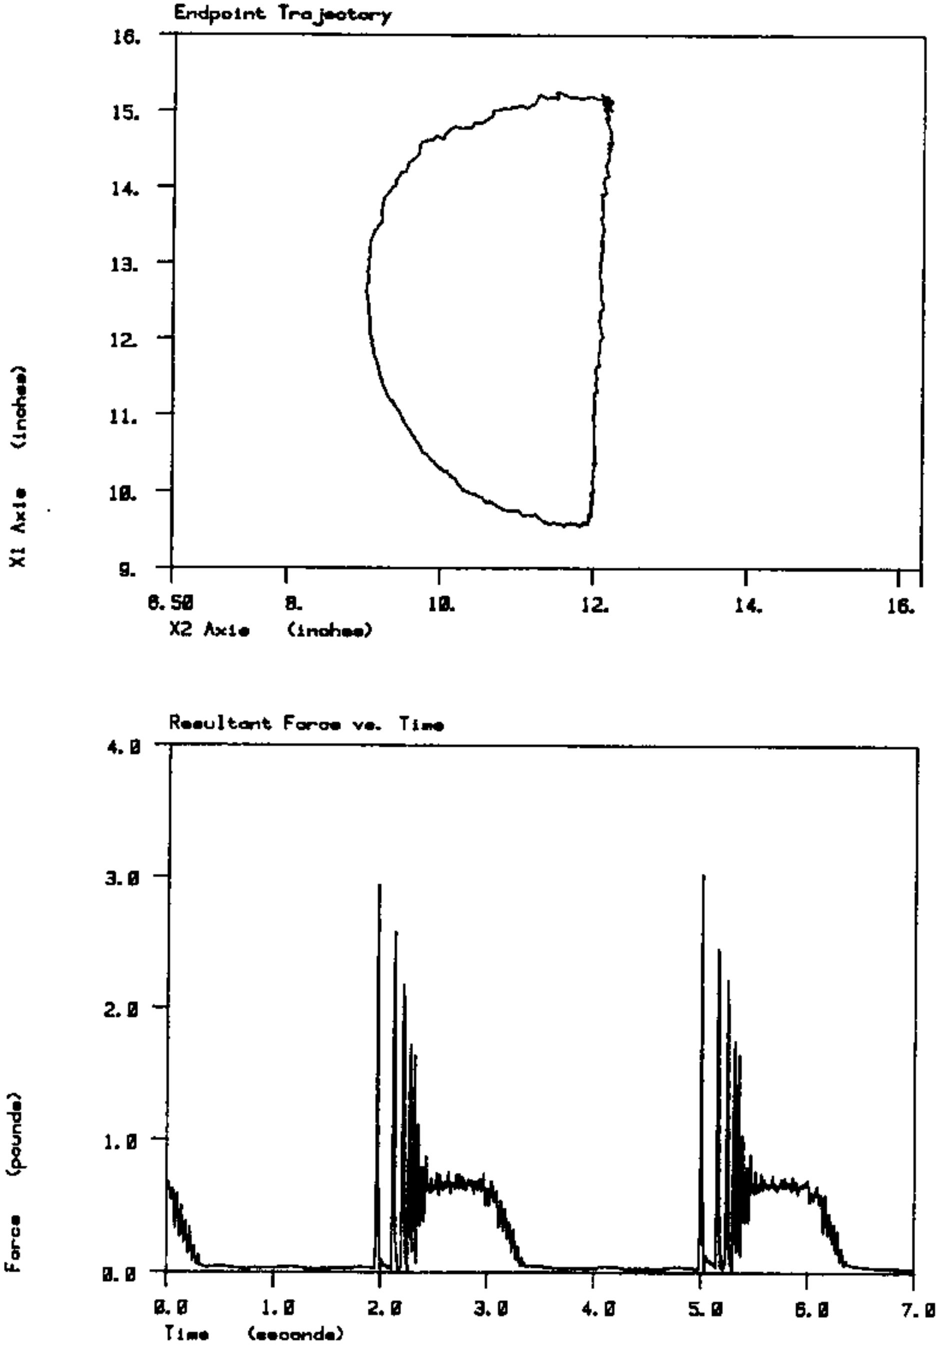

Figure 4. Actual performance of the experimental apparatus on the test task when the force feedback was disconnected. The path of the end-point is shown on the top. The time history of interface force is shown on the bottom.

Telerobotics: Neutral Buoyancy Simulation of Space Applications

David L. Akin
Assistant Professor
Space Systems Laboratory
Department of Aeronautics and Astronautics
Massachusetts Institute of Technology

ABSTRACT

Telerobotics represents the use of teleoperators and robotic devices to perform tasks in space which have to date required humans. From among the limited set of earth-based simulations of space, neutral buoyancy is selected as one of the most favorable for telerobotic operations. Two current systems under development at the Massachusetts Institute of Technology (MIT) Space Systems Laboratory are reviewed: the Multimode Proximity Operations Device, used for studying vehicle mobility, and the Beam Assembly Teleoperator, used to perform dexterous operations. A brief overview is also given of a newly-developed navigation system used to provide three-dimensional position to within an accuracy of 1-2cm, based on sequential acoustic pulses emitted from fixed locations, and received and processed on the vehicles. Test results indicate that both of these study areas have great potential for incorporation of telerobtic devices into planned space requirements.

INTRODUCTION

From the outset, space exploration involved the use of machines to represent humans in environments beyond their physical or economic reach. With the ever-expanding increase of required operations in space, the demands for operational capability far exceed the planned growth rate of human presence in orbit. To meet this need, there will be a real requirement for *telerobotics*: mechanical systems capable of many of the routine tasks which have been performed to date by humans in space. These devices, depending on need and capability, may be controlled in a master-slave teleoperation mode, or may have low-level routines capable of performing without immediate human command authority, or may operate totally autonomous from humans in space or on the ground. This field is a new and exciting extension of technologies from space, earth-based industry, and the undersea and nuclear robotics communities. Along with the required development of these technologies is the need to find an adequate basis for simulating the space environment, in order to develop the technology without the requirement for lengthy and expensive space flight experiments. In the MIT Space Systems Laboratory (SSL), much of this work has been well served by a program of neutral buoyancy simulation.

SPACE SIMULATION OVERVIEW

In comparison to the norms of undersea engineering, the space environment may be viewed as essentially benign. Primary design implications of space include the wide range of thermal environments (typically from -150° to +200°C), vacuum, and microgravity (more commonly referred to as weightlessness). The thermal and vacuum effects are essentially of interest in the design of the actual mechanisms; only the zero-gravity feature of space directly impacts the function and effectiveness of robotic devices. Some form of simulation must therefore be used to provide information on the interaction of vehicles and manipulators with the tasks to be performed in space.

Actual microgravity can be produced on earth in the cabin of an aircraft following a parabolic trajectory. This technique generally produces acceleration levels of .01-.03g, for periods of up to 25-30 seconds in transport-category aircraft. Although useful for specific tests of limited scope, this technique is not appropriate for most robotic simulations. Another form of simulation is to mount the vehicles and tasks on air-bearing pads on specially-built flat floors, providing three essentially frictionless degrees of freedom, with a requirement for providing mechanisms under computer control to simulate the last three degrees of freedom (DOF). Since some active control is required for a typical 6DOF task, another approach is to put all six degrees of freedom under direct active control, generally referred to as a motion carriage facility. A more passive approach to the unrestricted DOF problem is to place all of the simulation hardware underwater, carefully ballasted so as to be neutrally buoyant, and neutrally stable in all three rotational axes.

The neutral buoyancy environment allows a number of simulation features not available in alternative space simulation media. Unlike flat floor or motion carriage simulations, the vehicles and targets in neutral buoyancy have unlimited rotational degrees of freedom, and translational motion bounded only by the physical size of the simulator. Much of the MIT development work has been performed in the Neutral Buoyancy Simulator at NASA Marshall Space Flight Center, which is 23m in diameter and 12m deep. The neutral buoyancy environment also allows maximum flexibility in relative configurations of multiple mockups and vehicles. For example, maneuvering over a shuttle bay and around a target vehicle would be difficult in a flat floor facility, as only one motion base may be physically located on a single piece of floor. This problem is exacerbated in a motion carriage facility, as there is typically only one motion carriage, and all other mockups are limited to simple rotational degrees of freedom.

The drawbacks of neutral buoyancy include the effects of water drag, which tend to add an unrealistic element of damping. Control system development can offset this effect to some degree. In addition, hardware design must be performed to account for operation underwater: while this is not a novel problem in the context of the undersea robotics community, it is very different from the design practices typical of space flight. The problem is significantly mitigated from undersea standards, however, by the use of highly filtered tanks of fresh water, the largest of which has a maximum depth of 12m. The devices discussed in this paper therefore are designed to entirely different external pressure requirements than deep-sea submersibles.

VEHICLE MOBILITY STUDIES

Routine operations around a space station will require a substantial amount of vehicle mobility. For example, satellites to be returned to the station for refurbishment or repair must be tracked, approached, captured, returned to the station, and berthed. It would obviously be impractical to use the space shuttle orbiter for such a task; indeed, use of the orbiter would preclude revisiting the vast majority of satellites. For this reason, vehicles such as the Orbital Maneuvering Vehicle (OMV) and Orbital Transfer Vehicle (OTV) have been proposed for early development, and indeed form one of the essential pieces of space station infrastructure.

In order to maximize the productivity of the limited crew of a station, these orbital vehicles are designed to be flown unmanned. It has been proposed to use teleoperator control of these vehicles in the terminal stages of rendezvous and docking. Teleoperation can result in a variety of problems, due primarily to inaccuracies in communications, such as limited band width and time delays for signal propagation. These effects exist to varying degrees whether the operator is onboard the space station or on the ground. In addition, human factors aspects such as controllability, mental workload, and control station design for the weightless environment are other necessary research areas. An alternative approach to vehicle control is to fully automate all vehicle operations, including the docking phase. In order to further study the effectiveness of the various concepts of vehicle control, the MIT Space Systems Laboratory has developed the Multimode Proximity Operations Device (MPOD).

The multimode portion of the name refers to operating control modes of the vehicle. It is apparent that the lack of sufficient bandwidth and quality of sensory cues to the operator of a remote vehicle will result in the degradation, or total loss, of control capability. In order to examine the widest possible range of sensory information to the operator, MPOD was designed to be an optionally manned vehicle. Shown in Figure 1, the MPOD contains a cockpit for a single operator. Onboard video systems provide a forward view, transmitted to a video monitor on the control panel. In the manned mode, the operator sits internally, wearing common scuba gear supplied from an on-board air supply. In teleoperator mode, the control panel and control interfaces are removed from the vehicle, and the operator in scuba

Figure 1
Multimode Proximity Operations Device

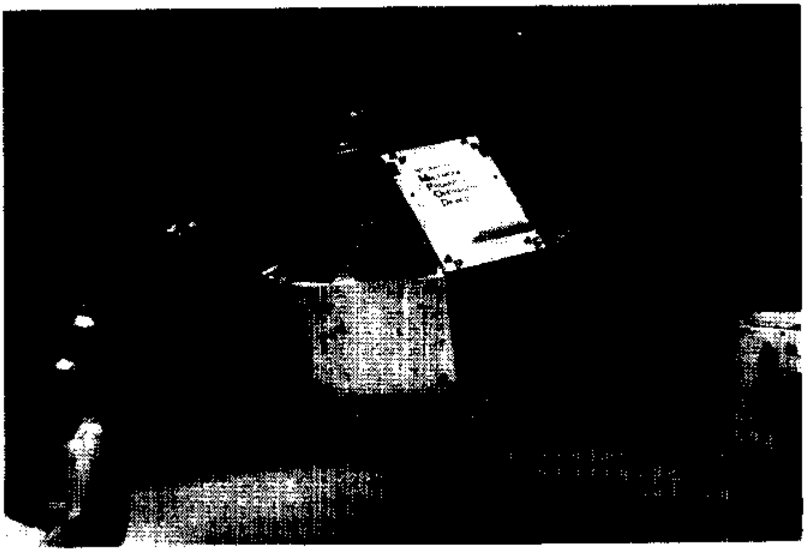

gear can control the vehicle externally using the same control panel and controls as in the onboard case. This eliminates the effects of operator interface differences on comparative runs. With the addition of a navigation system described later, MPOD is also capable of fully autonomous operations.

The external configuration of MPOD is designed specifically for the maneuvering task of interest. As a symmetrical, pseudo-spherical vehicle, the device exhibits uniform drag characteristics in any direction. With ducted propellers mounted on each of the six principal faces, all six degrees of freedom may be controlled with a set of six thrusters. Two such sets are used on MPOD, both to provide redundancy and to insure that sufficient excess thrust is available for modeling the mobility characteristics of any particular space vehicle.

For any body in motion underwater, the predominant forces are those externally applied (for example, vehicle thrust) and those due to hydrodynamic forces, such as drag. The acceleration of the vehicle is due to the sum of these forces, or

$$m\ddot{x} = T - \frac{1}{2}\rho v^2 A c_D$$

Since acceleration is the first derivative of velocity, this may be rewritten as

$$m\dot{v} = T - \frac{1}{2}\rho v^2 A c_D$$

This nonlinear first-order differential equation may be solved to find an equation for v as a function of time and the initial conditions

$$v = \frac{v_o\sqrt{a_o\beta} + a_o \tanh\left[\sqrt{a_o\beta}\,(t - t_o)\right]}{\sqrt{a_o\beta} + \beta v_o \tanh\left[\sqrt{a_o\beta}\,(t - t_o)\right]}$$

where the constants a_o and β are defined as

$$a_o = \frac{T}{m}$$

$$\beta = \frac{\rho A c_D}{2m}$$

If rest initial conditions are assumed ($v_o=0$ at $t_o=0$), this equation simplifies to

$$v = \sqrt{\frac{a_o}{\beta}} \tanh\left(\sqrt{a_o \beta}\ t\right)$$

By comparison, a vehicle in space is operating in a conservative field, with no damping at all; although this would lead to unlimited velocities, flight rules generally specify limits on velocities in the vicinity of the shuttle orbiter or the space station. Figure 2 shows a typical example of vehicle motion in space and underwater. The sample case chosen is for an existing proximity maneuvering vehicle: the manned maneuvering unit, here assumed to have a mass of 250kg and a thrust of 50N. Assuming the mass and thrust remain constant in the neutral buoyancy simulation, Figure 2 shows that the underwater version cannot replicate the maneuvering performance of the flight version, even with space velocity limits. The water drag terms slow the rate of acceleration as velocity increases, because an increasing proportion of the thrust must go to counteract water drag. This situation is aggravated by the presence of virtual mass terms, due to the acceleration of water in the proximity of the neutral buoyancy device. The effect of this is an increase in the effective inertial mass of the vehicle beyond its displacement, further decreasing the acceleration rate. It is clear from this analysis that excess

Figure 2
Manned Maneuvering Unit Performance

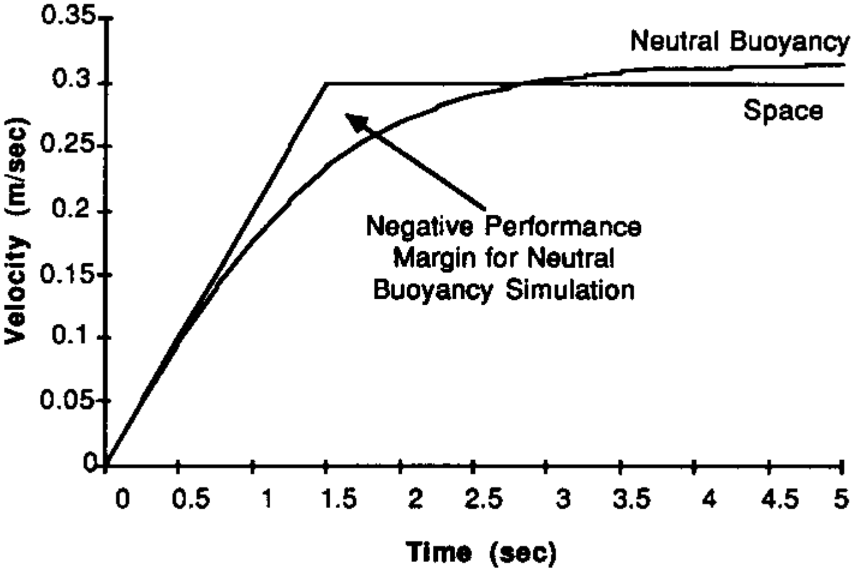

thrust capability is required for a neutral buoyancy simulation of a space vehicle. The velocity profile for the underwater vehicle must lie outside the space vehicle curve throughout the time history, so that excess authority in the propulsion system may be used to make the underwater simulation fly as the actual vehicle would in space.

In order to verify the mass and drag properties of MPOD, the vehicle was towed with a lightweight steel cable under a constant tension. The position of the cable was measured by passing it over a pulley attached to an optical encoder, which theoretically measured cable travel at the rate of 7.74×10^{-5} m/pulse. By examining the resultant data, the MPOD effective mass (including virtual mass) and drag coefficient could be directly calculated. The results indicated an inertial mass of 2919kg (as compared to a calculated displacement of 1888kg), and a drag coefficient of .674 referenced to a cross-sectional area of 1.705m^2. Figure 3 shows the MPOD performance limits, referenced to a generic Orbital Maneuvering Vehicle velocity profile. It is apparent from this figure that MPOD has sufficient thrust authority to simulate this OMV, given the proper instrumentation and control system to calculate and negate the effects of water drag.

Figure 3
MPOD Performance Characteristics

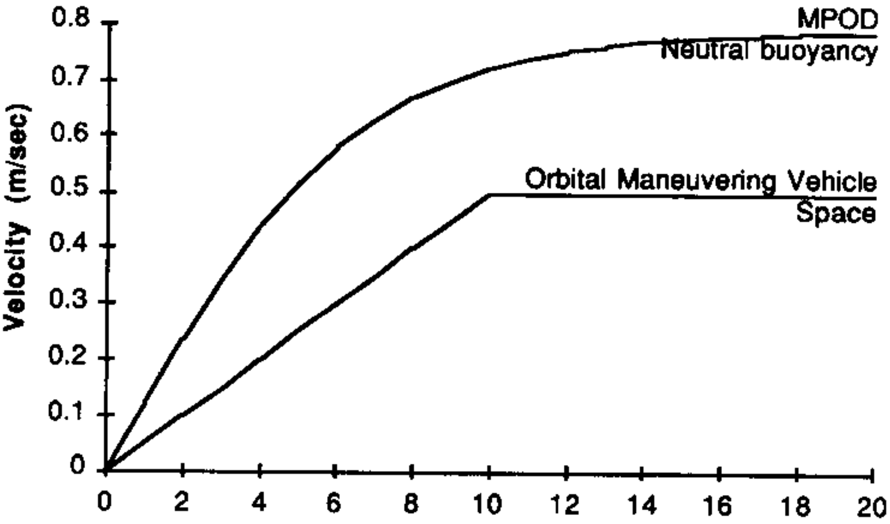

It is important in any such simulation vehicle to allow for fine-tuning of the buoyancy system; by closely adjusting the center of buoyancy to be coincident with the center of mass, water-induced torques on the vehicle will be minimized. A three-axis rate gyro system on MPOD provides data on vehicle rotational motions, for both the on-board control system, which flies the vehicle in closed-loop mode, and for a data logging computer on the surface. Numerical integration of the rate signals provides attitude angle estimates as a function of time, for quantitative measurement of crew performance. A three-axis system of orthogonal 360° pendulum inclinometers provide update information for the rate gyro integration, as well as direct data for the two rotation axes not aligned with the local gavity vector.

The current research task for MPOD involves those component activities required for maneuvering up to a free-flying module, such as a satellite, docking to it, maneuvering with it to the vicinity of a space station mockup, and berthing it into retention latches on the space station. The specific task currently being studied consists of starting from an initial position some distance in front of a satellite mockup, and maneuvering up to the target and docking with it. A probe-drogue type docking system was developed and installed on the vehicle to provide the fidelity required in the terminal phases of a docking approach.

Such a system requires some navigation data, both for quantifying performance of the operator in the onboard and remote control cases, and for obtaining information required for autonomous docking runs. To address this need, the MIT Space Systems Laboratory has developed a three-dimensional acoustic positioning system, or 3DAPS. A set of electromechanical pingers ("thumpers") mounted at known locations on the tank walls emit sequential acoustic pulses. These pulses are received by hydrophones on the vehicle; the vehicle also receives a timing pulse, denoting the start of the acoustic signal, via a fiber optic data path. The transit period of the acoustic path, timed onboard the vehicle, thus produces a linear range from emitter to receiver. Four emitters are the minimum required for non-ambiguous position determination: the eight pingers normally used in the system allow for shadowing of some of the sources based on vehicle attitude and position, and provide for error reduction through the use of the redundant data. Test results indicate that each receiver mounted on the vehicle is capable of resolving three-dimensional position to within 1-2cm. By using three or more receivers per vehicle, the vehicle attitude may also be directly measured by this system, allowing full position state feedback in the control system. There are no theoretical limitations on the number of receivers in this system, and multiple vehicles may use the system simultaneously for navigation data.

With the completion and integration of the first production 3DAPS navigation system into MPOD, the vehicle will be capable of a wide range of operation, from fully manual to fully autonomous. Detailed test results await the next opportunity for free-flight testing in the NASA Marshall Neutral Buoyancy Simulator: these tests are currently scheduled for January, 1987.

MANIPULATION STUDIES

The initial efforts of the Space Systems Lab in neutral buoyancy simulation involved the manual assembly of space structures by test subjects wearing space suits. This work produced an extensive data base on extravehicular activity (EVA) structural assembly in neutral buoyancy, and led to the Experimental Assembly of Structures in EVA (EASE) structural assembly flight experiment of shuttle mission STS 61-B in December, 1985. As an extension of the neutral buoyancy data base, it was decided in 1982 to design and construct a teleoperator, capable of the same structural assembly which up to that time had been performed manually. The device was therefore called the Beam Assembly Teleoperator, or BAT.

The Beam Assembly Teleoperator, shown in Figure 4, is designed primarily for dexterous manipulation. In its current configuration, it consists of a dexterous 5-degree of freedom manipulator arm and two specialized grappling arms, mounted on a mobility unit with six unlimited degrees of freeedom. Feedback consists of two stereo camera pairs, as well as normal manipulator and vehicle feedback parameters. Primary application of BAT has been in the area of large space structure assembly. BAT has successfully assembled four different structural connectors, including those used in the EASE and ACCESS space structures. It has also built complete structures underwater, although this effort has been limited to date by the restricted size of the MIT Alumni Swimming Pool used for development and verification testing. Due to these restrictions, much of the testing so far has involved teleoperator system studies (e.g., role and placement of video sources), satellite servicing taskboards, integrated systems checkout, and other tests not requiring extensive water volumes. Some effort has also been placed on the initial incorporation of supervisory control and subsystems automation into BAT, reducing operator workload and increasing the trend towards system autonomy.

In a typical test, BAT was started 5m from a structural component, with which it was supposed to grapple. The operator had to maneuver the vehicle to within the grapple envelope, and successfully grapple the beam. This test was performed repeatedly, with a variety of camera configurations. The results of these tests are shown in Figure 5. As can be seen, the presence of a stereo camera pair provided depth cues to the operator, and yielded the minimum time to grapple with maximum reliability (defined as percent of attempts resulting in sucessful grapples). Of particular interest is the fact that a variety of alternative views can actually decrease operator performance, as more time is spent switching between views for minimal or no net benefit. These and other tests showed that the presence of a video source fixed to the vehicle axis system was of significant help in sucessful navigation, as the transfer of coordinate frames between a moveable camera and the vehicle frame resulted in operator confusion and general degradation of performance. These tests have indicated a preferable video arrangement consisting of at least four cameras: a wide-angle stereo pair fixed to the vehicle frame, and a controllable stereo pair with close-up lenses to be used for controlling manipulator operations.

Figure 4
Beam Assembly Teleoperator

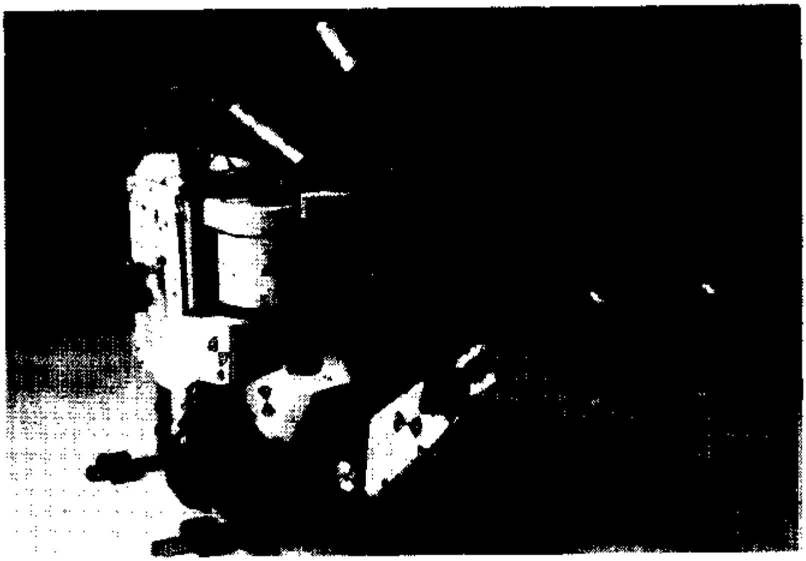

Initial results from the structural assembly task indicate that BAT is approximately an order of magnitude slower than a person in a pressure suit. However, this figure represents only the initial trials with relatively inexperienced operators. As was evidenced in the manual assembly tests, a rapid learning curve is expected to be shown when repeated assemblies have been performed. Due to the size of the baseline structure used, these tests will be performed at the NASA Marshall Space Flight Center facilities in January, 1987.

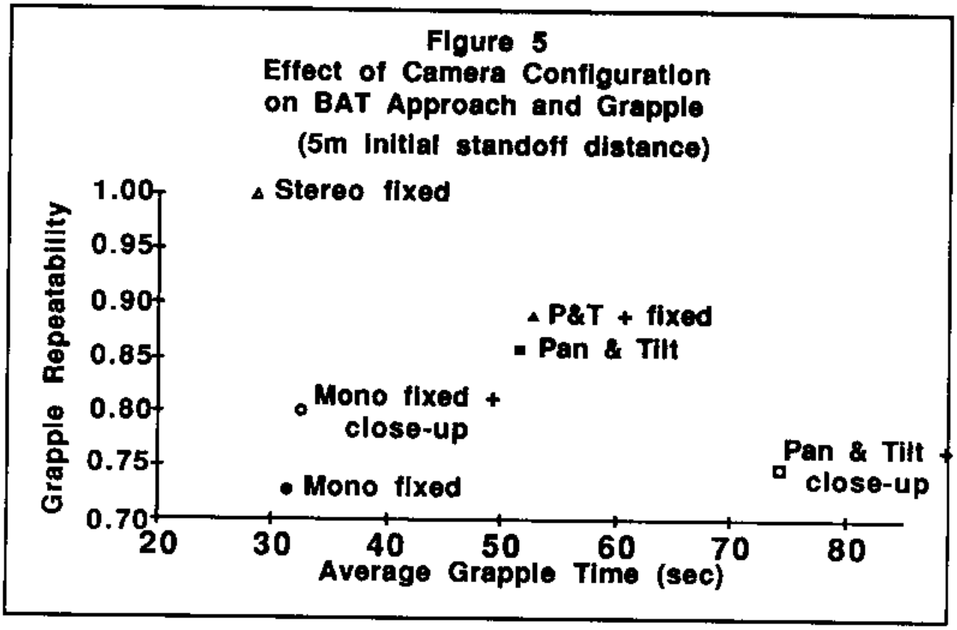

Figure 5
Effect of Camera Configuration on BAT Approach and Grapple
(5m initial standoff distance)

CONCLUSIONS

Despite its drawbacks, the neutral buoyancy environment offers significant advantages when performing integrated systems tests of space telerobotic concepts. There are many interesting parallels between undersea and space exploration: both are environments inherently hostile to unprotected humans, where human presence is expensive and limited in scope and duration. And in both environments, the progress of exploration and exploitation are dependent on identifying and using the most favorable mixtures of humans and machines.

ACKNOWLEDGEMENT

This work was sponsored by NASA Headquarters, under NASA Grant NAG-W21.

Overview of Current Vehicle Technologies in the Offshore Oil Business

F. Richard Frisbie
Senior Vice President
Ocean Systems Engineering

INTRODUCTION

The ROV (Remotely Operated Vehicle) is, in 1986, an essential part of the offshore oil business. Unreliable and limited in the mid 70s, it has grown into a complex family of systems providing a range of capabilities throughout much of the offshore industry. These capabilities address a range of tasks from coarse to sophisticated within the major areas of the industry: exploration, construction, production, and field shutdown. The present day capability has grown out of existing and projected needs and has caused, as well as being the result of, a series of technical and operational developments.

The offshore industry is driven by a need to perform, not to create. Although a great deal of offshore oriented research and development is done by the major oil companies, the ROV is a capability that has been developed by the efforts of small companies - manufacturers and operators - and not by the efforts of large oil companies, government agencies, etc. This paper will provide an overview of the present day capability by utilizing primary examples to illustrate the tasks, the equipment, and the technology that is presently being used in the industry. The paper will address what we do, and how we do it;

and will draw on a considerable amount of existing information. Technical development in our industry has no purpose except to enhance performance, capability, or cost effectiveness. The commercial ROV industry has pioneered a number of technological applications and developments but these cannot, and should not, be taken out of the context that the objective is to do the job. The objective of this paper is to show some of the capabilities of ROVs in 1986. The author has chosen to do this by selecting programs and equipment with which he is familiar. There are a number of other systems that address the requirements stated in his paper but the author was not in a position to address them.

ROV - WHAT IS IT AND ITS PURPOSE?

ROV - Remotely Operated Vehicle - is an acronym for a family, genus, and species of subsea equipment used to support underwater operations from survey through field shutdown. The term ROV represents a range of systems that covers over 30 types of conventional ROVs with totals exceeding 500 units. Within the ROV family are major categories for observation, inspection, light work, heavy work, and special applications. Each of the major categories further breaks down into specialties such as inspection - pipeline or platform, heavy work - manipulator or work package based, etc. In many instances there is overlapping or blurring of the specialities but there are clear cut requirements for each major category and the ROVs that are cost effective and productive are those which accurately address the requirements.

The ROV has evolved from a technical oddity to a reliable, productive capability fully integrated into many areas of the underwater support spectrum. This development has not been smooth and has seen more than its share of misplaced concepts and confidence. A major shortcoming has been repeated, and continuing, efforts to design and develop systems that do not address the needs of the marketplace. Out of these misguided efforts have come both unacceptably simplistic and unnecessarily complex ROVs that had no market. Inevitably these failures cast doubts on the ability of ROVs to address the requirements. In spite of this, the ROV's place in the offshore industry is secure and growing.

The purpose of the commercial ROV is to perform a task. A major mistake often made in discussing ROVs is the loss of perspective as to the purpose of the commercial ROV. Most presentations discuss topics such as technical developments, sophisticated subsystems, etc. and use words such as "dramatic, cutting, incredible, significant". These are pertinent and and often relevant but they do not address the real issue. What should be discussed are the tasks and how they should be performed. In order to perform the necessary tasks, some outstanding technical developments have taken place, and the industry can be justifiably proud of these achievements. However, the real issue is the task. In the following section I will use four examples to define the task and what it takes to perform it.

HIGH RESOLUTION PIPELINE SURVEY

Purpose: The purpose of a high resolution pipeline survey is to insure the integrity, stability, and safe operating conditions of the pipeline - for insurance, regulatory and investment purposes.

Typical tasks associated with a high resolution pipeline survey appear in Table 1.

The information collected is to be accurate with an allowable error of less than 5 percent for data and less than 2 meters for location.

The customer is paying for a report that details certain conditions of his pipeline. Although it requires a vessel, surface and subsurface positioning, a fairly unique class of ROV, sensors, sonars, and instrumentation to accomplish this task it is essential to remember that the customer is paying for the final report. In the final analysis the customer can be totally indifferent to the sophistication of the equipment and the complexity of the operation. The product is the report.

To describe only the product is to do a great disservice to the technical achievements associated with a comprehensive pipeline inspection. In 1986 special purpose units outfitted with a suite of task specific survey sensors, positioning hardware, and topside computers have evolved. High resolution bathymetric surveys are capable of generating profiles with bottom variations accurate to within centimeters and positional accuracies within a meter. ROV based bottom surveys provide the most accurate data available. The equipment utilized is sophisticated and requires considerable integration. Positioning accuracy and the need to accurately follow designated data paths further increases the complexity of the overall project. A properly defined, prepared, and executed program will provide extremely precise information. In pre-route survey the level of sophistication can range from visual to full bathymetric, although this detail is rarely required.

Post lay and periodic maintenance surveys generally incorporate the highest level of sophistication found in pipeline surveys. The ROV itself tends to be the most sophisticated in its class as it is called upon to provide a precise maneuvering capability, a stable platform, interfacing for a variety of precision sensors, multiple CCTV, 35mm photography, as well as fit into an accurate tracking envelope. Figure 1 shows a SCORPIO outfitted with sonars, pipe trackers, and positioning systems. The ROV unit may also require low acoustic and magnetic signatures so as not to affect the sensors. Representative sensors appear in Table 2.

Topside the ROV and sensor data is fed into positioning/survey computers as well as data collection and analytical computers. The final product is a series of detailed, high resolution charts showing all location points and critical data together with detailed reports by location, parameter, and defect. Figures 2 - 5 show representative report output. Integrated pipeline survey represents the most sophisticated, combination of ROV, sensors, and data output presently utilized in the ROV spectrum. ROVs have displaced the other methods for performing such projects and represent the standard by which new techniques will be evaluated.

As 1986 draws to a close work is being done on the development of improved sensors as well as on positioning accuracy. The objective is to provide more reliable and more accurate data. Steps are being taken to address final output that incorporates computer graphics to provide real time 3D pictures of the pipeline built up of a mosaic from the CCTV, trench profiler, pipeline tracker, etc. The output will also provide real time comparison against previous surveys and analysis against the acceptable limits of each parameter.

So, although the task is to generate a final report of accurate and meaningful data, the equipment used is sophisticated and covers a broad spectrum of technology. Added to this is the need to integrate the system into a surface vessel that must be capable of staying precisely on a predetermined path while determining its surface position accurately in real time within a 3 meter watch circle - all in up to a seastate 6 or 7.

DEEPWATER DRILLING SUPPORT

Purpose: The purpose of deepwater drilling support is to routinely carry out a series of inspection and work tasks that allow drilling operations to be carried out on a continuous and uninterrupted basis. Many of the tasks that have evolved from the deepwater drilling program can be applied to deepwater production support - which is discussed later. Typical Deepwater Drilling Support tasks appear in Table 3.

To understand the equipment and the capability it is essential to understand the requirements of a drilling program. The goal of a drilling program is to "make hole" in the most cost effective manner. Any specific drilling program is rarely developmental in nature - the object is to drill. However, the oil companies recognize that to drill in new areas, to drill more efficiently, to improve safety, to protect the environment, etc. they have to continually improve their drilling program. As a result new techniques and equipment are always being addressed - not because they are challenging, but to reduce the cost of "making hole" or to open up frontier areas for exploration.

Serious efforts at utilizing a work function ROV in drilling support began in 1976-77 time frame. In spite of the immaturity of the capability and the absence of operating guidelines, this application of the ROV began to expand rapidly in the 1980-82 time frame. By the end of 1982 the need had been well established, the ability to perform the work had been demonstrated - although it varied dramatically amongst operators. 1982 saw the development of an ROV system developed exclusively for drilling support. This was followed closely by the development of the HYDRA 2500, the deepest diving commercial ROV in the world and a system developed specifically for deepwater drilling support. Figure 6 shows the HYDRA 2500 vehicle, Figure 7 shows the HYDRA 2500 system installed aboard the deepwater Discoverer Seven Seas.

A number of these tasks require special tools or work packages which are either manipulator deployed or delivered and docked by the ROV due to their size and weight. Special tools have been developed for cutting and replacing guidewires, hydraulic override functions, water jetting, explosive charge operation, cable cutting, and a host of other tasks. A typical drilling support ROV includes a suite of some 8-10 special purpose tools and/or work packages. For instance, beacon placement and recovery may address beacons weighing 1200 lbs. in water that require placement 400 ft. from the stack itself. This requires heavy lift capability, accurate positioning, shackle/pin manipulation, load transfer, and visual inspection as a minimum.

It is worth looking at the removal and replacement of the AX/VX or BX rings (metal ring seals). This task can be fairly routine in shallow water using a tool deployed in the manipulators. However, in deepwater drilling the VX ring in the Lower Marine Riser Pack (LMRP), for instance, is locked in place 6 ft. up inside in the LMRP itself. Not only can it not be directly accessed, it cannot even be seen. The VX ring must be released and replaced up inside the LMRP because the lower connector location is down inside a series of spring loaded petal cones that require extremely high energy to open. The VX ring can only be lowered through the closed petal cones with difficulty and even if it could be placed on the connector face, the lowering of the LMRP would probably dislodge or damage it. As the LMRP may have 6 ft. of vertical motion during this mating operation a positively buoyant dart was inserted into the LMRP and when floated up in the LMRP it captured the VX ring. In order to replace the VX ring the ROV had to key itself to the heaving LMRP and release four separate set screws. When the set screws were released the ring could be removed by pulling the still buoyant dart down out of the LMRP. The replacement VX ring reversed this process. It can be seen that this one operation is a combination of visual, spatial, manipulative, and operational tasks of some difficulty.

What technical advances have allowed an ROV to make over 250 dives deeper than 5000 fsw and to work routinely in 7000 fsw? These developments not only benefitted deepwater drilling support but provided the basis for future deepwater production support. As a result of the importance of these developments they will be looked at in some detail. One important step was the decision to develop ROV systems to address specific areas of work. It was this move away from general purpose systems that led to the most significant increase in performance. The drilling support ROV had to be cost effective and had to do the work - reliably and consistently. The characteristics of a successful deepwater drilling support ROV are:

- o Reliability
- o Redundancy
- o Cage Deployed
- o Compact
- o Powerful
- o Work Proficient

These same characteristics are appropriate for the later section on Production Support. The development down to 3000 ft. was one on integration, testing, development, improvement, and experience. However, the jump into deepwater required some significant technical developments. A major issue was the transmission of video, signal, and power over long lengths of umbilical up to 10,000 ft. Although video/signal can be transmitted over coaxial cables in these cable lengths the type must be RG8, not RG 58 or 59. In addition crosstalk, noise/signal action, etc. becomes an increasing problem. The solution is optical fibers. In 1982 no one was operating an ROV on optical fibers and there was no data associated with optical fibers incorporated in electromechanical strength umbilicals subjected to repetitive cyclic loading over long periods of time. We chose to build a quasi conventional umbilical that had 3 x No. 8 power conductors, 3 x RG 8 coaxial cables, 3 x optical fibers plus a drain; wire all incorporated into a contra-helically wound double armored umbilical with a minimum break strength of 120,000 lbs. Figure 8 shows a section of this umbilical.

The optical fibers were put in to test the design criteria associated with cyclically loaded strength umbilicals. The system would be able to use the fibers for baseband video and telemetry once their integrity was proven. Over the past three years we have proven this critical concept and it has had a dramatic impact on umbilical performance. At present the offshore ROV industry probably leads the world in knowledge and experience associated with such technology. Over 500 dives have been made, some under near snapload conditions, with good reliability for the optical fibers.

However, while the use of optical fibers opens up new opportunity for deepwater performance it also creates a problem - how to store and operate 10000 ft., or more, of working umbilical containing optical fibers and subjected to high working loads. Conventional drum winch designs cannot protect against crushing, overheating, and highly variable loading, although there are claims to the contrary. Again, the object of the system is to work efficiently, reliably and routinely. This led to the development of a traction winch/take up drum approach - details of which appear in the literature. The traction winch allows high loads at the head sheave, and highly variable ones, to be absorbed and decreased to where the takeup drum experiences a constant predetermined load. The load is dissipated evenly over a length of some 100 ft.of umbilical that is on the traction winch. The takeup drum stores the umbilical smoothly, evenly, and at a constant packing density which prevents point loading, crushing, overheating, and premature failures.

Reliability and maintainability were addressed by the development of the first built-in diagnostic and troubleshooting system to be incorporated into a workbased ROV. Today this capability appears in many of the new ROVs, but it represented a bold step forward in 1982. The combination of new umbilical and handling system design, combined with built-in diagnostics, was instrumental in the success of the system. Much of the subsequent acceptance of ROVs for deeper applications has stemmed from the success of this deepwater capability. In 1986 drilling support in 6000 fsw is a proven capability, although far from routine and limited to a single operator. The task was to inspect a riser, or replace a VX ring, or place a beacon - seemingly straightforward ones. However, some fairly dramatic technical advances were necessary for this to take place - of which only a few of the more apparent have been addressed here. The step from 3000 fsw to 7000 fsw required a combination of evolutionary and revolutionary changes.

DEEPWATER PRODUCTION SUPPORT

Purpose: The purpose of deepwater production support is to carry out a series of inspection and work tasks associated with the installation, maintenance, and repair of the subsea and near surface components of production systems.

Deepwater drilling support has provided an extensive operational and technical training ground for workbased ROVs. The definitive nature of drilling support allowed the development of a successful system. Deepwater production has a totally separate, and often divergent, set of requirements. Deepwater production is a broad field that must be solved as there is a strong push into this area and the remote support capabilities are either immature or unproven.

Again, what are the tasks? It is here that the situation becomes complex. Representative production support will address the following subsea equipment:

- Templates/Manifolds
- Trees - Templated, clustered, satellited
- Controls - pods, manifolds
- Piping - flowlines, sales lines, manifold
- Risers
- Umbilicals

It will also be required to support the main phases:

- Installation
- Hookup/commissioning
- Long term maintenance and repair

The result is a matrix of tasks, many of which are repetitive, many of which are offshoots of drilling support, and many of which are being addressed for the first time. There are a number of critical elements to successful deepwater production support but unless the tasks are defined and the interface engineering between subsea equipment and support system successfully addressed, there will not be satisfactory long term performance. Figure 9 shows an example of ROV docking ports and valve tool slots on a subsea tree. Table 4 details Deepwater Production Support Tasks.

What are some of the tasks? The operation of a main line 20 inch ball valve on a flowline by an ROV is one. This appears technically straightforward until one realizes that 120,000 ft-lbs. of torque must be developed by a tool that is no more than 12 inches in diameter and 12 inches long, must be operated and controlled by the ROV's hydraulics, and must be flown into place, landed, and operated on a template 200 ft. x 80 ft. x 20 ft. containing up to 24 trees, 96 guideposts, two risers, and a host of cables and control lines - all with potentially less than 5 ft. of visibility and a current regime of 2.4 knots at the surface and 1 knot at the bottom.

A second task is to deploy and operate a replaceable pig trap containing up to six spherical balls that are used to clean flowlines. The ROV must transport, and install the manifold, which will weigh 300 - 400 lbs. operate a series of valves, up to 14 inches in diameter, to launch each pig separately. This must be done under the same conditions stated above, except that the valves are buried within the template and operation of the valves may be through a series of push/pull rods operated by the ROV after it has set the pig trap. It is worth noting that ROVs routinely operate multi-turn valves up to 1000 ft. lbs - in fact 11 such valves have been located, docked to, and fully operated - opened and closed - in 22 minutes in 1000 fsw. Figure 10 shows an example of this tool. In addition, guideposts have been landed and locked in position using a special tool in 3000 fsw in 30 minutes with visibility of less than 4 feet.

Many of the tasks specific to Table 4 have not been performed to date. Deepwater production, water depths greater than 1500 fsw, is restricted to a handful of projects - all of which are in the developmental stages. In order to support the tasks a new class of work based ROVs has been required. I will briefly describe one of them - the HYDRA^{TM} AT. The HYDRA^{TM} AT has been selected because of the author's familiarity with the system and the fact that the AT was developed for a known series of projects, rather than as one aimed at the market in general. The HYDRA^{TM} AT has been designed to operate in a unidirectional current regime of 3.0 knots head on and an omnidirectional current regime of 2.4 knots. Figure 11 shows the HYDRA AT. In addition the system is designed to perform both drilling and production support in up to 6000 fsw. This dual requirement imposes design constraints requiring special solutions.

The HYDRA^{TM} AT utilizes many of the lessons learned and concepts developed and proved on the HYDRA 1000s and 2500. Its lineage is easily recognized. However the HYDRA^{TM} AT is a new class of ROV with new systems and subsystems from its hydraulics, through its telemetry and control, to its handling system and control console. The HYDRA^{TM} AT generates its performance from a pair of 20 HP electrohydraulic power packs. System efficiencies have been improved by reducing losses in the overall hydraulic system by up to 40 percent. A new Siemens microprocessor based control and telemetry system was developed for the AT. The system is more efficient and reliable than its predecessors but the more visible change comes in the enhanced diagnostic, graphics, and checkoff systems.

The handling system consists of a new umbilical design and a microprocessor controlled traction winch and takeup drum. The HYDRA™ AT umbilical diameter had to be significantly less than conventional design allowed. The HYDRA 2500 had a 1.6 inch diameter and this was significantly less than the standard umbilicals at that time with equal capacity. However, the current regimes and designated working envelopes associated with the HYDRA™ AT's program required a smaller diameter. The HYDRA™ AT umbilical consists of 3 power conductors and 3 x 3 cluster of optical fiber in a double armored umbilical of 1.1 inch diameter. This is the first electromechanical strength umbilical that has no conventional copper elements for telemetry or video. The system can be operated on 3 fibers using a frequency stacking capacity but initially the system will operate on 6 fibers. Future improvements can be dramatic by the use of laser driven signals.

The umbilical is handled by a traction winch/takeup drum approximately 40 percent smaller than the HYDRA 2500 because of the umbilical size. Figure 12 shows the takeup drum. The winch uses microprocessor based controls with manual overrides. This control technique greatly simplifies the operation and reduces manning requirements. Other parts of the system, such as the control console, incorporate changes and developments as significant as those in the propulsion, control, and handling systems. Figure 13 shows the level of development typical in ROV consoles in 1986.

The HYDRA™ AT will carry out tasks listed earlier in Tables 3 and 4 described thereafter. In addition it will be called upon to carry out the replacement of control pods - weighing up to 8000 lbs - in up to 6000 fsw. The hookup and testing of inductive couplers are a new operation presently being addressed for operational feasibility. Again, deepwater production support is a new area for large scale ROV support but one that must be solved.

STRUCTURAL INTEGRITY INSPECTION SUPPORT

Purpose: The purpose of structural integrity inspection is to insure the safe operational condition of the structures over its life cycle. The specific area addressed in this example is the detailed inspection of the nodal welds.

Present day techniques require intensive precleaning of nodal welds prior to the video, photographic and NDT examination of the weld and the HAZ (Heat Affected Zone).

Therefore, the most challenging inspection and maintenance work centers around certain of the NDT requirements that are dependent on prior cleaning and visual inspection. A number of existing NDT techniques, and ones in development - such as Time of Flight Diffraction (TOFD), require some degree of surface preparation. More importantly, existing inspection requirements throughout the world require detailed visual inspection of selected welds and this requires a high degree of cleaning prior to detailed video and photography. This work has historically been the work by the inspection diver.

Some of the more critical factors that must be addressed in performing remote cleaning and inspection are listed in Table 5.

Each issue is critical in itself and, as in so many other areas, shortcomings in any one area results in an overall capability that will be of reduced practical value. Efforts to date have been carried out with ROV based systems that incorporate the cleaning, manipulation, and inspection systems in a variety of configurations. Results have been generally disappointing from a productive viewpoint, with encouragement coming from the improvements in overall performance. In addition to the items in Table 5, one must address the types of cleaning systems, the degree of cleaning/surface preparation required, and the location of the high pressure water and/or grit unit. If the decision is taken to locate the high pressure pump on the ROV then one is committed to using high pressure water exclusively. The drawback relates to the high electrical loads on the umbilical, the impact of pump vibration on the ROV, and the fact that water is sole source of cleaning energy. The other major approach is to place the cleaning equipment topside where it can be maintained and accessed and where its impact on the ROV is reduced. This method also lends itself to grit injection. The disadvantages are that the high pressure water or lower pressure water/grit must be pumped from the surface to the ROV, often through slip rings and tethers.

In recent years three ROV systems have been developed exclusively for cleaning and inspection - PROES 100 and 200 and PIC. In addition there have been three ROV systems which have been outfitted with specifically designed cleaning and inspection packages - PIONEER, SOLO, and CHALLENGER. PROES and Challenger have carried out many months of work but productivity has yet to reach commercially acceptable levels. Limitations have covered the entire range of critical factors listed in Table 5. Each of the above mentioned systems has suffered from some combination of these factors although attachment with suction cups has been effective under a wide range of conditions. Accessibility to welds is a factor that cannot be resolved in every instance. Certain nodal configurations can be neither cleaned nor inspected by anything but a diver, and even this coverage may be incomplete. In addition items left in place from yard fabrication - platforms, ladders, etc. - combined with anode placement, grout lines, etc. eliminates some percentage of nodes from reasonable access. Such limitations do not diminish the potential cost effectiveness of this remote capability but represent limitations.

An alternative to the ROV based approach is of a special purpose system designed to be set by divers or ROVs. The purpose of this system, of which DYNACLAMP is a prime example, is to perform the task and to leave delivery to existing capabilities. This approach results in a simpler, less sophisticated system while at the same time offering improved capabilities and cost effectiveness. DYNACLAMP is shown in Figures 14 and 15.

DYNACLAMP achieves good accessibility to the welds by the reduction in size and in the extended reach envelopes of the manipulators - which can handle 300 lbs. at 8 ft. reach. As the DYNACLAMP rotates a full 360 degrees about the tubular to which it is attached it is able to complete a full tubular weld, as well as portions of adjacent tubulars, without being repositioned. As DYNACLAMP is approximately 30 percent the volume of the ROVs noted earlier, the accessibility constraints are reduced.

Attachment, orientation, and repeatability are achieved by clamping to the tubular and rotating about it thereby insuring a fixed and accurate reference system - something not available to conventional based cleaning ROVs. The clamping mechanism is designed to provide a self centering capability over 6 inches of marine growth on a range of tubulars. Using an inclinometer system, as well as position feedback on the rotating ring, the position of specific weld defects can be accurately located and can be quickly accessed in later inspections. Again it is the fixed reference system of DYNACLAMP that provides this accuracy and repeatability.

Cleaning is achieved through any of the conventional methods: wire brushing, high pressure water, or grit entrained water. DYNACLAMP was designed to utilize these methods as different areas of the world tend to favor different methods. DYNACLAMP utilize a conventional 20 HP power pack to provide sufficient hydraulic energy to support any anticipated requirements. The energy budget of DYNACLAMP is satisfied with 3 HP except for wire brushing. The transition from one cleaning method to another is straightforward as the unique elements - high pressure pumps, grit hoppers, etc. - are mounted topside where they can be maintained.

DYNACLAMP's utilization of existing power packs, telemetry and control, and video reduces the technological risk of the system. The areas that required new developments were: attachment/rotation and the manipulators. In order to achieve cost effectiveness the system must be able to operate routinely for 16 hours per day throughout the operating season. Up to this time manipulators were used for only a fraction of the operating time of the ROV. On cleaning and inspection manipulators would be operating close to 100 percent of the time. Increased strength, ruggedness, reliability, and precision had to be achieved. The manipulators incorporated rate and spatial functions with the capability for teach-learn and task programmability.

Remote cleaning and inspection is a complex series of tasks that represent the most difficult matrix undertaken by remote work systems. The development and successes of the ROV based systems to date may not suggest significant progress but this is misleading. The difficulty in performing remote work has hinged on two basic needs: the need to be able to define the task specifically and unambigiously as it applies to an ROV; and the need to break the task down into basic operations that reflect current ROV technology. The first generation cleaning and inspection ROVs helped define the task as they apply to an ROV and now improvements and parallel developments can take place.

CONCLUSIONS

ROVs are indispensable tools in the commercial offshore business of the 80's and the trend is towards an increase in ROV capabilities and participation. The offshore industry is driven by performance and cost per unit of performance and these two factors shape the ROV and the services it performs.

In 1986 ROVs have developed solid track records in virtually every phase of the offshore industry, with pipeline and platform inspections and drilling support dominated by their performance. Deepwater production support is a developing area that favors the ROV because of its depth independence and the ability of technical groups to design and interface subsea equipment to be compatible with work based ROVs. Remote cleaning and inspection's future is one of growth once productivity and reliability has been demonstrated.

Over the past ten years the technological and operational developments of ROVs have led from a general concept to a capability that spans the entire spectrum of underwater involvement. In some cases the commercial ROVs incorporate design and equipment that is at the leading edge of subsea technology. This is the result of necessity rather than intellectual curiosity. Commercial ROVs are problem solvers, nothing else.

HIGH RESOLUTION PIPELINE SURVEY SENSORS

TABLE 1

a high quality video tape that includes up to 270 degrees of the pipe simultaneously with precise and accurate location data;

a continuous CP chart along the length of the pipeline with precise and accurate location data;

depth of burial charts;

trench profiling;

span profiling;

debris and damage assessment.

PIPELINE SURVEY SENSORS

TABLE 2

Precision depth and acoustic altitude bathymetry

Pipetracking system including height above pipe

Lateral profiling scanning sonars

Potential gradient CP monitoring sensors

Obstacle avoidance sonar

Responder/transponder for underwater tracking

Sidescan sonar

Subbottom profiling sonar

Temperature and salinity probes

Current meter

DEEPWATER DRILLING SUPPORT TASKS

TABLE 3

Survey and inspect drilling site

Debris clearance

Place and recover beacons

Inspect and visually guide TGB, PGB, BOP

Establish guidebase orientation

Guide drill bit reentry and casing operations

Observation of riser and template leveling/verticality

Establish and replace guidewires

Gas sampling

Observe cementing and drilling operations

Verify and test H-4 latching

Inspect BOP stack and riser

AX/VX ring removal and replacement

Recover dropped items of equipment

Hydraulic quick connect override capabilities

Water jetting

Explosive charge placement

Cable and soft line cutting

Current meter readings

DEEPWATER PRODUCTION SUPPORT TASKS

TABLE 4
(All Table 3 Tasks)

Overside/operates of valves from 1/2 inch to 20 inch

Operation of chokes

Replacement of choke and valve modules

Control pod replacement

Flowline/pipeline pull in and connection support

Control lines/umbilical installation/replacement

Anode replacement

Riser connect/disconnect

Mooring installation support

CRITICAL CLEANING AND INSPECTION FACTORS

TABLE 5

Accessibility to welds

Attachment capability

Orientation and repeatability

Cleaning and inspection performance

Manipulator requirements

Technological limitations/developments

Complexity and reliability

Operator training

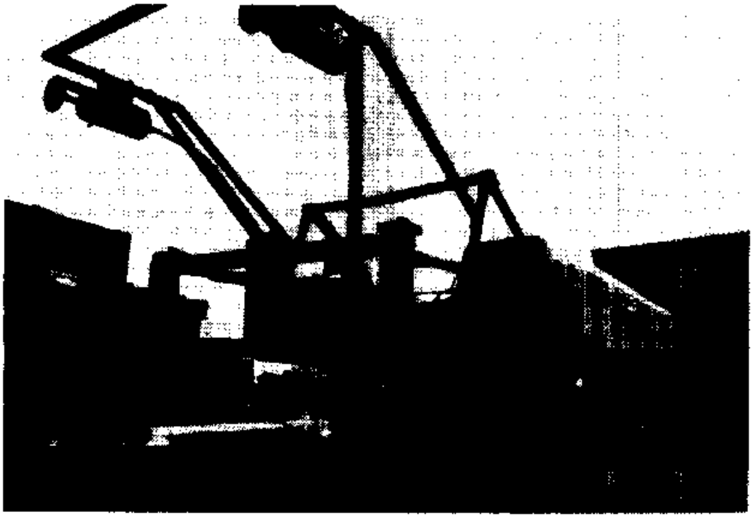

Figure 1
SCORPIO Outfitted with Sensor,Pipe Trackers
and Positioning Systems

Figure 2
Pipe Detail

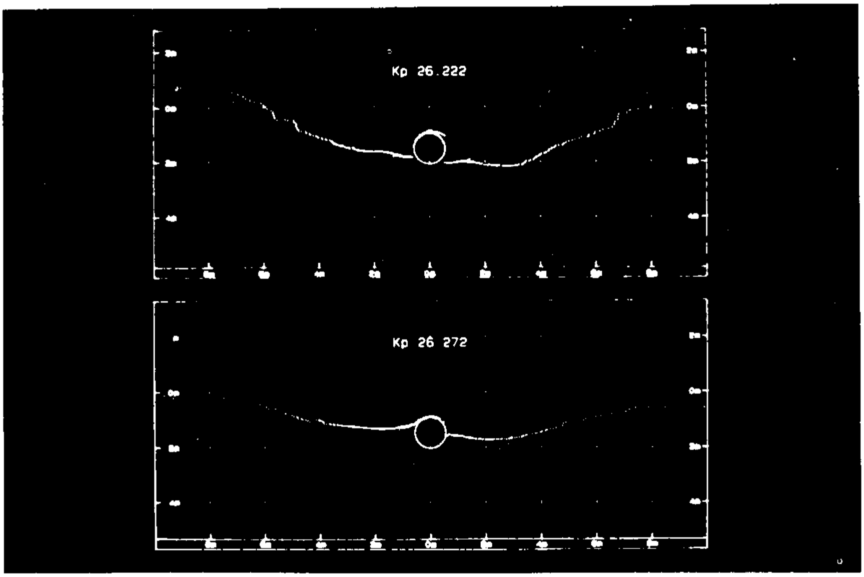

Figure 3
Trench Profiling Output

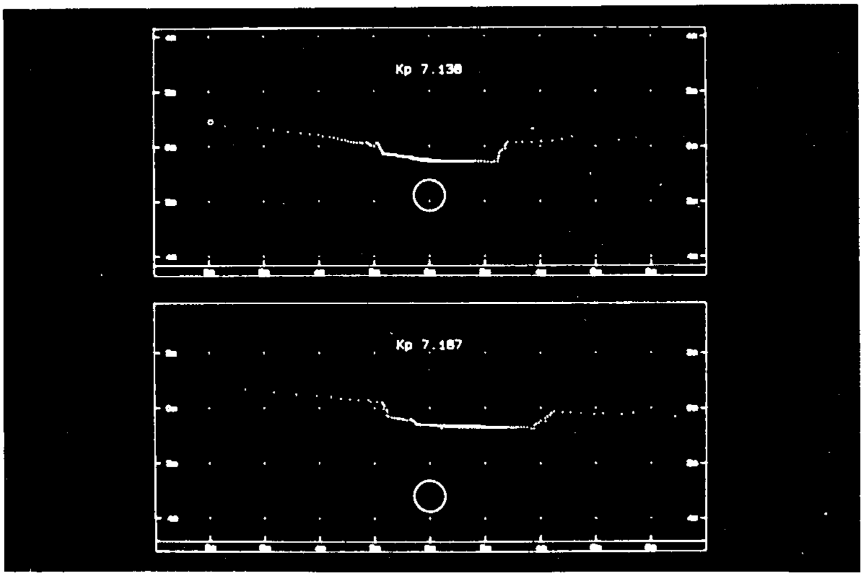

Figure 4
Depth of Burial Output

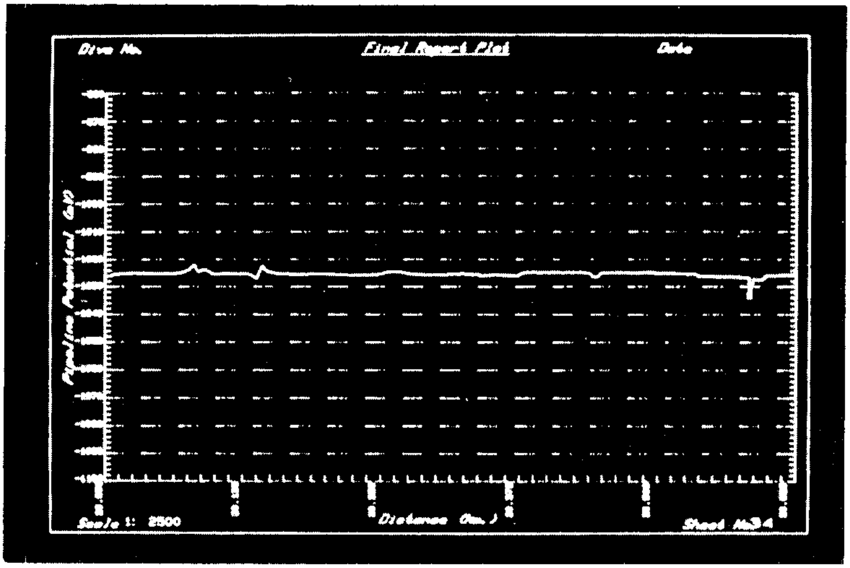

Figure 5
Cathodic Potential Output

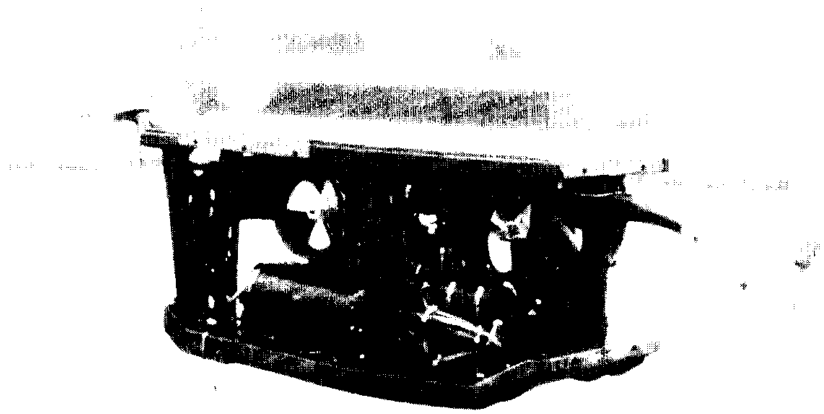

Figure 6
HYDRA 2500

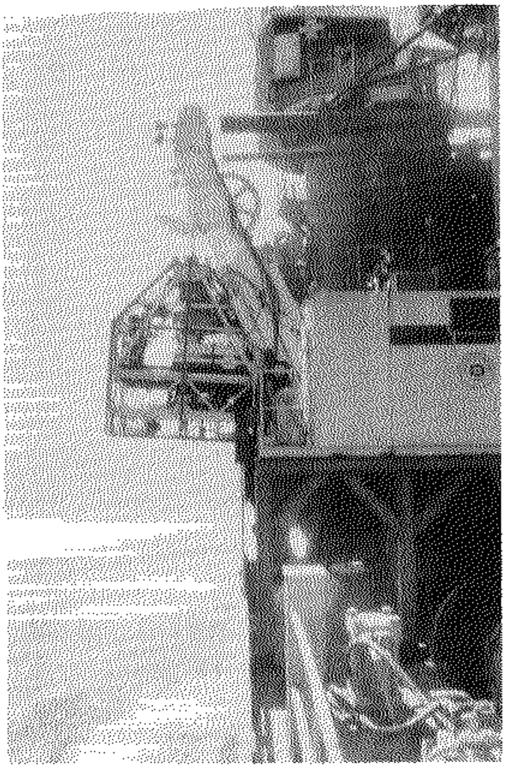

Figure 7
HYDRA 2500 Installed Onboard
the Discoverer Seven Seas

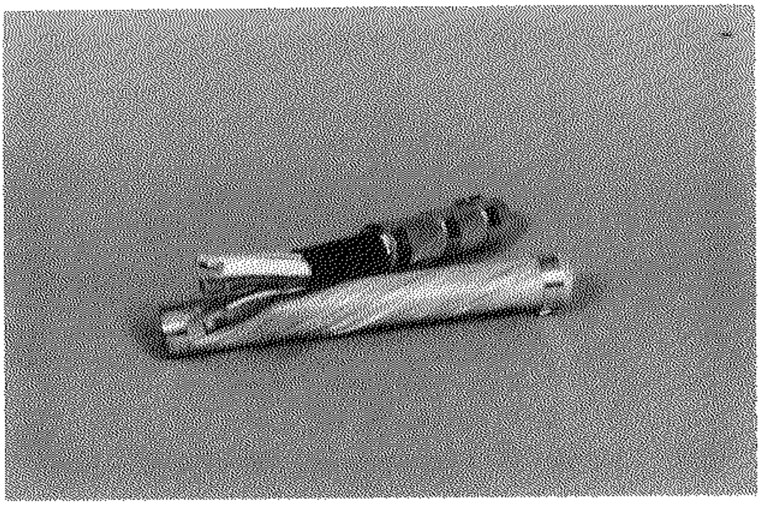

Figure 8
Section of Umbilical

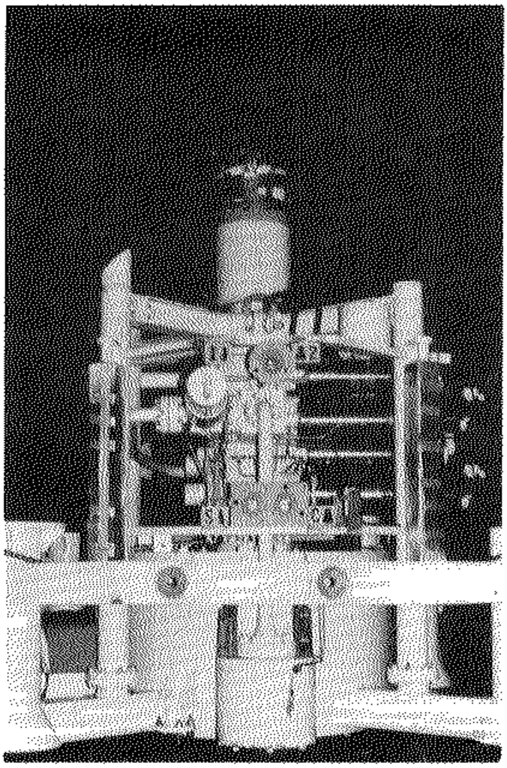

Figure 9
ROV Docking Ports and Valve Tool

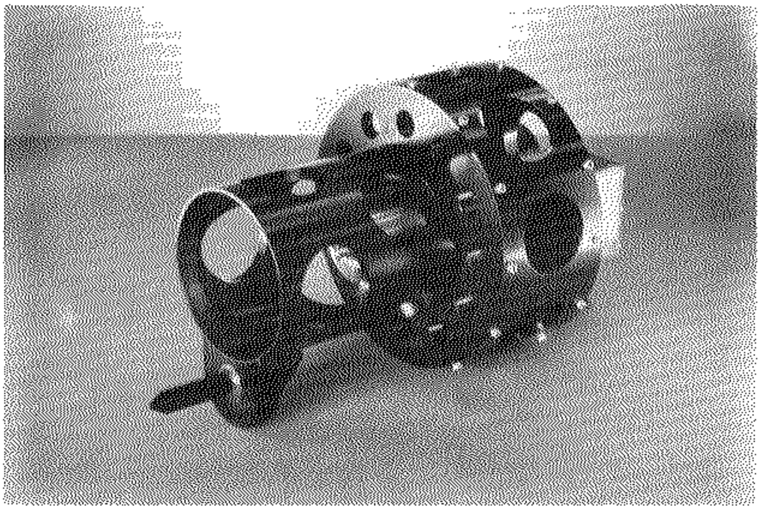

Figure 10
Hydraulic Tool for Multi-turn Valves

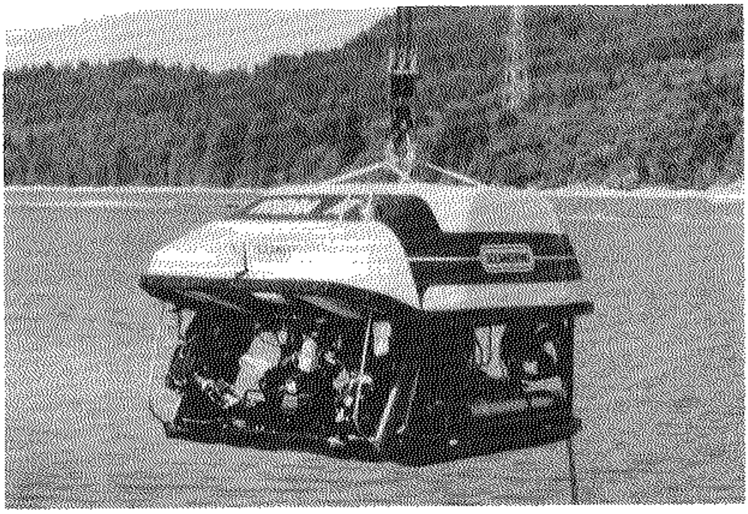

Figure 11
HYDRA AT

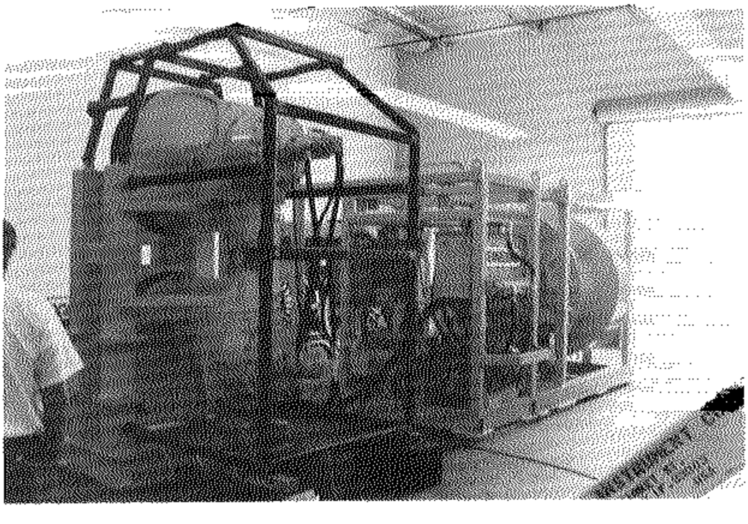

Figure 12
HYDRA AT Winch

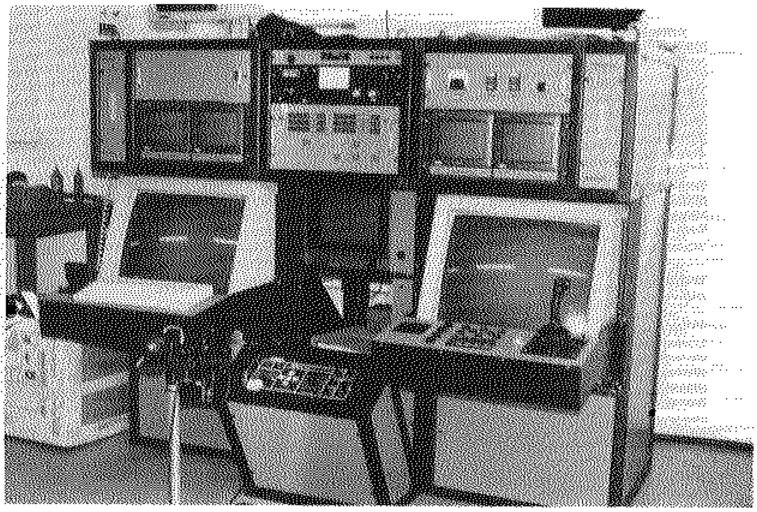

Figure 13
HYDRA AT Console

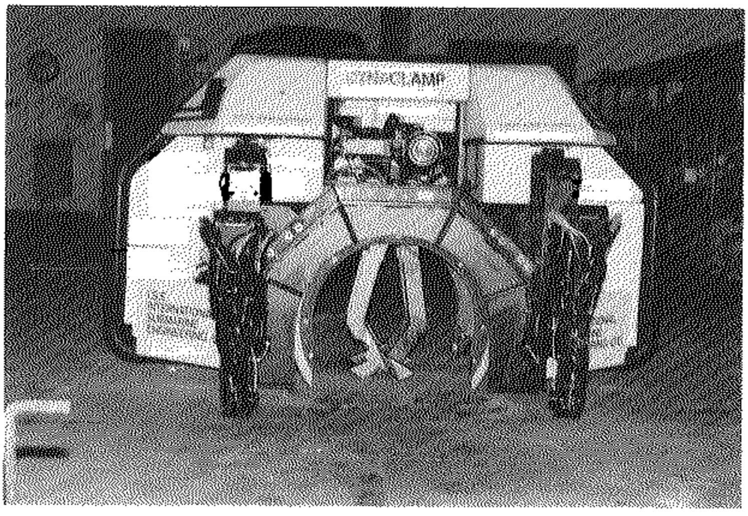

Figure 14
DYNACLAMP

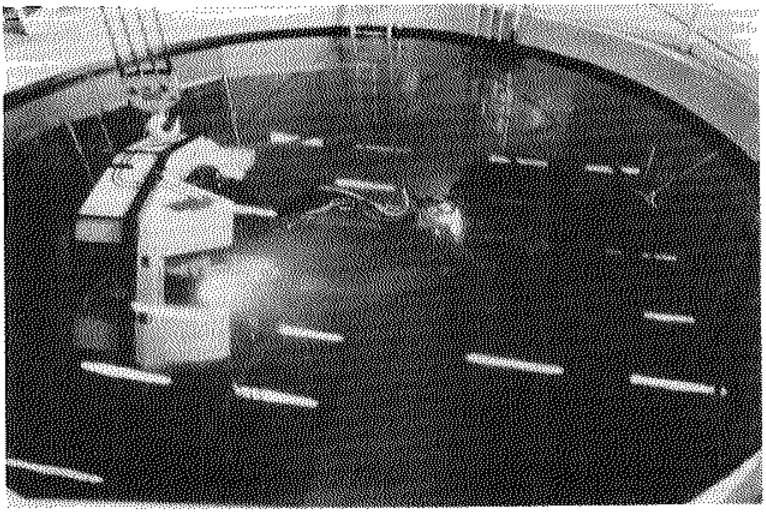

Figure 15
DYNACLAMP

Teleoperation for Marine Biology and Oceanography Research

Applications of Unmanned Submersibles in Benthic Marine Ecological Research

Kenneth P. Sebens
Associate Professor
Marine Science Center
Northeastern University

INTRODUCTION

Unmanned submersibles, or remotely operated vehicles (ROVs), are undergoing a surge in variety and versatility. At the same time, certain models are becoming so inexpensive that they can be purchased by individual researchers. Some tasks that were traditionally accomplished by dredging, camera sled or manned submersible may be better accomplished using unmanned vehicles. They offer the advantages of real-time video image feedback, multiple-task capability, ease of transportation and launching, and low cost. Even in shallow water, in cases where unsafe conditions prevent SCUBA diving (i.e. severely polluted areas), unmanned submersibles have already become the system of choice. Fortunately, the development of research functions for unmanned submersibles follows similar developments in diver carried tools and in manned submersible apparatus and can thus draw on several decades of previous experience.

Ecologists, oceanographers,marine biologists, and limnologists working on benthic (bottom-dwelling) plants and animals are severely limited at depths below those safely explored by SCUBA divers. Even in deep water where divers can still function efficiently (20-30m), they are restricted to less than an hour per day of actual bottom time by the need to avoid decompression problems. Underwater habitats improve on this situation by allowing divers to saturate, thus multiplying their bottom time by an order of magnitude or more at a given depth. However, habitats are very expensive to maintain and are generally not mobile. At present, there are no research-oriented saturation systems operating in North America, following the retirement of N.O.A.A.'s *Hydrolab* in 1985. Manned submersibles are currently the best alternative system for research in water from 30 to several thousand meters depth. They allow a bottom time of three to eight hours per day, sometimes more, depending on the battery life and recharging time of the system. Submarines, however, cannot carry out some of the simplest tasks undertaken by divers or, if they can, it takes them many times longer to accomplish the same goal.

Before manned and unmanned submersibles were routinely available, most benthic ecological research was conducted from surface ships by dredging, bottom grabs, net tows, and towed camera systems. This provided many specimens of deep-sea life but was unsatisfactory for quantitative analysis of faunal distribution or for behavioral studies of important species. Towed camera sleds, on or near the bottom, were among the first unmanned submersible vehicles and are still in general use; they have improved benthic surveys significantly (Machan and Fedra 1975, Holme and Barrett 1977, Uzmann et al. 1977, Hessler and Smithey 1984, Holme 1985, Smithey and Hessler 1985). Such vehicles allow quantitative assessment of macrofaunal populations over large transects of sea bottom, often with some capability for animal size estimation. These vehicles cannot be classified as 'remotely operated' since they provide little or no data feedback to an operator and their operation is not 'controlled' except by the speed and heading of the surface ship. Towed bottom sleds and modified dredges have also been used for quantitative collections of animals living on (epifauna) and within (infauna) the substratum for the past 20 years or more.

In this paper I review the capabilities of several types of unmanned submersible, noting in each case how these vehicles are, or can be, adapted to benthic sampling and experimentation. The major tasks that benthic biologists routinely undertake in shallow water will also be outlined as a 'statement of needs' for deep water research. Capabilities of SCUBA divers, manned submersibles, and unmanned submersibles will be compared to determine how the apparatus and techniques developed for the first two groups can be adapted for use by unmanned vehicles.

UNMANNED VEHICLES: WHAT IS AVAILABLE?

There are approximately seven types of unmanned submersible vehicles (Table 1.), of which there are over 900 in existence (Busby Assoc., 1985) with depth ranges to over 6,000 meters. The first group is the tethered free swimming vehicles;

Table 1.

ROV CATEGORY	**NUMBER (1985) VEHICLES**	**DEPTH RANGE (M)**
TETHERED FREE SWIMMING	740	50-3048, most 300-400
TOWED MIDWATER	23	100-6000, most > 4000
TOWED BOTTOM	27	60-1000, most 100-200
AUTONOMOUS UNTETHERED	24	90-6000, evenly distributed
BOTTOM CRAWLING	25	5-1370, most ≤ 100
STRUCTURE RELIANT	17	35-610, most ≤ 100
HYBRIDS	32	50-6000, evenly distributed

Source: Undersea Vehicle Directory, (1985) Busby Assoc. Inc.

in 1985 there were 740 of these covering a very broad depth range. This is certainly the group that has proliferated the most rapidly and has seen the most utility in recent years. Tethered free swimming vehicles probably have the most general promise for benthic research as well. They are excellent tools for photographic surveys, but moderate to poor for benthic sampling or for any kind of *in situ* experimentation. A small rack of experiments might be placed by such a vehicle, although not much of that has been attempted yet in field research.

The proliferation of new, extremely inexpensive, ROVs (less than $20,000) has just begun in the past year or two. Such vehicles are likely to be increasingly useful to individual researchers, especially those working in the hundred meter depth range, even if they can be used only for video or photographic surveys. These vehicles will become regular working tools for benthic marine ecologists over the next decade. The inexpensive ROVs are extremely limited at this point as to what tasks they can accomplish. They do not have the manipulation capabilities that, for instance, the commercial ROVs do. These new units are essentially photographic tools but they represent a major improvement in this area.

All other categories of unmanned vehicles combined totalled less than 150 vehicles in 1985. The second group, towed midwater vehicles, is also appropriate for certain photographic surveys, although less useable, if at all, for other needs. Towed bottom vehicles, the third group, are useful for photographic surveys as well, although they might be able to take certain types of bottom samples at the same time. An epibenthic sled or modified dredge (Sanders 1965), with quantitative photographic capability or even with video feedback to check habitat type, could collect voucher specimens of the same animals that appear in the photographs.

The fourth group, untethered autonomous vehicles, can become the system of choice for large scale photographic surveys because they can cover large areas of bottom, a meter or two off the substratum, taking timed photographs. A survey could probably cover a much larger area by using such a vehicle than using a manned submersible. These vehicles would also take less ship support, space and crew than do the manned vehicles.

Table 2.

	BENTHIC SAMPLING UTILITY		
ROV CATEGORY	**PHOTO SURVEYS**	**BOTTOM SAMPLES**	**SIMPLE EXPERIMENTS**
TETHERED FREE SWIMMING	HIGH	MOD	MOD
TOWED MIDWATER	MOD	NO	NO
TOWED BOTTOM	HIGH	LOW	NO
BOTTOM CRAWLING	MOD	MOD	MOD
STRUCTURALLY RELIANT	NO	MOD*	MOD*
AUTONOMOUS UNTETHERED	MOD	NO	NO
HYBRIDS	HIGH	MOD	MOD

* If structure (e.g. guide rail) were emplaced for this use

Bottom crawling vehicles, the fifth group in Figure 1,may have some potential for benthic research because of their ability to hold position rigidly. I have noticed, using manned submersibles, that if the submersible can sit in one place on the bottom, sampling is easier. If the submarine is moving back and forth in current or surge, it is much harder to take samples. There is thus some utility for a crawling ROV to take photographs and samples and also to place fairly large experiments and instruments. A large and heavy package could be carried on a bottom crawling vehicle in a way that might not be possible for a free swimming one.

Structurally reliant vehicles, the sixth group, may also have some utility in this respect. A vehicle able to follow a rail or cable laid on the bottom could move between experiment sites with great accuracy and could relocate fixed quadrats for long-term monitoring without stirring up the bottom. If a site must be revisited frequently, the investigator could put down a cable or a rail that would allow the vehicle to be oriented very exactly and to travel back and forth to take photographs of the same populations each time. There might, in fact, be a use for a vehicle that could descend, locate and attach to a structure *via* a sonic beacon, then follow a predetermined route on the bottom. Benthic communities could thus be monitored with great precision year after year or even on a monthly basis if necessary. There are also a few hybrid vehicles that have characteristics of one or more of the other groups. Hybrid vehicles of the future may be small, semi autonomous vehicles that are lowered into place in a tethered cage then work on their own until they return to the cage to relay information to the surface, to recharge, or to return to the surface ship.

BENTHIC MARINE ECOLOGY: NEEDS FOR RESEARCH

The needs of benthic ecologists and marine biologists conducting resource inventories, bottom surveys, population studies, and experiments are varied. These tasks are components of fisheries research, environmental impact studies as well as basic research on the makeup of benthic communities and on the roles of their component species. It is useful to compare how these needs fit the capabilities of a SCUBA diver, a manned submersible and an unmanned vehicle.

Mapping. Benthic researchers must have a detailed map of the site that they are working in. Most current work is done without accurate maps except for those constructed from navigational charts or from depth readings taken from a surface ship on site. On the smaller scale, the location of particular rock ledges, hydrothermal vents, or other surface features may be critical. Photographic mosaics on the order of a few hundred square meters are necessary to achieve this level of detail, comprising many individual photographs each covering one to a few square meters area. Divers can do some of this mapping directly in shallow water. Manned submersibles can produce such maps as well (Smithey and

Table 3. CURRENT CAPABILITIES FOR BENTHIC SAMPLING

TASK	**DIVER**	**SUB**	**ROV**
MAPPING SITES	SOME	SOME	NO*
RELOCATE STATIONS	SOME	YES	SOME
MACROFAUNA/EPIFAUNA:			
VIDEO SURVEYS	YES	YES	YES
VIDEO(with scale)	NO	YES	NO*
35 MM QUADRATS	YES	YES	SOME
POPULATION SURVEYS:			
COLLECTIONS (LARGE N)	YES	SOME	NO
INFAUNA:			
CORE SAMPLES (MULTIPLE)	YES	SOME	NO
35 MM PHOTOS (SUBSTRATE INTERFACE)	YES	SOME	NO*
EPIFAUNA:			
SUCTION SAMPLES (MULTIPLE)	YES	SOME	NO
EXPERIMENTATION:			
AFFIX MARKERS TO SUBSTRATE	YES	POOR	NO
EMPLACE INSTRUMENT/ EXPERIMENT PACKAGES	YES	SOME	NO*
MANIPULATION (*IN SITU*)	YES	POOR	NO

*can be added with existing technology

Hessler 1985). Unmanned vehicles, so far, are probably the least useable for this purpose. With more accurate positioning systems, unmanned vehicles could be used to photograph a series of parallel transects, at a fixed height off the bottom, thus producing an accurate small scale map, possibly without a need for surface guidance once the area is chosen.

Positioning systems: permanent bottom stations. Many benthic research projects depend on the ability to relocate study sites and resample the same populations or communities at regular intervals (Lundälv 1971, Sebens 1985ab, Witman 1985, Smithey and Hessler 1985). In some cases, this means relocating and photographing marked quadrats less than one square meter in area. Accurately positioning an underwater vehicle, either during a site visit or on a subsequent visit, is a difficult problem. Relocating stations without surface buoys can be accomplished by SCUBA divers using hand-held ultrasonic locaters, or by manned submersibles with ultrasonic receivers on board. There is potential for this type of positioning using unmanned submersibles as well. Finding a station from the surface, and then letting the vehicle home in on it independently using a small ultrasonic receiver such as those developed for divers, would allow an unmanned vehicle to perform this operation almost as well as a diver. An accurate positioning system might include one or more sonic transponders and a rail, or cable 'highway' system that could be followed by a vehicle once it had found the general research area. Stops along the line, or grid of lines, could be programmed into the vehicle's memory such that photographs or bottom samples were taken, experiments placed or collected, and instruments read, deployed, or retrieved with great spatial repeatability.

Surveys of benthic communities. Large scale surveys consist of photographs and/or bottom samples collected along a transect less than one, and up to several, meters wide by hundreds of meters long. Such surveys can be carried out using towed camera sleds (Uzmann *et al.* 1977, Holme 1985, Southward and Nicholson 1985) although such systems lack direct feedback to the researchers on the surface. Similar surveys by manned or unmanned submersibles allow scientists to investigate habitats, bottom features, or specific animals of interest once they are located. Video surveys provide inexpensive coverage of a large surface area, although the image lacks the resolution of 35mm still photographs and thus video transects are useful for quantification of only the larger macrofauna and physical features. Still photographs taken along a transect or grid provide much better resolution, often allowing identification, to species, of animals or plants a few millimeters in diameter (Green 1980, George 1980). Video systems for unmanned vehicles are limited by the need to observe either on a large scale (wide angle) or close to the substratum (macro). When a video system is in use for one, it cannot be used for the other. An ideal configuration would be to have two video systems, a top-mounted wide angle video providing orientation to the pilot, and a side or front mounted close-up system for benthic photography.

Sediment interface photography is accomplished by mounting a water-filled box with a glass wall in front of a camera and pushing the box edge into the bottom (Rhodes and Young 1970, Stewart, this volume). This is now accomplished remotely from surface vessels, or can be done using the manipulator arm of a manned submersible to hold and position the apparatus. This technique provides photographs of sediment grain distribution and layering in the top few centimeters of substratum, as well as images of some of the common infaunal animals or their burrows. Unmanned submersibles could be outfitted with a similar apparatus, which would be positioned using the video system and movement of the whole vehicle. Penetration into the substratum could be achieved by rapid descent from a meter or so above the bottom with the box on the front of the vehicle hitting the bottom first, or by installing a hydraulic device that pushes the box edge downward on command from the surface. The great advantage of positioning the camera by submersible is that the operator can choose appropriate habitat types, avoid obstacles, and sample within aggregations of particular dominant species.

Fixed Quadrat Photography: Photography of areas less than one square meter provides fine resolution, and can be used for permanent quadrats that are to be revisited. Data on an individual animal's growth rate, recruitment in benthic populations, and interactions between sessile species can thus be obtained (Lundälv 1971, Sebens 1984, 1985ab, Witman 1985 1987). Fixed quadrats are relatively easy to mark, relocate, and photograph during SCUBA dives. It is more of a problem when done from a manned submersible because of submersible movement in surge or current and because thrusters or the vehicle's undercarriage stir up sediment during positioning. This task is made easier when the submersible can become heavy, sit on the bottom, and then use the manipulator arm to move a camera over the fixed quadrat. Divers generally use underwater epoxy to affix quadrat markers to rock; submersibles mark quadrats more easily by placing weighted markers on the bottom, at depths where the markers cannot be moved even by storm surge.

A coupled system of video and 35mm cameras taking the same image is now available, allowing a submersible operator to observe, in the video, the exact photograph that the 35mm camera will take (Holme 1985). Higher resolution can be achieved by using a 2 x 2" format camera (e.g. *Hasselblad* system developed for the *Johnson Sea-Link*). These cameras, however, are bulky and may be more difficult to adapt to the smaller unmanned vehicles. Stereo photography has been used by divers (Lundälv 1971, Torlegard and Lundälv 1974, Green 1980) and from submersibles (Smithey and Hessler 1985) to provide additional information on the vertical relief of bottom features and sessile animals. The basic problem in fixed quadrat studies is how to find and photograph large numbers of previously marked quadrats rapidly, using a camera system that gives the greatest possible resolution.

The problem of scale: Until recently, there was no way to extract quantitative size and population density information from videos. A manipulator arm held a ruler or a quadrat in the field of view, providing only a very crude estimate of size or surface area or the bottom area was calculated from known camera height and angle. None of these methods worked well on irregular surfaces. As a partial solution to this problem, a parallel double laser system mounted above the video camera on the *Johnson Sea Link* was developed by the Harbor Branch Oceanographic Institue engineering group in response to the needs expressed by researchers working on benthic studies in the Gulf of Maine (N.O.A.A./N.U.R.P. project) and was first put into use in summer 1985. The paired laser beams provide a reference scale independent of distance or angle of the surface they hit. When the camera pans back and forth between near and far objects, the scale automatically moves with it. With this laser system, only animals near the center of the video frame can be measured accurately; it would be better to have scales spread out over the whole frame. A typical video image of an irregular bottom is a landscape where some animals are in the foreground, others are a bit further back, and there is a large background area. It should be possible, from a video image, to measure animals or plants wherever they occur using a grid of scale markers.

Still photographs (e.g. 35mm) of irregular terrain present the same problem. The paired laser scale is visible in still photographs as well, so a combined video and still camera system using the same lens could have the same scale in both types of image. Multiple parallel lasers, possibly using a beam-splitter to provide several pairs of beams from one set of lasers, might provide a grid of comparators covering the photograph. A sonic ranging system, already in use in automatic focusing cameras, for example, could scan the field of view in the photograph and store the data in a form that could later be used to generate a three dimensional grid onto which the image could be projected. For example, a computer program could then be used to integrate surface area in a photograph of a rocky bottom; population density per unit area on irregular terrain could then be calculated. That calculation is essentially impossible at present except when photographing a very flat bottom.

Time-lapse photography: Various time-lapse photography systems have been used to study the movement of benthic animals, and the behavior of predators and scavengers (Fedra and Machan 1979, Christie *et al*. 1985, Witman 1985). Such systems can be emplaced from the surface or from a submersible and left for days to months before retrieval. Submersible emplacement is preferable to surface deployment because the investigator can choose the appropriate habitat type and orientation of the camera, just as can be done by SCUBA divers. Unmanned submersibles, even in their present configurations, could be used to place time-lapse camera packages on the bottom and retrieve them later. A future development might be a dedicated unmanned vehicle that can carry several such camera packages, find a marked site, place the cameras at predetermined stations along a cable or rail, then return to the surface until needed for future retrieval.

Collections for population studies: Large scale population surveys, for which researchers must collect or measure large numbers of animals, are best done by divers in the depth range where they can work safely. It is harder to do this using a submersible if one must actually go along and pick up individual large animals; it takes a long time per individual. Such a collection is almost impossible with an unmanned vehicle at this point except for certain kinds of animals that can be collected in a suction device.

One major problem is that the manipulator dexterity is still extremely crude on most of the manned submersibles and the few unmanned ones that have arms. It is better on the manned submersibles, partly because the operator is sitting right there, but it is still amazingly time-consuming compared to what can be done in SCUBA diving research. Macrofauna samples for population studies are best collected using dredges, epibenthic sleds, or box cores. Box cores can be deployed from a submersible, and samples can be kept separated in multiple containers (e.g. *Johnson Sea Link I, II*). Core samples can be taken directly by plunging a coring device into the sediment, an easy task for divers but difficult for manned submersibles. Unmanned vehicles, as far as I know, do not yet have this ability. The fauna that grows on rock surfaces, often termed the 'fouling community', and can be sampled easily by a diver. It is much harder to sample this community quantitatively from a submersible. An area of rock surface must be scraped, and a suction device used to collect everything from a known area of rock surface. Vacuum devices have been used by divers for the past two decades (Hiscock and Hoare 1973, Sebens 1985a), and several of the current manned research submersibles have this capability (e.g. *Alvin, Johnson Sea Link I, II*), although the unmanned vehicles do not.

Experiments in situ: Benthic experimentation includes tasks such as affixing markers to the substratum so that sites can be relocated and photographed. This is relatively simple for divers to accomplish but cannot be done by submersibles with any efficiency. Placing instrument or experiment packages is also easiest using divers. Manipulation studies, where the researcher changes something *in situ* (scraping a surface, removing sessile fauna, placing predator exclusion cages) is again very difficult to accomplish with anything but divers, who have much greater dexterity by comparison to submersibles.

One of the techniques used commonly to study benthic invertebrate larval recruitment is the placement of settling panels on a rack to hold them off the bottom. Large heavy racks require a large submersible to put them into position. The data from such experiments shows which species are able to recruit into the system rapidly or over longer time periods. Rock or carbonate blocks are another type of experimental settling plate which mimics the natural bottom surface. These are heavy but can be lowered from a surface vessel and the submersible used to orient them on the bottom using the submersible's locating capabilities as the main tool. Certain experiments can be lowered to the bottom and then manipulated into place with an unmanned vehicle as well; it would not necessarily have to carry them. Large cages on the bottom are used for predator inclusion and exclusion experiments, to see the effect of specific predators on their prey populations when they are maintained together or apart for a period of months to a year. Such caging experiments are relatively easy to emplace by divers or by manned submersibles; larger unmanned vehicles could accomplish this as well.

Physiological experiments, such as respiration chambers, have been emplaced by divers and by submersibles. Researchers must transfer animals in and out of the chambers and check such chambers for proper operation. Physiological experiments also have been conducted on the bottom in deep sea habitats, collecting animals *in situ* and placing them in specially constructed chambers. Such studies probably could not be done efficiently using an unmanned vehicle until their collection apparatus improves.

In addition to, or combined with, benthic experiment arrays, researchers often desire to collect long-term data on the physical environment concurrently. Instruments such as current meters, thermographs, salinometers, irradiance meters and others can be attached to experimental arrays or deployed separately nearby. Again, manned submersibles are ideal for this task, which involves complex manipulation as well as accurate placement. A major limitation for remotely operated vehicles is the large size and weight of both experiment arrays and oceanographic instrument packages.

Rapid modification of vehicle for task-specific apparatus: The installation of a large rack on the front of a submersible has turned out to provide an excellent platform for benthic sampling because it can support many different types of sampling devices. The *Johnson Sea Link*, for example, has a Lazy-Susan device that rotates sampling containers so that multiple suction or grab samples can be taken on one dive. Benthic research depends on quantitative information with large numbers of replicates. Replicate samples should come from similar habitats and must be kept separate, and prevented from flushing out on the way up. It is actually quite a feat to design something that will do this and it usually ends up very large and awkward. A multiple sample collection device has yet to be developed for the unmanned vehicles.

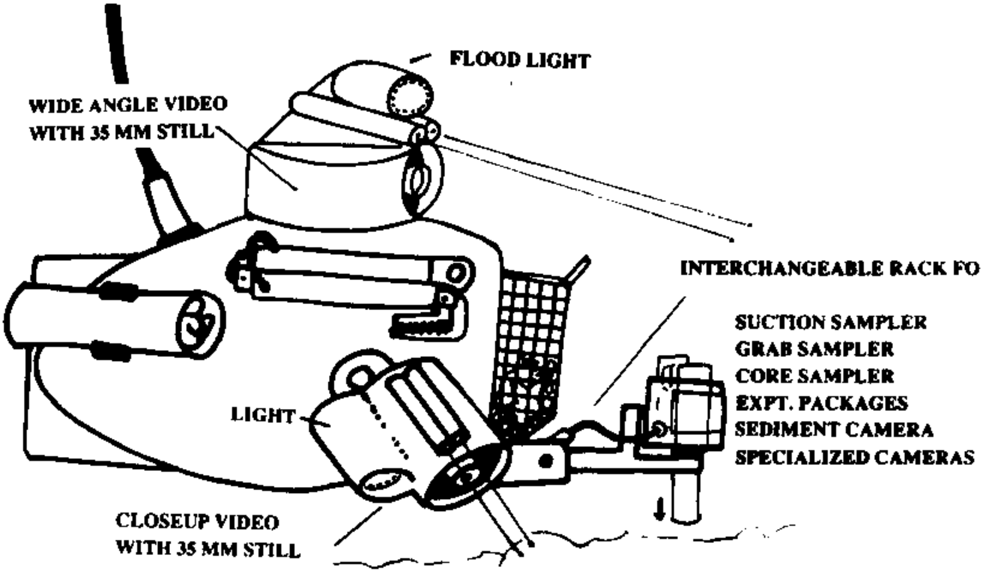

Figure 1. An unmanned submersible (ROV) designed for benthic ecological research, not based on any existing vehicle. The submersible has two video systems, one for large scale surveys and vehicle navigation, the other for quantitative benthic photography. Both video systems have paired laser scaling devices and coupled 35mm still cameras using the same lens as the video. The top-mounted rotating video turret allows the submersible pilot to scan the surroundings without moving the submersible. The interchangeable front rack allows rapid modification for specific sampling tasks.

A CASE STUDY: SUBMERSIBLE BASED RESEARCH IN THE GULF OF MAINE

The National Undersea Research Program (N.O.A.A.), operating out of the University of Connecticut at Avery Point, began funding and coordinating a submersible-based research effort in the Gulf of Maine in 1984. This multidisciplinary research has continued during the summers of 1985 and 1986, investigating habitats from 50 to 250 m depth. The hard substratum research group (K. Sebens, R. Steneck, R. Vadas, J. Witman) has focused on rock ledges in the central Gulf of Maine, primarily Cashes Ledge, a very important bottom fish feeding ground. This research is an attempt to examine hard bottom areas in the Gulf of Maine because these areas have been very poorly sampled in the past. They cannot be sampled effectively by towed photographic sleds. Such habitats are best studied by SCUBA diving if shallow, by manned or unmanned submersibles if deep.

The Gulf of Maine project has, so far, made use of four manned submersibles, the Mermaid II, the Johnson Sea Link I and II, and the Nekton Delta. These vehicles vary considerably in their applicability to benthic sampling. The unmanned vehicles that have been used are the RECON IV and the Deep Sea MiniRover. The RECON IV is an intermediate capability machine. It was outfitted with several different photographic systems (video, 35mm still camera), and a compact manipulator arm. Scale for photographs could be provided only by taping a small ruler to the manipulator jaw. The RECON IV was used to survey sites on Cashes Ledge before choosing a site for the manned submersible dives. Still photographs and videos were used to determine depth distributions of sponges (Witman and Sebens 1987), various other invertebrates (Witman and Sebens 1985), and macroalgae (Steneck 1986ab, Vadas and Steneck 1987) and to search for dense scallop populations (Shick et al 1987, Shumway et al. 1987), lobster, and bottom fish populations. The only non-photographic task that was attempted with the RECON IV was to collect voucher specimens of sponges for species identification. This was an extremely difficult procedure resulting in a very few specimens. The MiniRover has only photographic capability, and was used by Townsend and Stevenson to extend their surveys of herring spawning areas and larval dynamics (Townsend et al. 1986).

The manned submersibles, on the other hand, allowed us to take good quality video transects and still photographs and to collect numerous specimens by grab and suction techniques. The *Johnson Sea Links* (I and II) were equipped with rotating sample container arrays to keep twelve or more samples separate and undisturbed after collection. We were also successful in deploying large predator inclusion and exclusion cages, granite blocks as settling plates, racks of plexiglass settling plates, time-lapse cameras and bottom-mounted recording current meters (Interocean S-4). Scientist observations from the submersible provided the first information on predatory activity of certain benthic invertebrates and on the abundance of bottom fish utilizing this habitat. Permanent quadrats for still photography were established and marked with lead weights, although we found that the bow-mounted still camera was not satisfactory for repeated photography of marked quadrats; an arm-mounted camera will be used in the future.

Of all the tasks carried out, none except the photography could have been done efficiently, if at all, using the available unmanned submersibles. An ideal combination, at present, is to use an unmanned vehicle for surveys, possibly while a manned submersible is being recharged or refitted, followed by manipulative tasks using the manned submersible. The remotely operated vehicle that we used in 1984 (*RECON IV*) saved us from deploying the manned submersible in habitats that were not appropriate for our research needs. In the future, remotely operated vehicles may approach the capabilities of the manned submersibles but, until then, they are cost effective and time saving alternatives for the only photographic and survey components of benthic research programs.

The future: what can we expect? Part of what must be done, and this development is in progress, is to take the technology that has been developed for the manned submersibles and adapt it to the ROVs. There may also be things that the ROV can do that, in fact, the manned submersible has a harder time doing. One of these is simply spending a lot of time on surveys (Stewart, this volume) where many hours on the bottom are necessary and where there is not much manipulation necessary. The expense of such a project using a manned submersible can be prohibitive but is quite tolerable using an ROV. It is also impossible for divers below their safe working depths. There is thus a unique niche for ROV's in benthic survey work.

Figure 1 is a diagram of what an unmanned submersible for benthic sampling might look like if it were designed just for that purpose. Two video systems could be installed, one for observing on the large scale (wide angle), for finding specific sites, and for surveys of fish and large invertebrates, and one that can be used for close-up work. A scaling device such as the paired laser system would be available for either video. A rack on the front would accept interchangeable units comprising suction samplers, grab samplers, core samplers, experimental packages, sediment cameras, and other types of specialized cameras. Rapid manipulator arms, now being developed, are necessary as well.

This design configuration represents a transfer technology from some of the existing manned submersibles, and some of the diver sampling technology, to improve unmanned submersible sampling capabilities. This is a short term solution, however. The long range possibilities are even more exciting as unmanned submersibles take on new and presently unimagined autonomy, with new instrumentation for environmental sensing, data storage, data transmission and decision making. Unmanned submersibles will soon be able to work far more efficiently, and for longer periods, than divers can today. In the deep sea, they will work alone or in tandem with manned submersibles or habitats. They may even find a useful role in shallow water research as companion vehicles for divers, carrying instruments and experiment packages as well as conducting surveys on their own while divers are occupied with manipulative tasks. Inexpensive vehicles, that can be deployed from small boats by one or two persons, are bound to proliferate and become increasingly sophisticated. Marine ecologists can only benefit from these exciting developments in submersible engineering.

References:

Anonymous. 1985. Undersea Vehicle Directory. Busby Associates Inc., Penn.

Bohnsack, J. A. 1979. Photographic quantitative sampling of hard-bottom benthic communities. Bull. mar. Sci. 29:242-52.

Christie, H., R. A. Evans and O. K. Sandnes. 1985. Methods for *in situ* subtidal hard bottom studies. pp 37-47 in: J. D. George, G. I. Lythgoe and J. N. Lythgoe, eds. Underwater Photography and Television for Scientists. Clarendon Press, Oxford.

Fedra, K. and R. Machan. 1979. A self-contained underwater time-lapse camera for *in situ* long-term observations. Mar. Biol., Berlin 55:239-246.

George, J. D. 1980. Photography as a marine biological research tool. pp. 45-115. In: The Shore Environment, Vol. 1: Methods. eds. J. H. Price, D. E. G. Irvine, and W. F. Farnham. Academic Press, London.

Green, N. W. 1980. Underwater stereophotography applied in ecological monitoring. Report 1. Methods and preliminary evaluation. Norwegian Institute for Water Research. Report OF-80613, 99 pp.

Hessler, R. R. and W. M. Smithey, Jr. 1984. The distribution and community structure of megafauna at the Galapagos Rift hydrothermal vents. pp. 735-770. In: Hydrothermal processes at sea floor spreading centers, NATO Conference Series, Series IV: Marine Science. eds. P. A. Rona, K. Bostrom, L. Laubier and K. L. Smith, Jr. Plenum Press, New York.

Hiscock, K. and R. Hoare. 1973. A portable suction sampler for rock epibiota. Helg. wiss. Meeresunters. 25:35-38.

Holme, N. A. 1985. Use of photographic and television cameras on the continental shelf. pp. 88-99. In: J. D. George, G. I. Lythgoe and J. N. Lythgoe. Underwater Photograpy and Television for Scientists. Clarendon Press, Oxford.

Holme, N. A. and R. L. Barrett. 1977. A sledge with television and photographic cameras for quantitative investigation of the epifauna on the continental shelf. J. mar. biol. Ass. U.K. 57:391-403.

Lundälv, T. 1971. Quantitative studies on rocky-bottom biocoenoses by underwater photogrammetry: A methodological study. Thalassia jugosl. 7:201-208.

Machan, R. and K. Fedra. 1975. A new towed underwater camera system for wide-ranging benthic surveys. Mar. Biol., Berlin 33:75-84.

Rhoads, D. C. and D. K. Young. 1970. The influence of deposit-feeding organisms on sediment stability and community trophic structure. J. Mar. Res. 28:150-178.

Schick, D. F., S. E. Shumway and M. Hunter. 1986a. Allometric relationships and growth in *Placopecten magellanicus*: The effects of season and depth. In press.

Schick, D. F., S. E. Shumway and M. A. Hunter. 1986b. A comparison of growth rate between shallow water and deep water populations of scallops (*Placopecten magellanicus*) in the Gulf of Maine. In press.

Sebens, K. P. 1984. Water flow determines colony size in a temperate octocoral. Proc. Natl. Acad. Sci. 81:5473-5477.

Sebens, K. P. 1985a. Community ecology of vertical rock walls in the Gulf of Maine sublittoral zone: small-scale processes and alternative community states. Chapter 23. pp.346-371. In: P. G. Moore and R. Seed, eds. The Ecology of Rocky Coasts. Hodder and Stoughton Educational, Kent.

Sebens, K. P. 1985b. Spatial relationships among encrusting marine organisms in the New England subtidal. Ecol. Monogr. 56:73-96.

Shumway, S. E., D. F. Schick and S. Naidu. 1987. Synopsis of biological data on the scallop, *Placopecten magellanicus*. NOAA Technical Report/NMFS. Manuscript.

Smithey, W. M. Jr. and R. R. Hessler. 1985. Megafaunal distribution at deep sea hydrothermal vents: an integrated photographic approach. pp. 165-176. In: J. D. George, G. I. Lythgoe and J. N. Lythgoe, eds. Underwater Photography and Television for Scientists. Clarendon Press, Oxford.

Southward, A. O. and D. Nicholson. 1985. Photography of the deep-sea bottom fauna with remotely operated cameras. pp. 122-152. In: J. D. George, G. I. Lythgoe and J. N. Lythgoe, eds. Underwater Photography and Television for Scientists. Clarendon Press, Oxford.

Steneck, R. S. 1986a. The ecology of coralline algal crusts: convergent patterns, and adaptive strategies. Annual Reviews of Ecology and Systematics. 17:235-255.

Steneck, R. S. 1986b. The tie that binds technology and research. Explorations: a journal of research at the University of Maine. 2:36-38.

Stewart, L. 1987. Underwater Research needs of fisheries. This volume.

Torlegard, A. I. and T. L. Lundälv. 1974. Underwater analytical system. Photogramm. Engng. 40:287-293.

Townsend, D. W., J. J. Graham and D. K. Stevenson. 1986. Dynamics of larval herring (*Clupea harengus* L.) production in tidally mixed waters of the eastern coastal Gulf of Maine. In: Tidal mixing and plankton dynamics. M. Bowman, C. Yentsch and W. T. Peterson eds. Lecture Notes on Coastal and Estuarine Studies. Springer-Verlag, Vol. 17:253-277.

Sanders, H. L., R. R. Hessler and G. R. Hampson. 1965. An introduction to the study of the deep sea benthic faunal assemblages along the Gay Head-Bermuda transect. Deep Sea Res. 12:845-867.

Uzmann, J. R., R. A. Cooper, R. B. Theroux and R. L. Wigley. 1977. Synoptic comparison of three sampling techniques for estimating abundance and distribution of selected megafauna: submersibles *vs.* camera sled *vs* otter trawl. Mar. fish. Rev. 39:11-19.

Vadas, R. L. and R. Steneck. 1987. Zonation of deep water benthic algae in the Gulf of Maine. Marine Ecology Progress Series. In Press.

Vadus, J. R. and R. F. Busby. 1979. Remotely operated vehicles - an overview. Proc. 15th A. Conf. Mar. Technol. Soc., New Orleans, 1979. 120-129. Marine Technology Society, Washington.

Witman, J. D. 1985. Refuges, biological disturbance and rocky subtidal community structure in New England. Ecological Monographs 55:421-445.

Witman, J. D. 1987. Subtidal coexistence: storms, grazing, mutualism, and the zonation of kelps and mussels. Ecological Monographs. In press.

Witman, J. D. and K. P. Sebens. 1985. Benthic faunal zonation at a subtidal rock ledge in the central Gulf of Maine. Amer. Zool. 25:10.

Witman, J. D. and K. P. Sebens. 1987. Distribution and ecology of sponges at a subtidal rock ledge in the central Gulf of Maine, U.S.A. In: K. Reutzler, ed. Proc. Symp. Sponge Biol., October 1985 Woods Hole, MA. Smithsonian Inst. Press.

Remote Manipulation Systems for Research ROVs

Andrew M. Clark
Head, Mechanical Engineering Department
Harbor Branch Oceanographic Institution, Inc.

Tyler Schilling
Director
Schilling Development, Inc.

INTRODUCTION

Harbor Branch Oceanographic Institution (HBOI) has for the past 15 years been developing and operating two of the most active and productive (Alldredge and Youngbluth, 1985; Littler et al., 1986) undersea research vehicles, the 805 m operating depth JOHNSON-SEA-LINKs (JSL) I and II. These four-man submersibles were designed with lock-out capabilities enabling two researchers to exit and re-enter the craft at depth. It was originally envisioned that manipulative tasks, and those functions requiring finesse, would be carried out by these lock-out divers. Engineers at HBOI have, however, developed an inventory of tools systems (Tietze and Clark, 1986) which enable the sub's occupants to perform a wide variety of research from within the safety of its hull, and well beyond the depth limits of safe lock-out diving.

Having eliminated the need for the researcher to venture outside the submersible, the next logical evolutionary step, as evidenced by industrial counterparts, is to remove the researcher from the submersible altogether. While the oil industry's replacement of manned submersibles with ROVs is all but complete, the subjective nature of such qualitative research as behavioral studies, in which HBOI is engaged, dictates the continued use of manned submersibles.

There are, however, a number of tasks for which JSL I and II have been employed which could be performed equally well or better with a properly designed and equipped ROV.

Among the more obvious of these tasks are the delivery and retrieval of unattended *in situ* measurement devices. Beyond these brutish chores, however, lie a number of much more complex research tasks which require the presence and attention of the investigator to perform. The identification and collection of benthic and midwater organisms has been reduced to routine from the manned JSL I and II. The systems described herein afford the investigator sufficient "telepresence" to enable him to perform such work with an ROV.

Demand by the research community for JOHNSON-SEA-LINK I and II has exceeded their availability. To meet this demand, and extend the Institution's 805 m operating depth, HBOI purchased a 1,000 m HYSUB-40 ROV (Figure 1). The vehicle has been fitted with a bolt-on undercarriage (SCIENCE SLED), 30.5 cm high, 104 cm wide and 145 cm deep. Mission-suited tool systems are designed to slide easily into the SCIENCE SLED. This not only facilitates ease of reconfiguring the vehicle for varying tasks within a cruise, but it is further hoped that this concept will be adopted by other research ROV operators, thereby providing the capability for interchanging tool packages among operators, thus facilitating collaborative efforts.

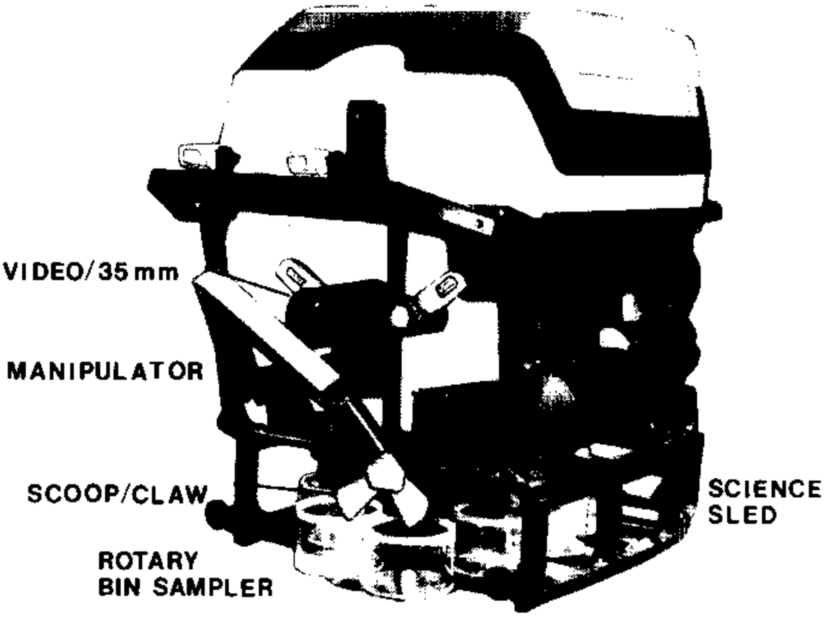

Figure 1. HBOI's HYSUB-40 ROV

SAMPLING SYSTEMS

Manipulator

The portion of the work package which serves to unite the system is the manipulator arm, (Figure 2) designed and manufactured by Schilling Development, Inc. Careful allocation of task responsibility is of key importance in any properly functioning system. The purpose of the manipulator in this scientific work package is to provide the operator with an interactive means of carrying out collection tasks while freeing him of the tedious operations, such as sample storage, which do not necessarily require his attention.

These requirements established certain design criteria for the manipulator. The machine must be subject to a tight control loop for the performance of interactive as well as pre-programmed operations. In addition, in the interactive mode of operation, the slave arm must perform at a velocity high enough to allow the operator's own human control system to play an active roll; the operator should never have to wait for the slave arm. The ability of the arm to maximize performance is determined by its mechanical, as well as its electrical design. While the dynamic performance of the system can, to a great extent, be shaped in software, proper hydraulic design was not

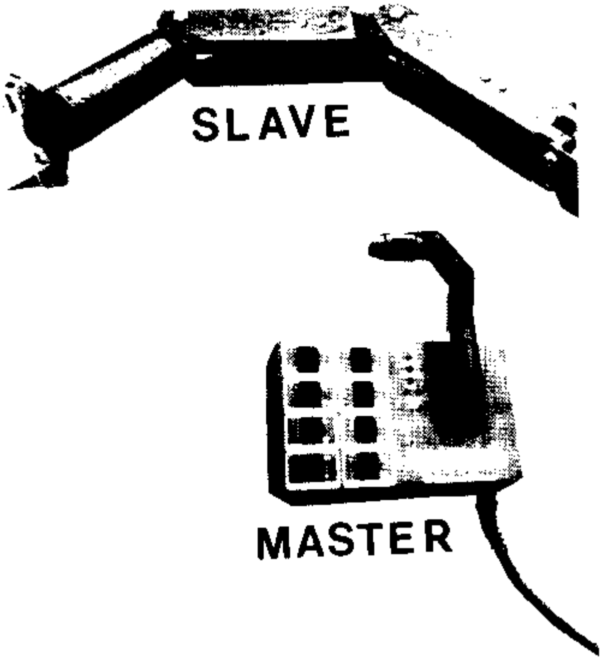

FIGURE 2. Schilling Manipulator and Controller

overlooked. The manipulator system yields this high performance through careful attention to mechanical design and the integration of a powerful microprocessor control system, in addition to a hand held terminal assembly to provide maximum flexibility. An important aspect of the manipulator portion of this system, is the design of its master control assembly. Its miniature design affords the operator with a comfortable means of controlling the five-degree of freedom arm. Unlike conventional master control arms, this design does not require the operator to hold his or her arm outstretched at any position in the operating envelope. From the front panel of the master control assembly, the operator can select various modes of arm operation. The overall control dynamics can be altered by the operator to facilitate the arm's operation in air or submerged. Additionally, the wrist rotate function can be operated in a slaved mode or a velocity mode to provide for tasks requiring continuous wrist rotation. Likewise, the jaw function has two modes of operation; it can be used in the toggle mode, which requires two actuations to open and then close the jaw, or in the open mode, where the jaw remains open unless the button at the tip of the master arm is held depressed. The slave arm can also be electronically frozen in any position from the front panel of the master control assembly.

A variety of detachable appendages which affix to the arm have been designed to perform various types of collecting, grasping or manipulation. Figure 1 depicts a multi-purpose scoop/claw designed especially for grasping and dislodging sponges and similar organisms. The multiplexed control system (described below) for the manipulator also provides proportional control for the suite of collection and data acquisition devices.

Control System

Beyond the mechanical design of the arm lies its control system (Figures 3 and 4). The control system for this manipulator was designed with several criteria in mind: power, reliability, flexibility and expandability.

For the package to possess power, it must have sufficient processor band width. Hence, the system was based around an Intel 8088 microprocessor. While this is not a brute-force chip, it provides more than adequate power.

The reliability of the package is assured by maximizing the use of digital components throughout the system, along with the fact that all control system parameters are defined in software. This provides for a system that requires no hardware adjustments.

System flexibility is obtained through the software-defined control parameters. This allows the user to actually alter the control algorithm through the hand held terminal.

The issue of expandability was addressed in the design by providing additional analog interface circuitry beyond that which the arm itself requires. The manipulator control consists of two

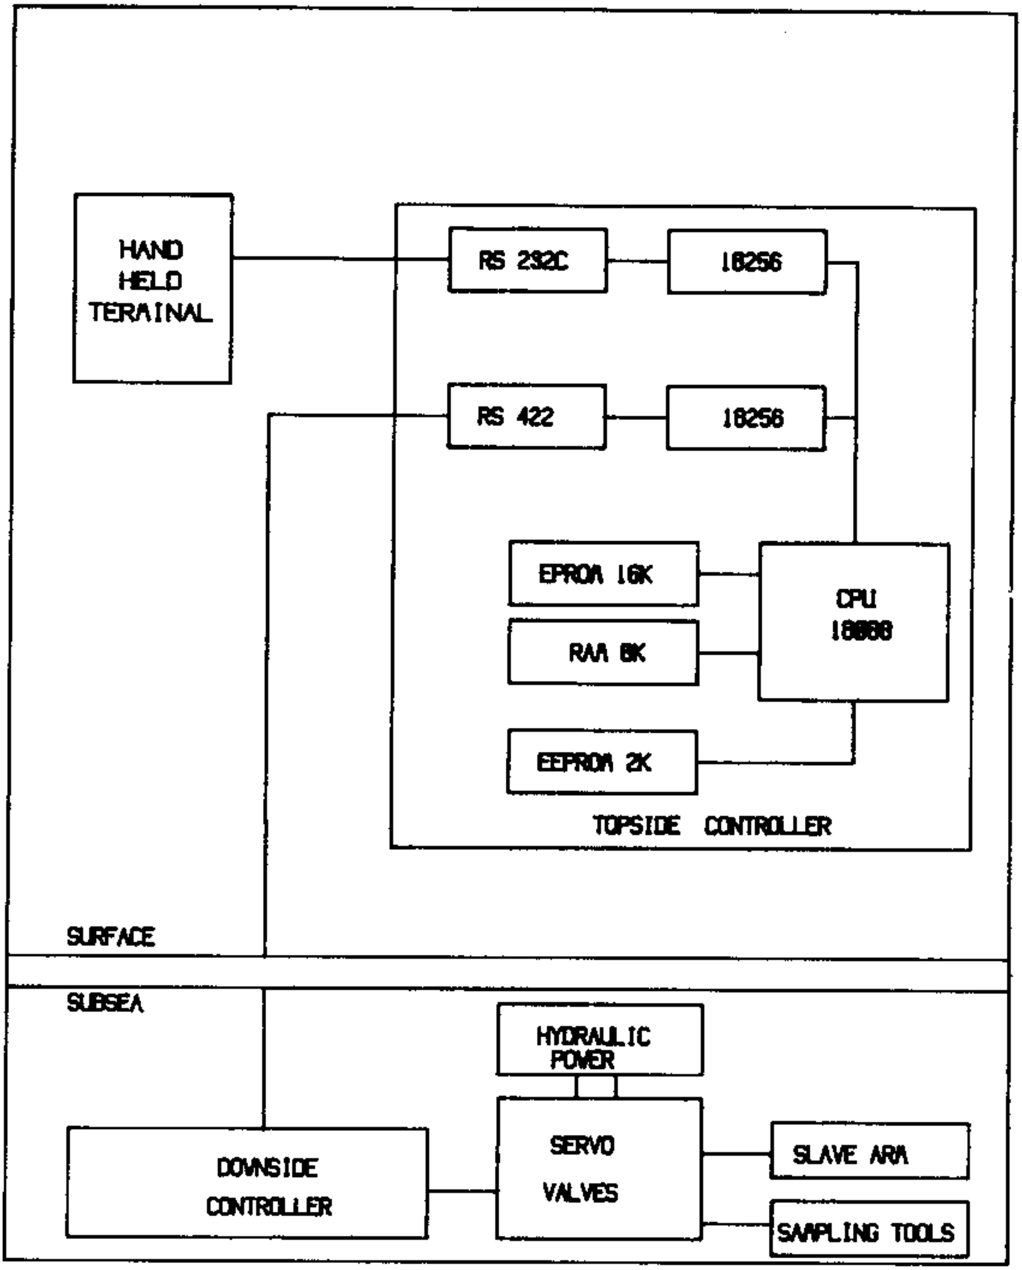

Figure 3. Control System Schematic

identical computer packages, one top-side and the other housed within the subsea valve package (Figure 5). These assemblies communicate over a twisted wire pair by an RS 422 standard. They utilize a data packet transmission scheme to provide maximum fault tolerance.

The function of the down-side control package is the closure of the control loop, while the top-side assembly contains the "personality" of the system. The top-side software generates the data used to update the tables in the down-side control loop. This data is generated by either reading joint angle values from the master arm, or

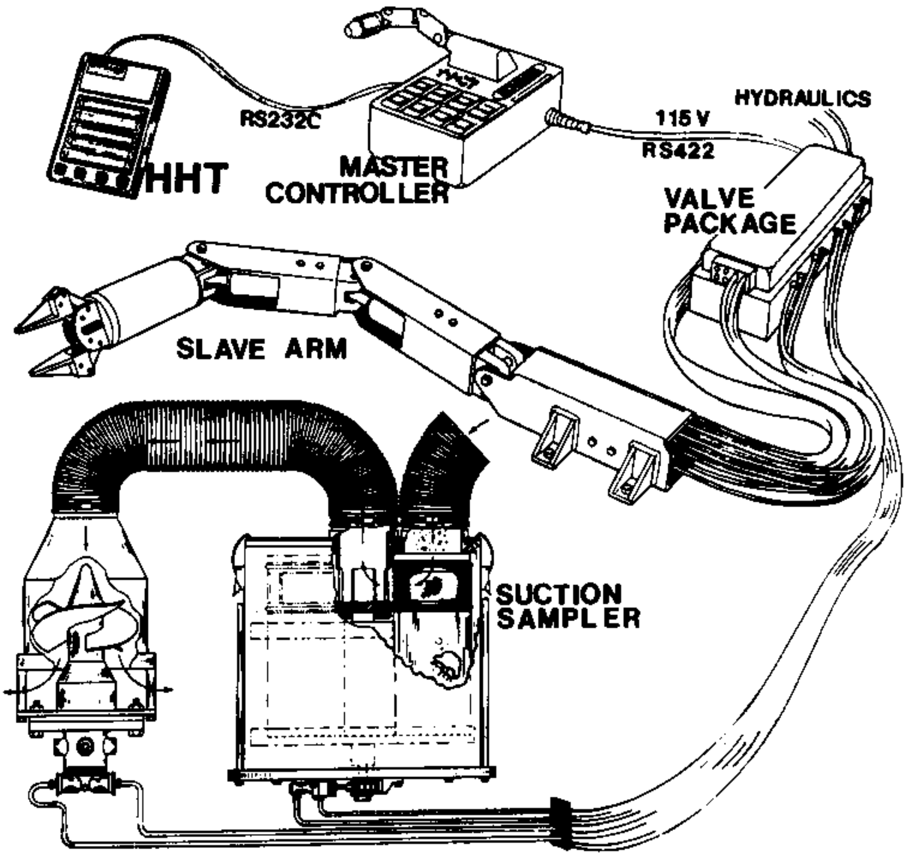

Figure 4. Remote Manipulation System

drawing this information from memory in the case of pre-programmed operation. Also contained in the top-side unit are the values which represent the limits of the operational envelope. These limits can be defined to prevent the manipulator from making unintended contact with other vehicle systems.

As a result of close attention to design of the mechanical portion of the machine, the actual control scheme implemented is very classical, yet still provides the performance objectives established at the onset. Further performance gains are expected to be yielded through the implementation of more creative techniques. These gains are envisioned to provide optimum performance in the pre-programmed modes of operation.

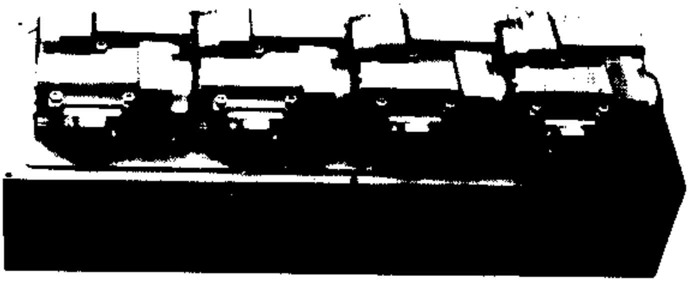

HYDRAULICS (TOP VIEW)

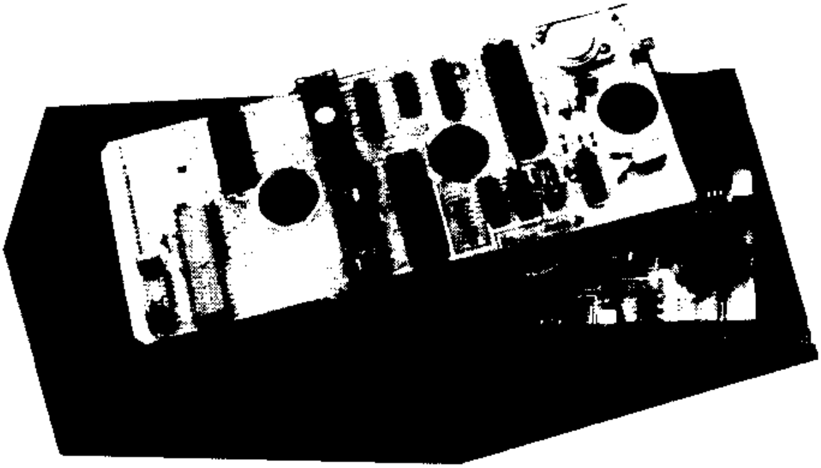

ELECTRONICS (BOTTOM INSET)

Figure 5. Subsea Valve Package

Hand Held Terminal (HHT)

The practicality of pre-programmed operation is provided through the use of the Hand Held Terminal (HHT). This menu-driven terminal unit allows the operator to input a program sequence to the control computer in two ways. The operator can select a point-to-point mode for simple program paths, or a streaming mode with pre-selected record increments for more involved paths. In both modes, the operator moves the arm in a normal master/slave manner. He can then store these paths, and execute them on command. The HHT is composed of five LCD screens, each with fifteen touch sensitive areas (Figure 6). To operate the device, the operator need only touch the desired command prompt. The unit also affords the capability to read auxiliary sensory data: hydraulic pressure, temperature, water intrusion, etc., as well as the capability to control auxiliary functions. These functions may include the operation of the ROTATING BIN SAMPLER, suction pumps, still cameras or other data acquisition systems. An additional benefit of the HHT is its ability to perform diagnostic functions on the manipulator system. The ultimate power of the control system and HHT lies in their flexibility. The low probability of the designer predicting fully the needs of the operator in the field has been established through experience. Therefore, the ability of the operator to change the rules "on the fly" is very important.

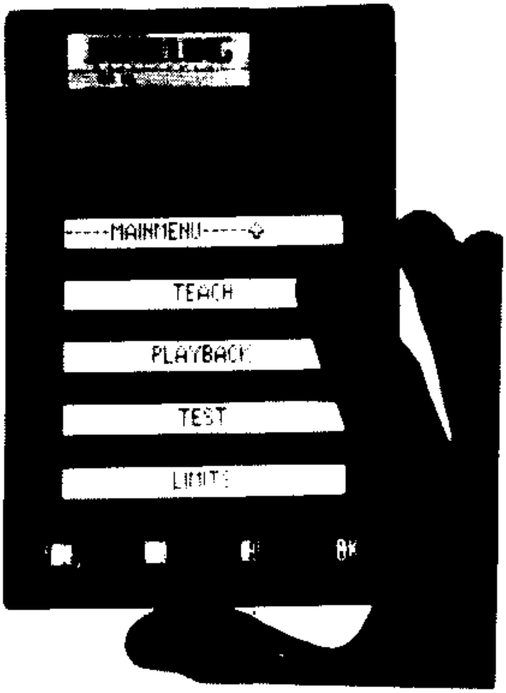

Figure 6. Hand Held Terminal (HHT)

Rotating Bin Sampler

The ROTATING BIN SAMPLER slides into the tracks of the SCIENCE SLED. Basically, the device indexes and stores discrete samples in sets of up to 18 (17.8 cm diameter x 17.8 cm) bins or 12 (25.4 cm diameter x 17.8 cm) bins. The bins are acrylic cylinders which are conveyed to the sampling location by way of a titanium chain. When collecting large, solid specimens, as with the manipulator's scoop appendage, a numbered bin is indexed beneath a funnel into which the manipulator drops the specimen. The operator can effect this through either the supervisory or pre-programmed mode. The system can also be used to collect and index macro- or microscopic, and liquid samples. A suction plenum is mounted at the top of the tray, adjacent to the funnel. Samples are drawn in by means of a hydraulically driven "slurp" pump through a flexible hose fixed to the manipulator. The specimen is deposited in the bin positioned beneath the suction plenum. The titanium chain is driven and indexed by a reversible, low speed-high torque hydraulic motor.

The suction force is infinitely variable through the multiplexed control system, and flow is monitored topside by the output of an in-line flowmeter. Thus, the device can be adjusted to low flow, for capturing extremely delicate organisms, or to high flow for dredging sediment.

Photographic/Video

Much of the data collected by researchers using HBOI submersibles is visual documentation. As such, high quality video and still photography are paramount in making this ROV a viable research tool. Sharp focus of 35 mm still photographs is insured by a system in which both a color video imaging tube and a 35 mm film cartridge share the same lens (Figure 7). Thus, images brought into focus on the topside monitor by the operator can be recorded on 35 mm film by a single trigger which activates both the shutter and strobe. The microprocessor-controlled system also provides the researcher alphanumeric write-on capability for specimen documentation.

In order to maintain the highest possible resolution, video signals are transmitted up the lift umbilical via 50/120 µm graded index silicon optical fibers. Analog fiber optic transmitters and receivers are nested within both the subsea and the surface slip rings, thereby eliminating the need for an optical swivel joint.

The vehicle is equipped with an auxiliary transformer which provides 1500 W at 24V, 28V, 32V DC and 110V AC to facilitate various other 35 mm or large format camera systems and their strobes. Compact dual laser systems have been developed to serve as range-finding and absolute measurement scales.

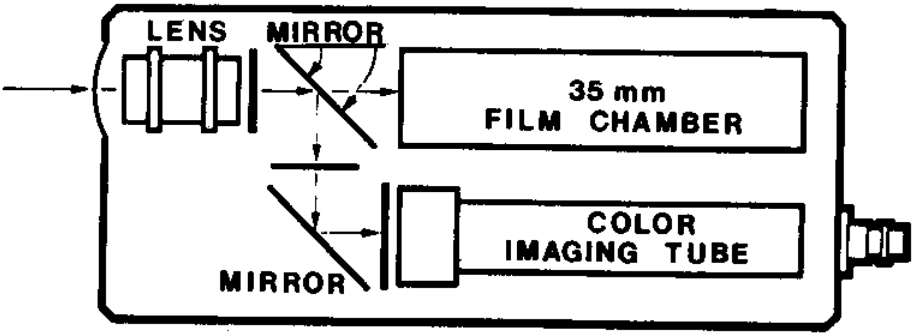

Figure 7. Video/35 mm System (Courtesy of Osprey Electronics)

FUTURE ENHANCEMENTS

Further optimization of the manipulator portion of this work system is anticipated in the area of interchangable end effectors. Man often takes for granted the fact that the human hand functions as dozens of tools while the jaw portion of a manipulator has very limited capabilities. Therefore, the ability to change out remotely and automatically specialized end tooling is of utmost importance. A system to provide for the interchange of hydraulic tools which may be either open or closed loop type is in the initial stages of design.

A number of different schemes are being evaluated for providing pseudo 3D viewing by means of color stereo video systems. Such a system, coupled with the easily mastered and highly dextrous Schilling manipulator, will provide the researcher with further telepresence. Thus, the co-pilot's seat at the operator console is readily occupied by investigators with little or no ROV experience, yet with the ability to identify, collect, index and record a vast array of specimens at depths to 1,000 m.

Data acquisition systems, standard to the JSL manned submersibles, which enable monitoring and recording such parameters as conductivity, temperature, and transmisivity are being adapted for use on this vehicle. Perhaps the single most significant enhancement anticipated is an upgrade of the ROV's operating depth. Through modifications to the vehicle's telemetry system and replacement of the 38 mm diameter lift umbilical with an umbilical more reliant on optical fibers and subsequently, of significantly decreased diameter, it is intended to ultimately increase the vehicle's operating depth to 2,000 m. The vehicle's manufacturer, International Submarine Engineering, Ltd. (ISE) has built two other vehicles to-date which meet or exceed this design depth, and are currently building a 5,000 m system. As such, it is not anticipated that this upgrade will present insurmountable technological problems.

REFERENCES

Alldredge, A.L. and M.J. Youngbluth. 1985. The significance of macroscopic aggregates (marine snow) as sites for heterotrophic bacterial production in the mesopelagic zone of the subtropical Atlantic. Deep Sea Research 32: 1445-1456.

Littler, M.M., D.S. Littler, S.M. Blair and J.N. Norris. 1986. Deepwater plant communities form an uncharted seamount off San Salvador Island, Bahamas: distribution, abundance and primary productivity. Deep Sea Research 33: 881-892.

Tietze, R.C. and A.M. Clark. 1986. Remotely Operated Tools for Undersea Vehicles. Current Practices and New Technology in Ocean Engineering. ASME OED-Vol. 11: 219-223.

The Sixth
Robert Bruce Wallace
Lecture

The Genesis and Metamorphosis of Underwater Work Vehicles

James R. McFarlane
President
International Submarine Engineering Ltd.

HISTORY

Man's interest in establishing a capability to work underwater goes back to the beginnings of recorded history. However, it was not until the 1600's that concepts which would shape this capability began to evolve. For example, in the 1600's diving bells were conceived by Halley and Kessler. The August Siebe suite followed 65 years later. In the late 1800's and early 1900's, diving physiology associated with increased diver activity was investigated by Paul Bert and John Haldane; rigid suites by Camagnole brothers, Galeazzi were developed in Italy and were followed by those of Neufeldt and Kuhnke in Germany. Subsequently, in 1934, Bebee gained fame with his bathysphere. In the mid forties Cousteau and Gagnon developed the breathing regulator, which enabled man for the first time complete freedom to maneuver in the water.

At the outset of the 60's, developments began to accelerate to extend depth, duration and capability. Investigations which had previously proceeded on a sporadic unstructured evolutionary basis, began to evolve rapidly in a manner which is characteristic of a revolution. This revolution was fueled by the needs of the offshore petroleum industry and a keen USN interest.

On January 23, 1960, Piccard and Walsh descended 35,800 feet to the bottom of the Challenger Deep. In 1962 divers Hannes Keller and Peter Small descended to 1,000 feet. Keller survived, Small did not. In 1964 Captain George Bond, USN, and Captain Cousteau were conducting mixed gas diving experiments in the USA and France respectively. At the same time the development of manned submersibles commenced, habitats were constructed and lived in and the remotely controlled vehicle CURV was developed by the USN. Man left the 60's and entered the 70's with a real capability to work in the ocean. Some of the systems were not very reliable, but they had all the attributes of work systems. Shortcomings were readily amenable to refinement.

As is common in a revolutionary situation, the period since the 60's has been a period when entrepreneurs and innovators alike have had spectacular successes, as well as failures. In most instances, the failures can be traced to poorly identified objectives, lack of attention to commercial return and the production of laboratory curiosities. The manned submersible business had many examples of both success and failure. In the commercial setting, success and survival over the longer term depends on being aware of the true requirements of the market place, and being responsive to these requirements. One must also pay careful attention to government agency planning and regulations which can dramatically affect the market size and sometimes, unfortunately, impede technical development.

As the industry has matured, systems have been developed to fill the rows and columns of the work systems matrix shown in Figure 1. The spaces described as manned-tethered, manned-untethered, manned remote-tethered have numerous data points. The remaining spaces manned remote-untethered have only recently began to receive attention. One can expect substantial effort to be expended on these last three spaces in the next half decade.

Another way to consider the work space is to represent the capabilities of work systems by cylinders; See Figure 2. This figure allows one to visualize the space characterizing unique and competing capabilities among the work systems. The intersections of the cylinders indicate areas where systems compete directly. Areas which do not intersect represent activities where systems have substantial uniqueness. These intersecting and non-intersecting areas are affected in a complex fashion by depth and other environmental conditions such as the current profile and visibility. Thus new intersections are created for each external variable considered. When one thinks of the work space in this manner, it becomes obvious that no individual equipment is a panacea for the spectrum of subsea tasks which exist.

VEHICLE WORK SYSTEM EVOLUTION

As the market place is addressed, system evolution proceeds in two ways, namely the evolution of vehicle types and the evolution of their associated subsystems. There are undoubtedly vehicles which are feasible which have yet to be built. These are not necessarily totally new, but can be variations on existing vehicles. Figure 3 attempts to provide a visualization re how future vehicle development might proceed. Figure 3 describes a space with three pure vehicle types on the x, y, z axes. The pure vehicles described by the three axes are: ROVs used for inspection and manipulation, autonomous subsea vehicles presently used only in survey and dynamically stabilized radio controlled vehicles used as instrument delivery vehicles. Any part of the vehicle space characterized by a vector where any two coordinates are greater than zero is a hybrid. Clearly the hybrid volume represents fertile ground for vehicle development. An example of a hybrid vehicle is the Heriot-Watt University ANGUS. It is easy to envisage another version of Angus which would allow it to disconnect from its umbilical. Angus has batteries which are trickle charged from the surface. Combinations of acoustic remote controlled vehicles ARCS and radio remote vehicles DOLPHINS are also obvious hybrids. One could also imagine a hybrid vehicle mutation being produced by mission requirements. The resultant family creates yet another vehicle space.

During the revolutionary period of the 60's, useable ROV technology was developed. Interestingly this was not immediately picked up by the subsea service industry. Industry of the 60's was caught up in the concept of man in the ocean, much as the space program today is focused on man in space. However, in the late 70's, work systems in the oceans began to move rapidly toward robotics, primarily due to commercial considerations. This trend is not as strongly in evidence in space, where man in space appears to have the highest priority.

ROVS

Once the notion of ROV use in the work place became obvious and profitable, the phenomenon referred to by noted industry author Frank Busby as the 'thundering herd syndrome,' took over. At one point in the 70's, more than thirty companies were involved in ROV development. Now as some areas of the market begin to mature, the classic theory of concentration of power obtains.

Typical of early developments in any revolution, each company touted their particular design as being universal in capability. This notion, which the proponents may have indeed believed in, was, and is, specious. There is no such thing as a universal vehicle. Families of vehicles and hybrids of these families are required to meet the needs of a multidimensional work set. Thus the attributes of a successful vehicle must, in the end, be mission driven. This has been clearly established in more mature industries such as the aviation industry and say the construction industry where work vehicles could be characterised in a space similar to the undersea vehicle space shown in Figure 3.

The first work to be undertaken by ROVs was pipeline inspection. This was followed by the more sophisticated task of drilling support. In turn, this will be followed by greater use in platform inspection and new uses in production field construction and maintenance sometimes referred to now as intervention.

Today ROVs span in size from very small vehicles weighing under 50kg to the large "TRAPR" weighing 8500kg. Current diving depths for most commercial ROVs range from 100 meters to 2500 meters with most distributed in the range of 400 to 1000 meters.

SUBSYSTEM DEVELOPMENT AREAS

The tasks which have been undertaken by ROVs have become increasingly sophisticated, and this has led to evolution in vehicle subsystems of which some have evolved more rapidly than others. The more mature subsystems have developed more slowly. One would not see the need for a major expenditure of effort on developments in vehicle structure or ballast, as marginal changes in these areas would not produce significant improvements in cost or performance. However, marginal changes or advances in propulsion, electric power, control, navigation sensors, manipulation and tools could dramatically impact vehicle performance. Obviously some attention has to be paid to the cost benefit of such developments.

The propulsion subsystem is a good example of where a significant revolution would obtain if low cost, light weight, quiet, easily controlled electric motors could be developed. These may indeed evolve out of brushless technology. Such developments could displace the electro-hydraulic strategies. Alternately, small low cost hydraulic components could displace electric motors now used primarily in small vehicles.

Developments in battery technology or other forms of onboard vehicle energy conversion, such as a small stirling engine, would change the transition from surface to onboard vehicle power. On board vehicle power, with sufficient endurance, would be attractive because of reduced umbilical cost and reduced umbilical drag. This would again adjust the transition. Given that onboard vehicle power endurance is relatively more finite than surface power, once the decision was made to implement this, one would notice evolutions in vehicle form to reduce drag, hence extend endurance.

Other developments which could make substantial improvements in work capabilities are;

- integrated navigation packages which include multi beam, multi spectral, electronically scanned 3 dimensional sonar to allow auto positioning in space
- 1000 line stereo video with panaroma flat screen.
- Manipulation with teach/playback combined with integrated vehicle positioning.
- Improvements in fibre optic data transmission, already a reality.
- Computer aided piloting (maneuvering).
- NDT packages.
- High definition and low definition acoustic imaging with particular reference to multibeam.

AUTONOMOUS VEHICLES

Autonomous submersible development is now proceeding on two fronts; subsurface and near surface. The notions are not new. For example Tesla conceived an autonomous vehicle in the late 1800's. However, with the exception of torpedoes, the untethered vehicles shown on the x and z axes of Figure 3 are relative newcomers to the scene. Although they have capacity for substantial autonomy, they are for the moment often supervised. At the present time, part of the reason for emphasis on hands on control/supervision is psychological. As time passes, the perceived need for hands-on control (or supervision) will diminish. Current examples of autonomous vehicles are enumerated in Figures 4 and 5. The driving force for autonomous vehicle developments is the potential for cost reduction of work tasks such as surveying. Cost reduction is important both in the commercial and military sectors. Two examples of these vehicles are DOLPHIN figure 6, and ARCS figure 7. The DOLPHIN is an 8 meter long, 2500 kg, 15 knot radio remote controlled and reporting mini snorkeling submarine. ARCS is a 7 meter long, 2000 kg, 5 knot battery powered vehicle.

Areas where technological developments will make substantial improvements in work capabilities are;

- additional memory
- greater processing speed and power
- vision
- better knowledge of networking
- communication strategies to allow for supervised manipulation
- general programs to provide primitive intelligence

Current iterations of autonomous vehicles are proceeding along the x and y axes of Figure 3. However, one can expect hybrids to be produced in the medium term. For example;

- ARCS/DOLPHIN hybrid-diesel battery for oceanographic use.
- Untethered manipulation both by teach and playback, preprogrammed plus machine vision bandwidth compressed video.
- Deep ocean vehicles for exploration and sampling which can act like ROVs or ARCs.

NETWORKING AND AI

Multiple vehicle operations have commenced and we can expect networked multi-vehicle operations such as shown in the sketch Figure 8. Using this approach, groups of relatively low cost vehicles might perform surveys, undertake harbour defence operations or carry out mine countermeasures tasks. Computer supervised autonomous multi-vehicle operations will not only have the advantage of lower cost, they will also have the luxury of being able to respond at computer speed, not man-in-the-loop speed. Man-in-the-loop slows response as the man always has to verify/confirm result as he attempts to ensure that he will not be injured by the contemplated action of the platform he is being transported on.

Recently, considerable attention has been given to artificial intelligence (AI) and its application to robotics. This of course extends to autonomous vehicles. But what does this really mean? Marv Minsky is purported to have defined AI as, "The science of making machines do things that would require intelligence if done by man". If we look in the dictionary, we find intelligence is defined as, "The capacity to acquire and apply knowledge", "The faculty of thought and reason". At this point, when we can not even duplicate the capabilities of insects, birds and/or lower animals with AI, it may be a bit grand to presume that we are hot on the trail of what can be done by man. Additionally, it appears that there is confusion concerning the differences among cognition, adaption, reflexes, prehension and maneuvering. It is not true to conclude that by having better cognition than our mates on this planet, that we can do everything better. Quite the contrary. Our reflexes, maneuvering capability and persistence of vision appear to be inferior to that of birds. Additionally, our preoccupation with cognition leads us to be confused re the transitions from parallel, massive parallel, and networking or logic to fuzzy logic. Some of Frank Busby's thundering herd are in evidence here.

SYSTEM IMPLEMENTATION

If a system is not a harmonious union of parts then it will not be successful, or it may only be marginally successful in the market place. When considering subsystems and elements in isolation, we often overlook this consideration. The systems integration which represents the physical manifestation of all notions is based on mission, available science and experience. These multidimensional integration problems are never uniquely solved. We must rely on successive iterations and judgement to achieve convergence. This approach in systems evolution and system mutation is one of our most useful tools. This means it is mandatory for us to exercise our capabilities if progress is to be made. Perhaps, in order for technology to evolve, successive generations have to be produced; analogues in some ways to Darwin's notion of evolution. Additionally, exercising one's capability is a sine qua non of development/evolution. Considering more iterations have occurred in the marine field, one could conclude that some of the future outerspace technology could be drawn from innerspace technology where "n"th generation telepresence and first generation robotics already exist. It is interesting to note that the subsea evolutionary approach to systems development has been relatively low cost because of commercial considerations.

In the space program, the cost of each iteration is so large and the consequences of failure so unfavourable that progress is being limited. Perhaps they could follow the lead of subsea engineers and improve on their time rate of implementation by moving more in the direction of unmanned vehicles which produce simultaneous reductions in cost and risk.

CONCLUSION

Although we have been in a revolution since the mid 60's, it appears that the revolution will not run out of steam in the near future. Those of us looking for a challenge have considerable opportunities left to contemplate. The proliferation of systems and derivatives will continue as we attempt to map the space, or perhaps more correctly, the innerspace. Thus the metamorphosis of underwater work suites will move relentlessly on.

FIGURE 1

	Tethered	Untethered
Manned	Habitat Sun Guppy Jim/Wasp Mantis Hard Hat Diving	Pisces PC16 Scuba Diver
Manned Remote	RCV 225 Scorpio Hydra	Eave Dolphin Arcs
Unmanned	Torpedo	Epaulard Torpedo Arcs

FIGURE 2

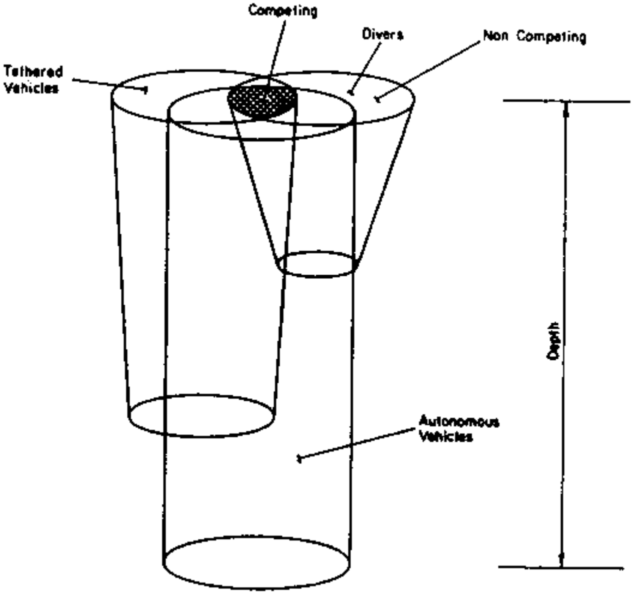

The intersections and forms of cylinders would differ with task set and environmental constraints.

FIGURE 3

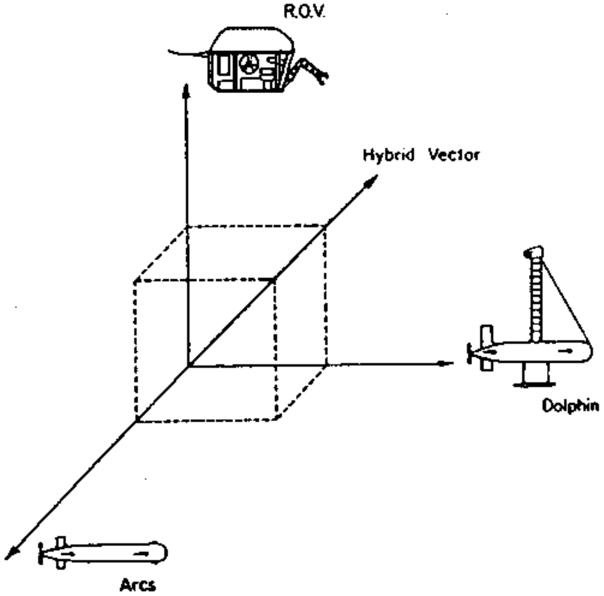

Figure 4

AUTONOMOUS VEHICLE PROGRAMMES

PROGRAMME	COUNTRY	COMPANY	STATUS	APPLICATION	REMARKS
AUV	USA	Rockwell-NRL	Inactive	experimental	
RHADS	UK	Various	Inactive	MCM	Vehicle never built
EAVE(EAST)	USA	UNH	Advanced Development	Scientific	Operational trials 1986
RUMIC	USA	NCSC (USN)	Inactive	Classified	Vehicle never built
AUSS	USA	NOSC (USA)	Operational	Salvage	Development complete June 1985
ARCS	Canada	ISE	Advanced Development	Under ice survey	Operational trials September 1985
DOLPHIN-SEALION	Canada-USA	ISE	Development	Hydrographic, Geophysical Survey, MCM and other countermeasures	Program start September 1986
SUPER EPAULARD	France	Societe ECA	Development	Seabed Exploration	

Figure 5

UNTETHERED VEHICLE PROGRAMMES

PROGRAMME	COUNTRY	COMPANY	STATUS	APPLICATION	REMARKS
EPAULARD	France	Societe ECA	Operational	Seabed Exploration	1 vehicle Development Continuing
PINGUIN A1	West Germany	Messerschmitt (MBB)	Advanced Development	MCM Route Survey	Operational in 1986
DOLPHIN SEA LION	Canada	ISE Ltd.	Operational	Hydrographic and Geophysical Survey	6 vehicles Development Continuing

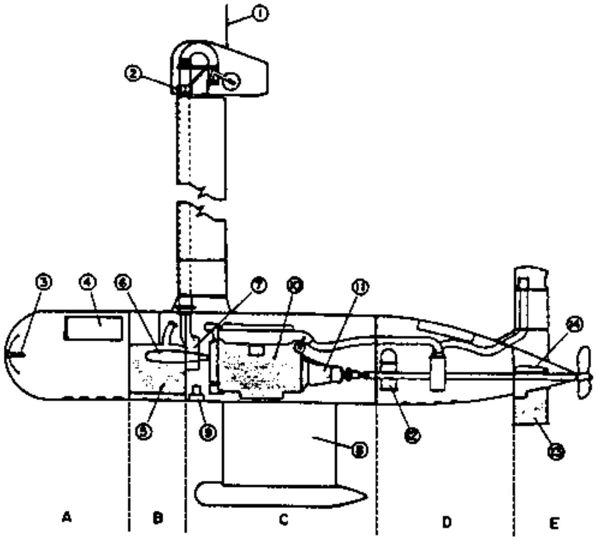

Figure 6

Schematic representation of DOLPHIN.

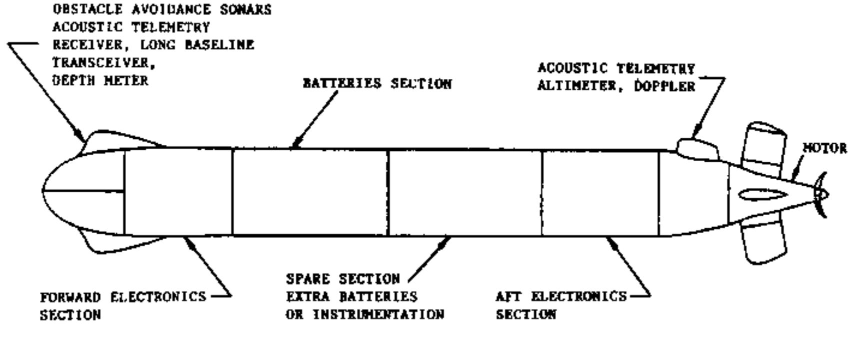

FIGURE 7
ARCS PROFILE

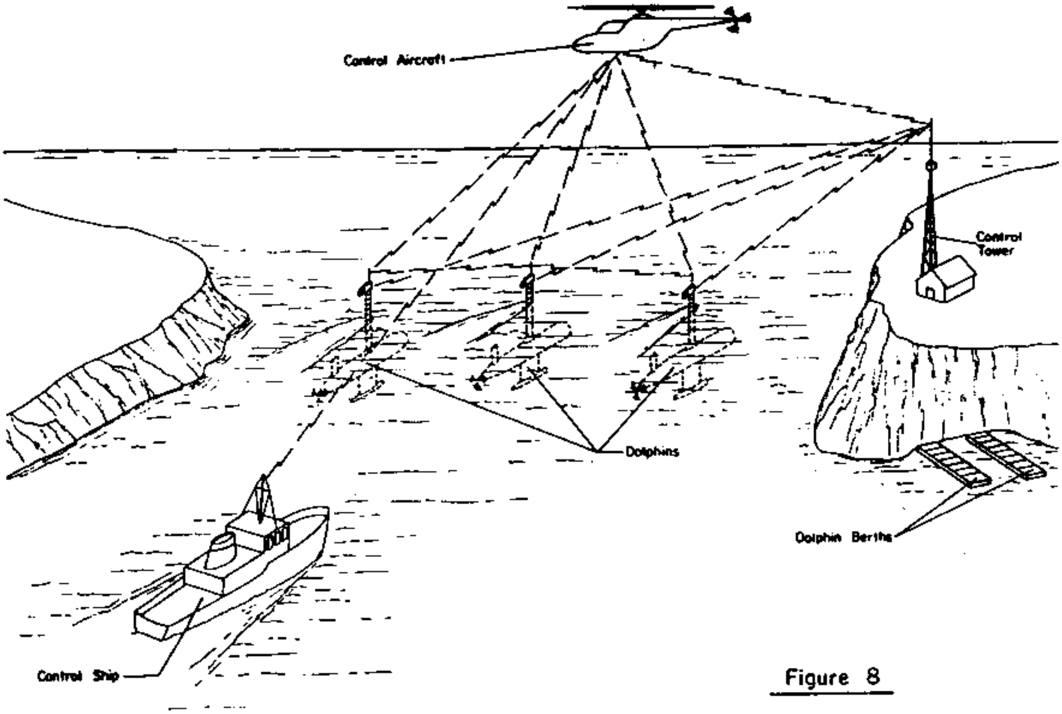
Control Aircraft
Control Tower
Dolphins
Dolphin Berths
Control Ship

Figure 8

The Fourteenth MIT Sea Grant College Program Lecture

Autonomous Underwater Systems: An Advancing Technology for the 1990's

Robert W. Corell
Director
Sea Grant College Program
University of New Hampshire

"In both high and low-technology products, success in the global market means creating and applying new knowledge--which is to say new technology -- faster than one's competitors. This is the fundamental law in this competitive world."

Erich Bloch, Director
National Science Foundation

"U.S. industrial undersea vehicle technology has come to lag behind other industrial nations. For example, the French, Japanese and British have developed a substantial lead to both shallow and deep ocean vehicle technology... The technologies to support the development of offshore mineral and gas/oil deposits have shifted from the U.S. to other nations."

Special 1986 Report to NOAA

I. A NATIONAL CONTEXT FOR ADVANCING UNDERWATER TECHNOLOGY DEVELOPMENT

The context for advancing underwater technology development is set in both an economic competitive framework and in a geopolitical environment. Probably at no time in the history of our nation has the need to think strategically about the oceans been so apparent as it is now. Economic, political, societal and scientific issues of critical importance to the nations of the world are increasingly global and ocean-connected (e.g., international corporations, Chernobyl and similar events, "instant" media coverage of international developments, etc.). This perspective demands a scope of thinking heretofore limited to a few issues. The oceans, particularly those areas within the Exclusive Economic Zone (EEZ), require of us just such a comprehensive point of view. The importance of the EEZ is well outlined in a recent (1986) Special Report to the President and the Congress by the National Advisory Committee on the Oceans and Atmosphere.

> "President Reagan's proclamation (EEZ) brings an enormous frontier region within the Nation's purview. Present sources of minerals and energy resources that are critical to our continued role as the leader of the free world are now contingent on volatile political conditions in regions such as South Africa and the Middle East. Concern is also focusing on the trade deficit resulting, in part, from increased foreign commerce in these vital minerals and energy resources. For economic and political reasons, the Nation must now address exploration of this frontier with a well-conceived and comprehensive scientific exploration plan."
>
> National Advisory Committee
> on Ocean and Atmosphere
> June 1986 Special Report

The need and opportunities for new autonomous underwater technology is set in and driven by such a perspective. The opportunity for economic and industrial growth, new developments in national security and international affairs and the creation of new EEZ legal and jurisdictional regimes provides unprecedented opportunities for and challenges to our Nation's abilities to exploit a broad range of oceanic environments. The lengthy and extremely costly search for the critical parts lost in the tragic Challenger accident, the inability to locate vital parts (i.e., the flight recorders) lost in the Korean Airline incident, the inability to effectively and immediately deal with the massive losses of crude oil in offshore drilling failures and the vulnerability to our national security created by an increase in Arctic operations; all are poignant examples of our need for new oceanic technologies. Our ability to work in remote and deep locations in the world's oceans is profoundly inadequate.

The explosive developments in complex and sophisticated microcomputer and microelectronic-based systems is probably the key to a more realistic appraisal of the useful potential of new underwater technologies. Microcomputer systems, with their attendant microelectronics and powerful soft/firmware provide a potential for addressing deep water and remote operations heretofore either difficult or impossible. Acoustic tomography, sophisticated navigation systems, underwater imaging systems with extended range and resolution, remotely or autonomously operated underwater vehicle systems, and complex "smart" instrument systems are all examples of the impact that this new and evolving technology is having on the ocean and its development.

To provide a more thorough perspective on this topic, let us review some of the key issues that, in my opinion, will drive underwater technology development during the next ten to twenty years. Exploitable natural resources and national security interests are probably the most significant factors driving our national interest in developing new and advanced underwater technology. The EEZ, for example, provides enormous opportunities for the nations of the world, as serious economic and industrial development problems and important societal needs can be addressed there and resolved. For the developed nations, these ocean regions provide new opportunities and new bases for future growth and development. For the developing nations, these regions hold considerable promise for natural resource development that has historically been a basis for the growth of many of the now developed nations of the world. In my opinion the EEZ will increasingly dominate international affairs, and the oceans will play a much more central role in the everyday lives of its citizens. The shear numbers of human population (1984 World Bank predictions suggest a doubling by the year 2100) will focus increasingly our attention on ocean resources and the contribution that the oceans can play in the normal conduct of world affairs. The consequence of this phenomenal population growth is new stresses on world affairs and a heightened role for the ocean, i.e., transportation, communication, disposal of wastes, and as a source of both food and raw materials. It is from these economic and political perspectives of thought that one can identify potential driving forces that will probably pace new ocean and offshore technology and drive its development. These include:

- OFFSHORE ENERGY DEVELOPMENT: The development of deep offshore oil and gas provinces and the establishment of some new and advanced technologies such as OTEC. New provinces and opportunities have been identified.

- DEEP SEABED WASTE DISPOSAL: Developing subseabed repositories for highly toxic nuclear and chemical wastes. The momentum of interest and geopolitical forces are increasingly focusing attention on this alternative to a critical problem for all nations of the world.

- OCEANIC MINERALS DEVELOPMENT: Development of oceanic sources for selected and highly sensitive strategic minerals and the longer term development of some economic minerals. While longer term in realistic potential, the nation's interest in this area of oceanic development is increasing rapidly.

- ENVIRONMENTAL ASSESSMENT AND NATURAL HAZARDS PREDICTION: The ability to assess and predict environmental quality (e.g., the assimilative capacities of near shore oceanic waters to toxic substances) and the impacts of natural hazards (e.g., storm surges, El Nino, etc.) that occur within the oceans will be important to the Nation.

- OCEANIC RESEARCH AND TECHNOLOGY FOR NATIONAL SECURITY: There are national security interests that ocean technology must address and they importantly impact our defense posture. Remote locations such as the ice covered arctic are dramatically changing our defense strategies, and require new and yet to be adequately developed undersea technologies.

- OCEAN SCIENCE RESEARCH: Scientific research to assess the physical, chemical, biological and geological characteristics and behavior of the oceans (particularly the EEZ regions) provides an essential foundation for the development of new oceanic technologies. The interplay between science and technology provides a rich environment for the enhancement of both.

- SURVEYS, MAPPING AND SITE ASSESSMENTS: Detailed and local surveys of bottom topography, fine scale mapping and comprehensive site assessments will increasingly need to be available to those working the EEZ and the deeper oceans. Improvements and increased cost effectiveness for these is essential.

- o ARCTIC PROGRAMS AND OPERATIONS: Offshore oil and gas development in the U.S. arctic regions, pipelines for such development and the national security needs press U.S. technology capabilities for arctic operations.
- o HAZARDOUS OPERATIONS: Hazardous underwater operations, ranging from the need to work remotely near underwater nuclear dump or accident sites to the need to defuse unexploded ordinance, increasingly requires undersea technologies that currently are inadequately deveolped.
- o ROUTINE SURVEILLANCE AND INSPECTION OPERATIONS: There a are variety of fine scale and detailed inspection operations (e.g., long and enclosed aquaducts and sluiceway tunnels) and oversight surveillance operations of underwater applications (e.g., photographic/TV reconnaissance of a malfunctioning deep water wellhead "christmas tree") that would be substantially improved if new underwater technology for operating deep and remotely were available.

These economic and socio-political issues set a context for establishing engineering research and technology development priorities for new and advanced autonomous underwater systems, and provide a driving forces to advance and develop these technologies. These issues, set in a context of our "national interest", provide a framework for our discussions of the more detailed aspects of autonomous underwater technology, and for the beneficial role they potential play in our nation, its industry and commerce, its governmental agencies, and its institutions of research and education.

II. STATE OF THE ART--A HISTORICAL PERSPECTIVE

Research and technology development for autonomous underwater systems has experienced explosive development in the past decade or so. While a comprehensive analysis of the state of the art might be helpful, I want to limit the discussion here to several thoughts that further set a context for the central purpose of this paper, which is to discuss current trends and patterns in research and development, and to suggest some thoughts about where this field is going over the next decade or so.

It is always difficult, and sometimes dangerous, to suggest how a new field got started and when. Taking a very general definition for autonomous underwater systems, we have been at this business for about one hundred years, as the first "autonomous" underwater vehicle was probably the Bristish Whitehead torpedo, developed in the 1860's.

There were earlier torpedos, but the Whitehead was probably the first that had some onboard control systems that qualified it as "autonomous" by our current standards and definitions. It had depth, trim, and directional controls, though like most torpedoes, it was pre-programmed and hence was a "dumb" autonomous underwater system. About the same time (the 1850's or so), the need for underwater telegraph and telephone cables set the stage for another class of underwater work system, the tethered bottom-crawler and/or towed bottom cable laying system. This technology expanded with the advent of buried cables, a concept of the 1930's. It is a highly developed and effective technology today, with the SEA PLOW series of the U.S. Bell System probably the best example.

Early torpedoes and cable laying systems, while autonomous underwater systems , did not provide a sufficient need for "smart or intelligent" autonomy to drive the development of autonomous underwater systems as we think of them today. Therefore, autonomous underwater systems with some on-board capabilites to adjust and adapt to changing conditions in the environment in which the system operates is a phenomena of the past decade or two. National security interests and the offshore industry provided the necessary driving force for the development of such tools. Those needs combined with available new technologies changed the development patterns markedly. Interests expanded dramatically to exploit new markets and take advantage of new technologies (i.e., u/w television, solid state electronics, etc.). Dimitri Rebikoff's POODLE, the U.S. Navy's CURV, and the development of a variety of work vehicles for the needs of offshore industries, particularly the oil and gas industry, paced the early development of autonomous underwater work systems. These remotely operated vehicles (ROVs) were almost always tethered. It is interesting to note the growth trends. In the twenty year period after Rebikoff's POODLE (1953), 85 percent of the tethered ROV's built were government funded and operated. However, during the next eight years, 96 percent of the 350 vehicles built were industrial efforts. The industrial grow was dramatic. In 1970, only one company produced a tethered ROV; today over 25 companies, worldwide, produce tethered ROV's. Three capabilites characterize these vehicles: (i) remote underwater television and photographic capabilities, (ii) remotely operated manipulator capabilities, and (iii) remote man-machine control systems. Virtually all the ROVs are a mix of these three technologies. ROVs are now operating in virtually all the oceans and coastal regions of the world, and as the TITANIC expedition demonstrated, at great depths. It is an operational and effective technology. These tethered, free-swimming vehicles pioneered autonomous underwater systems technology, as they filled a real need in the offshore industrial and miltary marketplace.

In the 1970's, the untethered underwater vehicle systems or autonomous underwater vehicle system (AUVS as some call them), began to appear, and the need to address the concept of autonomy became apparent and real. University research laboratories became interested in these more autonomous vehicles, through their work on such vehicles as the SPURVs at the University of Washington, the Robot Submarine at MIT, the EAVE series at UNH, the Rover at Heriot-Watt University in Scotland, and others. Government and military research laboratories expanded into the untethered arena, with work in the U.S. on the B-1, AUSS, NOSC/EAVE, RUMIC, UFSS, and others to note a few. International activities expanded as well, with the French EPAULARD pioneering an operational free-swimming and untethered vehicle. The USSR SKAT is an untethered vehicle, though we know little about its details. There has been industrial activity as well, such as the ARCS at ISE in Canada, the VERA by CEA & France-Dunkerque Shipyard in France and an unnamed vehicle by CSF Thompson in France. Maybe more importantly, there has been a dramatic growth of interest in industry in the underlying technology (e.g., command and control strategies, KBS architectures, etc.) essential to the development of operational autonomous underwater vehicle systems. One measure is that only four (4) companies participated in the First International Symposium on Unmanned Untethered Submersible Technology in 1980, while seventy (70) participated in the fourth in that series of symposia in 1985. In the U.S., the companies with active autonomous underwater technology programs reads like to "Who's Who in american industry. All of these programs of research and technology development are addressing basic questions essential to exploiting the basic concepts of autonomy and the ways to evolve from the concepts of and experience with pre-programmed and single mission directed systems to systems with expanded autonomy and on-board "intelligence".

There is another class of autonomous underwater systems that is evolving which may be of equal importance to the mobile vehicle systems, a class which we might call autonomous underwater platforms and instruments (AUPIs). There has long existed scientific, military, and industrial interest in underwater platforms and submerged instrument systems (fixed and generally not moved) that operate unattended and often quite remotely, e.g., deepwater seismometers, acoustic arrays, offshore wellhead and completion systems, and a variety of oceanographic instrumentation systems. These systems have increasingly become "smarter" and adaptive. While immobile, they often have many of the characteristics we assign to autonomous underwater vehicle systems, i.e., sensing and adapting to changing environmental or operational conditions. Like the AUVS, these systems are becoming more robot-like.

This historical perspective, suggests that autonomous underwater systems might be categorized into five generic types of systems.

- o Towed Vehicles (taut/load-bearing tethered systems)

- Seafloor and Structure Work Vehicles (tether often slack)
- Tethered, Free-Swimming Vehicles (tether slack)
- Untethered, Free-Swimming Autonomous Vehicles
- Autonomous Underwater Platforms and Instruments (untethered)

While each of these systems presents a series of important and critical research and development opportunities and problems, the remainder of this paper is devoted to a discussion of issues relating to generic research and development issues primarily associated with untethered, free-swimming autonomous vehicles and autonomous underwater platforms and instruments.

III. TECHNOLOGY ISSUES - PATTERNS IN RESEARCH AND DEVELOPMENT

Untethered and highly autonomous underwater vehicles systems, and to a lesser extent autonomous platforms and instruments, rely upon the availability and effectiveness of several key technologies. The ten technology areas listed below, provide a framework for our discussions and an opportunity to outline some important research topics central to the evolution of the autonomous underwater system.

- Work Systems and Mission Planning
- Control and Guidance
- Microcomputer and Microelectronic Systems
- Artificial Intelligence and the Computer Sciences
- Manipulation and Man/Machine Systems
- Communications
- Vision Systems
- Navigation and Sensors
- Power Sources and Systems
- Simulations, Modeling, Testbeds and Prototypes

The maturity of the technology in each of these areas, as it relates to untethered and autonomous underwater systems, is markedly different, and ranges from adequate to virtually undeveloped. The issue of autonomy in underwater systems requires major advances in several of these technologies before these systems will make significant contributions to the needs and opportunities outlined in Section I. These technology areas, in my opinion, are critical to the future of autonomous underwater systems, vehicles and platforms.

WORK SYSTEMS AND MISSION PLANNING

The design and development of autonomous underwater systems, whether vehicles or fixed platforms, is ultimately influenced by the missions, the applications, and the work scenario. Autonomous underwater systems present a variety of challenges in this regard. While we could discuss many topics that impact this aspect, three topics are of sufficient importance that they deserve special research and development efforts.

Automated Plan Generation

An important function of an intelligent system is that it be able to generate its own plan of action once it is given a high level description of it's mission. The system must then reason about mission goals, the system's own capabilities and limitations, and the environment. This planning process must be invokable at any time during a mission in order to provide for real time planning in the face of problems or unexpected events. In a real-time system, the planning process is complicated greatly by the constraints imposed by time critical functions. The problem is further complicated by the fact that the time constraints may vary due to a specific situation or in a certain context (e.g. the altitude maintained by a system is of far less concern if the bottom is flat and soft as compared to hard and variable). An important variant to automated plan generation is three dimensional trajectory planning. A mobile submersible system must be able to plan its own motions in three axes as well as its attitude about at least one of those axes. This ties in closely with the system's world model, which represents the system's current position, attitude, and velocities in the physical world.

Ability to Cope With Unanticipated Situations or Events

Missions previously considered too dangerous for manned systems are now being contemplated as new system concepts become available. The success of such systems will only be realized if they are capable of demonstrating sufficient on board decision-making capability to function in unexpected situations. This does not imply that an autonomous system must deal with all and every imaginable situation. It does suggest that the system be capable of coping with many unanticipated situations or events and be able to respond in a reasonable manner to achieve a desired goal. An unpredictable event might be defined as one of unknown timing, magnitude or relationship, but of a known class of events such that it can be adequately described prior to the operation. Strategies then can be created to cope with the event, risks and costs, and mission requirements.

Cooperation Between Multiple Agents

A potential benefit of autonomous underwater systems is that a support vessel or remote control station might support multiple autonomous underwater systems. Large area, single pass surveys, reconnaissance, surveillance, or other tasks where multiple intelligent interdependent systems can be used requires cooperation between multiple systems. If multiple systems can communicate and cooperate with each other autonomously, then there is the potential to increase the overall mission effectiveness. This force multiplication capability of autonomous systems may indeed be one of the more valuable characteristics and offer the biggest potential reward for intelligent underwater systems technology.

CONTROL AND GUIDANCE

The mobile autonomous underwater system presents fascinating control and guidance problems. The architecture for control and guidance of these autonomous underwater systems presents a set of special opportunities and problems. Much research and development effort has been directed at developing architectures for autonomous systems. More recently some attention has been given to the special problems associated with underwater applications. The control and guidance requirements for such systems present a special set of conditions, which must be taken into consideration:

1. The mission requirements for an underwater system,

2. The unique nature of the underwater environment,

3. The need to operate remotely in real time,

4. The state-of-the-art of appropriate artificial intelligence constructs, and

5. The realities the of the state-of-the-art hard/software.

Recent research suggests that the AI constructs of knowledge-based systems (KBS) and other constructs of knowledge engineering have the potential of advancing the state of the art in a variety of autonomous vehicle guidance and control applications. A number of architectures have been proposed for underwater vehicle systems. It is increasingly clear that these architectures will be hierarchical and probably based on knowledge-based system concepts. Further, the architecture should accommodate the fact that some functions (e.g., control outputs to effectors) must be accomplished on a scale of fractions of a second while high level path planning functions are several orders of magnitude slower. Finally, the architecture should accommodate the fact that information must be separated according to the level of importance and role. For example, some functions require virtually continuous data streams (e.g., control signals to effectors), others utilize threshold data (e.g., system response when the power source/plant reaches minimums), while still others might only consider trend information. The constructs of a knowledge-based systems architecture lend themselves to these requirements and systems. The general functions and components for a knowledge-based guidance and control architecture for intelligent underwater systems are suggested in Figure 1. A number of questions are still unresolved and will require research effort, such as, (i) means of efficiently communicating between the various modules of a hierarchical architecture, (ii) interfaces between symbolic and numeric languages as well as the implications of using symbolic languages in real-time environments, and (iii) dealing with incomplete and uncertain data.

MICROCOMPUTER AND MICROELECTRONIC SYSTEMS

Microcomputers and microelectronic components probably have had more impact on the potential for and the development of autonomous underwater systems than any other technology. The low power requirements of CMOS electronics, the remarkable compute power of the new 16 and 32 bit microcomputer chips, space efficient and megabyte capable solid state memory, and powerful languages and attendant software have made an enormous difference in our capacity to seriously consider the opportunities and problems associated with this field of underwater development. The prospects for future advances in this arena are incredible, with the speeds of computation, measured in million instructions per second (MIPS), increasing from a few MIPS now to several 10s of MIPS to possibly 100 MIPS or more in the future. The concepts of multiprocessing and parallel processing combined with major developments in chip capabilities are likely to make these predictions a reality in the next few years. These developments must become realities before any major breakthroughs in autonomous underwater vehicle technology is truly operational, reliable, and highly capable. Compute speeds must increase by an order of magnitude or so, and RAM must be greatly expanded. The developments in "smart" and "intelligent" systems require such microcomputer performance capabilities. These performance levels should be available to our research and development laboratories with the decade. Software and language developments are paralleling those of hardware, largely because of major developments in ubiquitous personal computer field. These developments combined, give rise to major opportunities in the hard/software area for autonomous underwater systems.

ARTIFICIAL INTELLIGENCE AND THE COMPUTER SCIENCES

While microcomputer developments have provided the tools for major advances in underwater systems capabilities, the conceptual framework for the autonomy we seek for autonomous underwater systems lies in recent developments in the fields of artificial intelligence and the computer sciences. The advances in computer soft/hardware technologies and the maturity of selected artificial intelligence concepts make it possible to address the requirements of autonomous underwater systems, particularly real-time systems. These possibilities will be achieved when the scientific principles of artificial intelligence are "engineered" into real world applications. The emergence of the fields of knowledge-based systems and knowledge engineering give credence to that goal and potential.

The contemporary artificial intelligence literature is replete with discussions of new and powerful mathematical routines, computer languages, data manipulation processes, and powerful problem-solving techniques. In general, these new concepts have been applied to demonstration and concept validation problems that are familiar to the academic community. There are few instances, however, that illustrate the successful use of AI techniques to enhance the performance of complex real world applications. The potential is clearly there. The research and development questions are many, a few are listed below to give a favor to the trends and possible patterns for AI related research for autonomous underwater systems, particularly vehicles and other mobile platforms.

Data Structures/Databases/World Models

The ocean is a complicated environment. If an intelligent system is to operate effectively in this environment, it must build, maintain and update an internal model of itself and the environment within which it exists (control state, system status, environmental conditions, and mission plans and goals). Research will be needed in the areas of data structures, knowledge representation, and database manipulations before such a world model becomes a reality. The model probably should express the physical realities of the environment as viewed through the system's sensor suite, as well as the assumptions and conclusions made by the system.

The information to be represented is very diverse: depth, lateral placement, velocities, objects and obstacles, system resources and limitations, conclusions based on sensor fusion, and on inexact reasoning (assumptions) or historical trends and dependencies. Although some effort has been directed at the problem of describing a world model for an underwater system (Westinghouse/CMU, MIT work area model), there has been little published.

Methodologies for Machine Representation of Knowledge

Research should address those issues which involve the acquisition (from on-board sensors), organization and encoding of system knowledge in a form and format suitable for machine processing. The representation scheme might be structured such that on-board knowledge can be modified or updated as new information is developed. Data structures, database management concepts, and developing an appropriate world model for vehicle system are the kinds of questions which should be addressed. Further, the types of knowledge required, how the knowledge is to be represented and how the knowledge is to be used are key problems which need to be studied. Research in our Laboratory has led us to consider three types of system knowledge. The first type is "knowledge" derived from data assessment, i.e., information from the sensors and from systems status. The second type describes those actions which are derived from the problem solving and planning functions. The third type seeks to describe the system and the world within which it exists. Methods will need to be developed to represent these kinds of "knowledge" for the system.

Symbolic to Numeric Language Interface

An autonomous system will have to utilize both numeric data (from sensors) and symbolic information. The numeric data can be processed in traditional algorithmic programming languages while the symbolic information is best processed in a symbolic environment such as LISP. In order to realize intelligent autonomous systems, a method of interaction and communication must be developed to interface these two dissimilar types of computations. The interface must provide for the ready exchange of information and control signals. The disparity between these two types of computation is great enough to favor an interface, as opposed to developing methods of forcing one type to do both.

MANIPULATION AND MAN/MACHINE SYSTEMS

Research in this area will be discussed by several others at this Symposium, but a comment seems appropriate. One of the central questions we must address is devising ways to effectively and efficiently achieve control functions autonomously in light of the fact that the vehicle and the manipulator subsystem are highly interactive and interdependent. The motion of the manipulator causes a displacing force on the vehicle, and conversely, vehicle movements directly alter the positioning accuracy of the manipulator. Supervisory control and recent developments by Yoerger, Sheridan, and others show considerable potential for addressing this serious problem for the quasi-neutrally stable underwater system.

COMMUNICATIONS

Communications between an autonomous underwater system and the surface, or with another system or bottom relay station is critically important to the usefulness of these systems. If there is a limiting technology in the application of autonomous underwater systems, it is the present inability to communicate with free-swimming vehicles or autonomous platforms (i.e., without a tether or other physical link). While a number of acoustic communication experiments have been conducted, no reliable and reasonably broadband underwater acoustic communication links exists. There are, however, acoustic releases and offshore oil field acoustic communication links that provide low data rates (few tens to hundreds of bits/sec.) for communications with remotely located autonomous underwater systems. The needs, however, for effective use of autonomous underwater vehicles and systems require effective baud rates of a minimum of a few hundred (300 to 500) for basic command and control information, and a more useful range of 10 to 20 kbits/sec for transmitting video and other high density data streams. An "acoustic tether", with these kinds of data rate capabilities, would dramatically increase the usefulness of autonomous underwater systems, and would "open-the-door" for a wide range of experiments with such vehicle systems. Some progress has been made. The trade offs that must be made are many. If range is kept to a few hundred meters, the acoustic carrier frequencies can be quite high (i.e., a few hundred khz to a megahz or so), while the need to work over larger areas or into deeper depths (i.e., 1000 meters and beyond) places severe restrictions of acoustic operating frequences. Field experiments have been conducted, over essentially vertical paths, in water depths to 14,000 ft. and with carrier frequencies of 14 khz. The results are very promising, with reported error rates of about 1 error per million bits transmitted within a 45 degree vertically oriented cone. While communications over essentially vertical paths is vital, communciations over horizontal paths is equally important. The vertical path problem, particularly in deep oceanic environments is tractable, and probably will be routinely available during the next decade or so. Horizontal communication paths are a much more difficult problem, particularly for distances which exceed the water depth. Multipath becomes a serious problem. Therefore, most of the work in this area has been limited to short path lengths (few tens to hundreds of meters) and at higher frequencies (above 50-100 khz). At these frequencies and ranges, it is possible to communicate acoustically over horizontal paths.

Two other communications strategies should be further studied. Electromagnetic communciation links are extremely attractive for certain applications (i.e., shallow fresh water lakes). Dunbar and others have reported encouraging results for these very short range and special applications. Further, work has been underway for almost a decade on fiber optic links, both of the type that can be installed in a ROV tether (such as is planned for the Argo-Jason vehicle at WHOI), and of the disposable type that might be used on an "autonomous" vehicle. These types of links are extraordinarily attractive for they solve the communication problems with their incredible bandwidths. The disposable tether must address most of the normal problems associated with a tether. These and other reported results give promise to markedly improved communication links between autonomous free-swimming vehicle systems and the surface.

VISION SYSTEMS

One of the most significant and important roles that ROVs play is there potential for and ability to provide remote vision capabilities. Obtaining such images from underwater scenery presents designers with several difficult problems, i.e., backscattering, high attenuations rates, color distortions, and poor and difficult lighting conditions. Further, deducing three dimensional features places severe restrictions on vision system requirements. There has been substantial progress in underwater video systems (e.g. low light level CCD cameras) and in cinematography and still photography. While evolutions in these technologies will substantially improve autonomous underwater systems capability, the central issue remains the inability of these systems to provide such images in real time. The exciting developments in u/w imaging are in recently developed methods to improve the effective bandwidths of the acoustic communication channel and to bandwidth compress the video images so that quasi-real time video images can be transmitted over long distances (several kilometers) on what might be called "an acoustic tether". Much progress has been made on this topic in the past several years. For example, a quasi-real time (4 frames per second) acoustic link video imaging system has been reported and laboratory evaluated. Several studies suggest that a 256 x 256 pixel image (with 6 bits per pixel) is adequate for piloting an autonomous vehicle, provided the frame rate is 4 frames per second or more. Further, it is possible to trade frame rate for resolution so that higher image quality can be obtained by simply reducing the frame rate. Several bandwidth or data compaction schemes have been reported and studies conducted to determine their feasibility for underwater applications, i.e., transform coding, spatial coding, entropy coding, and hybrid coding. All these techniques seek to reduce the redundancy in the raw video image. Experimental results indicate that compression ratios in the range of 10 to 100 are feasible. Analysis of error rates suggest that this problem is tractable, though error correcting schemes will probably need to be employed. The standard 4 Mhz TV signal, with 100:1 bandwidth compression, is transformed into a 40 khz acoustic signal, a realistic possibility for an acoustic tether. There are exciting possibilities for the use of scene analysis and other signal processing scenarios that could possibly bring the effective compression rate into the realm so that 15-20 khz acoustic channels could handle a video image with virtual real time performance (256 x 256 pixels, at 6 bits per pixel and frame rates a 4 frames or more per second). An acoustic communication channel with this kind of video performance would revolutionize the applications for autonomous underwater vehicles and systems.

The direct use of acoustics to image underwater scenes has long been a dream of underwater technology. Limited image capacity obstacle avoidance sonars come as close to acoustic imaging as is now available. These systems are improving in the resolution of detail, and numerous examples exist in which acoustics are used to classify underwater objects. The prospects of using acoustics to assist in path planning and obstacle identification and avoidance are exciting and progress in this area is likely to be substantial in the next decade or so. Such developments when combined with the prospects for quais-real time video images transmitted on an acoustic link, suggest major changes in the application and uses of autonomous underwater vehicles and platforms.

NAVIGATION AND SENSORS

These two topics have been of central importance to underwater vehicle development for a long while. The ability to navigate with reliabilty and accuracy has long been important to research and industrial submersibles, and to submarines and naval underwater systems. During the past decade substantial improvements have been made in this field. Short range (i.e., few hundred meters) underwater navigation systems are available and tested. Modest range (few kilometers or so) systems have been used extensively by the science community for some time, such as the ALNAV system developed for ALVIN. Long range navigation systems have been developed for national security purposes, such as the SIMS system for the nuclear submarine fleet. However, no operational long range (beyond 10 or 15 kilometers or so) underwater navigation systems is available that can be used by the autonomous underwater vehicle community. Research laboratory concepts have been advanced, but limited work has been conducted for such long range systems. The importance of reliable underwater navigation systems cannot be overlooked. The development of AUV systems is not likely to be delayed because of limitations in navigation technologies.

Sensors for underwater systems are critical to the goals of autonomy. However, many of the essential sensor systems are operational and available, e.g., the physical environmental parameters of temperature, depth, current velocity, attitude, compass bearing, water turbidity, etc.. The senors for vision are, of course, critical to thoses systems and recent developments in low light level video cameras provide real promise in this area (e.g., effective ASA rating of 50,000 to 200,000) have been tested. Research on light-camera separations has improved the quality of images, and post processing of images, using image enhancement techniques, is now starting to be used in underwater vision/imaging systems. Work in sidescan sonar has advanced to the stage where preliminary field studies have been conducted to test the operation of autonomously operated side scan systems. There is a host of sensor opportunities for autonomous vehicle operations, the discussion of which goes beyond the objective of this paper, and deserves a paper totally on that topic. While there are important needs and opportunities, the technology of sensors for autonomous vehicles is not yet limiting the research and development of these systems.

POWER SOURCES AND SYSTEMS

The power required to propel and power underwater vehicles has been a subject of some interest amougst the field of autonomous underwater vehicle technology. Research by the AUVS community in this field has been limited largely because the pacing problems have been field testable with existing power supplies (e.g., lead acid and other batteries). As the mission time and range requirements increase, the need for high energy density power sources will also increase. There is on-going research, driven largely by national security needs, in high energy density batteries, nuclear power sources, fuel cells, closed cycle internal and external combustion engines, and the recent report of a mechanical gill for supplying oxygen to a conventional power plant. Power is currently not a limiting technology to the future development of autonomous underwater systems.

SIMULATIONS, MODELING, TESTBEDS, AND PROTOTYPES

The variety, amount, and dynamics of the data that an autonomous underwater system has to consider is enormous. System development work must address a wide variety of situations and problems. The development of the autonomous systems would be greatly enhanced if more environment and mission simulations and full featured dynamic models existed for evaluation of autonomous underwater vehicle concepts and systems. A comprehensive environment and vehicle simulator, if developed, would be a helpful tool to allow stepwise development of autonomous systems. It would provide the means for intensive laboratory testing of the systems prior to building and field testing prototypes. It would also provide the capability to generate mock problem scenarios to which the autonomous systems must react. The resulting simulation tool would be analogous to the flight simulators. There would need to be man-machine interface to the simulator. The need for full scale field testing of testbed and prototype autonomous underwater systems can not be understated. The success of any major new technology rests on careful component and full scale testing. We, in ocean technology, have had a propensity for underestimating the importance of in situ testing and evaluation.

IV. CONCLUDING THOUGHTS - AUTONOMY AND UNDERSEA TECHNOLOGY

The ideas discussed at this Symposium are exciting and provide realistic potentials for oceanic exploration and development. The thoughts outlined in this paper concerning key technologies for autonomous underwater systems are intended to suggest trends and patterns that will likely apply during the next decade for so. There is, however, an assumption that underpins these discussions; that is that the state of the art of both undersea systems technology and the knowledge engineering/artificial intelligence is mature enough that substantive new technologies can be realized. Why do I feel confident that such is possible?

The growth of the technological base upon which autonomous undersea systems rest has been remarkable. The literature on this topic has grown dramatically, both in the refered journals and in the more informal publications of the ROV Conferences, MTS and OTC Meetings, and the more specialized symposia like this one and the International Symposium on AUVS technologies held at UNH. The proven usefulness of ROVs to the oil and gas industry, and the recent demonstration of ROV technology during the Shuttle search and TITANIC expedition all attest to the maturity of the underlying technologies critical to future growth in the more autonomous systems. The real question is whether the artificial intelligence technologies, maturing in the evolving field of knowledge engineering, will realize their potential and substantively assist in establishing the technical basis for autonomous underwater systems.

Knowledge based applications are developing rapidly, and are demonstrating that the conceptual frameworks are working. The history of knowledge-based systems is moving from one of demonstration to operational usefulness. The Digital Equipment Corporation uses a knowledge-based system to configure virtually all their orders for VAX computer systems. More recently, I saw KBS strategies used to substantially improve the management and processing of environmental acoustics data for the Navy and for research. Just the other day I saw a demonstration of a KBS approach to help weather forecasters better predict and anticipate the existence of microbursts, those powerful downdrafts that have caused several tragic airline accidents in recent years. Further, a fascinating article recently appeared in SCIENCE about an KBS model to assist in cotton crop management. The results were outstanding. The KBS model predicts the harvesting dates for cotton crops, and the economic benefits on one farm were "over $60 per acre on this 6000-acre farm". The list of such KBS applications is growing phenomenally, with demonstrated success in many fields.

The interest and investment by American industry in KBS related research and development for autonomous underwater systems technology is remarkable, and augers well for the success of the applications anticipated. With continued efforts on the research discussed at this Symposium, I am optimistic about the impact that these kinds of efforts will have on our technologies. The intellectual environment is rich with ideas, maturing technologies, and realistic potentials for autonomous undersea systems.

Mobile Robots in Unstructured Environments

Long-term Autonomy for Mobile Robots

Rodney A. Brooks
Assistant Professor
Artificial Intelligence Laboratory
Massachusetts Institute of Technology

Abstract-**The Mobot project at the MIT Artificial Intelligence Laboratory has built two mobile robots—one with a manipulator attached. Rather than being told what to do, these robots are built with a set of parallel processors running, each trying to achieve some task. The robot exists in its environment and reacts to its surroundings in a manner determined by its built-in nature. These insect-like robots are thus able to operate over long periods of time in people-populated environments without needing maps or world models. There are many applications underwater for robots with such control structures.**

1. Introduction

Robots have traditionally been viewed as general purpose devices which can be used to automate tasks previously done by humans. There is an inherent bias in this view. The way in which the robots are instructed to achieve the task often mimics the way in which a human would do the task.

Robots in structured environments (such as factories) have been given repetitive tasks, such as welding and spray painting, with very little sensing, and are often left running for long periods with little human intervention save for scheduled maintenance. Usually, however, there are people near by to help with error recovery.

Robots in unstructured environments (such as underwater) usually do non-repetitive tasks, such as maintenance tasks, but with almost constant human intervention, i.e., teleoperation. The human does most of the perceptual processing for the robot.

Now suppose we change the equations a little. Suppose it is possible to get a robot to do some non-repetitive tasks in an unstructured environment. Suppose furthermore that the robot can operate for long periods of time completely autonomously—this can radically change the economics of using such a robot, without the need for constant monitoring. Then the way in which we get robots to do tasks for us can be radically different.

In this paper we first explore a method for building useful robots with long-term autonomy. Then we speculate on the implications for underwater robots.

2. Requirements

We can identify a number of requirements of a control system for an intelligent autonomous mobile robot. They each put constraints on possible control systems we might build and employ.

- **Multiple Goals.** Often the robot will have multiple goals, some conflicting, which it is trying to achieve. It may be trying to reach a certain point ahead of it while avoiding local obstacles. It may be trying to reach a certain place in minimal time while conserving power reserves. Often the relative importance of goals will be context dependent. Getting off the railroad tracks when a train is heard becomes much more important than inspecting the last 10 track ties of the current track section. The control system must be responsive to high priority goals, while still servicing necessary "low level" goals (e.g. in getting off the railroad tracks it is still important that the robot maintains its balance so it doesn't fall down).
- **Multiple Sensors.** The robot will most likely have multiple sensors (e.g. TV cameras, encoders on steering and drive mechanisms, and perhaps infrared beacon detectors, an inertial navigation system, acoustic rangefinders, infrared rangefinders, access to a global positioning satellite system, etc.). All sensors have an error component in their readings. Furthermore, often there is no direct analytic mapping from sensor values to desired physical quantities. Some of the sensors will overlap in the physical quantities they measure. They will often give inconsistent readings—sometimes due to normal sensor error and sometimes due to the measurement conditions being such that the sensor (and subsequent processing) is being used outside its domain of applicability. Often there will be no analytic characterization of the domain of applicability (e.g. under what precise conditions does the Sobel operator return valid edges?). The robot must make decisions under these conditions.
- **Robustness.** The robot ought to be robust. When some sensors fail it should be able to adapt and cope by relying on those still functional. When the environment changes drastically it should be able to still achieve some modicum of sensible behavior, rather then sit in shock, or wander aimlessly or irrationally around. Ideally it should also continue to function well when there are faults in parts of its processor(s).
- **Extensibility.** As more sensors and capabilities are added to a robot it needs more processing power; otherwise the original capabilities of the robot will be impaired relative to the flow of time.

3. Levels and Layers

There are many possible approaches to building an autonomous intelligent mobile robot.

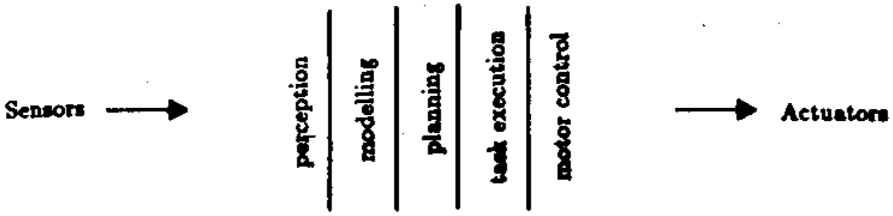

Figure 1. A traditional decomposition of a mobile robot control system into functional modules.

As with most engineering problems they all start by decomposing the problem into pieces, solving the subproblems for each piece, and then composing the solutions. We think we have done the first of these three steps differently to other groups. The second and third steps also differ as a consequence.

3.1 Levels of Competence

Typically mobile robot builders (e.g. [Nilsson 84], [Moravec 83], [Giralt et al 83], [Kanayama 83], [Tsuji 84], [Crowley 85]) have sliced the problem into some subset of:

- sensing,
- mapping sensor data into a world representation,
- planning,
- task execution, and
- motor control.

This decomposition can be regarded as a horizontal decomposition of the problem into vertical slices (see figure 1). The slices form a chain through which information flows from the robot's environment, via sensing, through the robot and back to the environment, via action, closing the feedback loop (of course most implementations of the above subproblems include internal feedback loops also). An instance of each piece must be built in order to run the robot at all. Later changes to a particular piece (to improve it or extend its functionality) must either be done in such a way that the interfaces to adjacent pieces do not change, or the effects of the change must be propagated to neighboring pieces, changing their functionality too.

We have chosen instead to decompose the problem vertically, figure 2 , as our primary way of slicing up the problem. Rather than slice the problem on the basis of internal workings of the solution we slice the problem on the basis of desired external manifestations of the robot control system.

To this end we have defined a number of *levels of competence* for an autonomous mobile robot. A level of competence is an informal specification of a desired class of behaviors for a robot over all environments it will encounter. A higher level of competence implies a more specific desired class of behaviors.

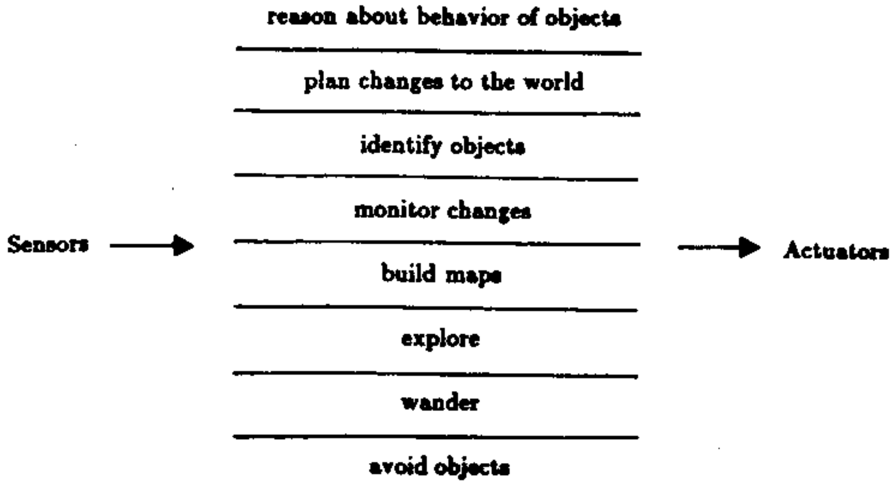

Figure 2. A decomposition of a mobile robot control system based on task achieving behaviors.

We have used the following levels of competence (an earlier version of these was reported in [Brooks 84]) as a guide in our work:

0. Avoid contact with objects (whether the objects move or are stationary).
1. Wander aimlessly around without hitting things.
2. "Explore" the world by seeing places in the distance which look reachable and heading for them.
3. Build a map of the environment and plan routes from one place to another.
4. Notice changes in the "static" environment.
5. Reason about the world in terms of identifiable objects and perform tasks related to certain objects.
6. Formulate and execute plans which involve changing the state of the world in some desirable way.
7. Reason about the behavior of objects in the world and modify plans accordingly.

Notice that each level of competence includes as a subset each earlier level of competence. Since a level of competence defines a class of valid behaviors it can be seen that higher levels of competence provide additional constraints on that class.

3.2 Layers of Control

The key idea of levels of competence is that we can build layers of a control system corresponding to each level of competence and simply add a new layer to an existing set to move to the next higher level of overall competence.

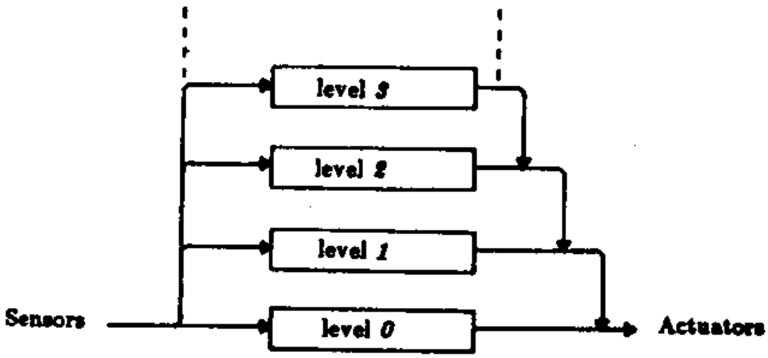

Figure 3. Control is layered with higher level layers subsuming the roles of lower level layers when they wish to take control. The system can be partitioned at any level, and the layers below form a complete operational control system.

We start by building a complete robot control system which achieves level *0* competence. It is debugged thoroughly. We never alter that system. We call it the zeroth level control system, Next we build a another control layer, which we call the first level control system. It is able to examine data from the level *0* system and is also permitted to inject data into the internal interfaces of level *0* ***suppressing*** the normal data flow. This layer, with the aid of the zeroth, achieves level *1* competence. The zeroth layer continues to run unaware of the layer above it which sometimes interferes with its data paths.

The same process is repeated to achieve higher levels of competence. See figure 3. We call this architecture a ***subsumption architecture***.

In such a scheme we have a working control system for the robot very early in the piece—as soon as we have built the first layer. Additional layers can be added later, and the initial working system need never be changed.We claim that this architecture naturally lends itself to solving the problems for mobile robots delineated in section 2.

- **Multiple Goals.** Individual layers can be working on individual goals concurrently. The suppression mechanism then mediates the actions that are taken. The advantage here is that there is no need to make an early decision on which goal should be pursued. The results of pursuing all of them to some level of conclusion can be used for the ultimate decision.
- **Multiple Sensors.** In part we can ignore the sensor fusion problem as stated earlier using a subsumption architecture. Not all sensors need to feed into a central representation. Indeed certain readings of all sensors need not feed into central representations—only those which perception processing identifies as extremely reliable might be eligible to enter such a central representation. At the same time however the sensor values may still be being used by the robot. Other layers may be processing them in some fashion and using the results to achieve their own goals, independent of how other layers may be scrutinizing them.
- **Robustness.** Multiple sensors clearly add to the robustness of a system when their results can be used intelligently. There is another source of robustness in a subsumption

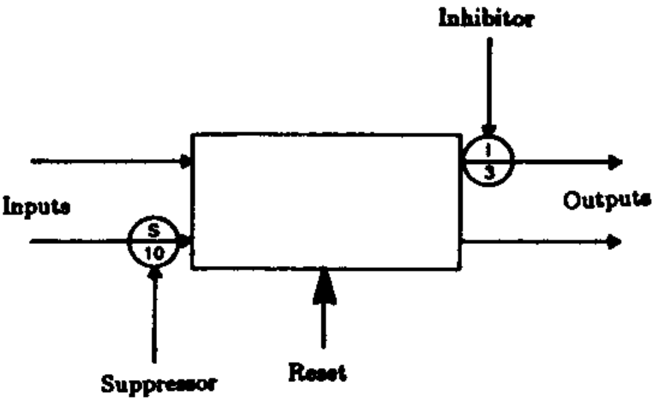

Figure 4. A module has input and output lines. Input signals can be suppressed and replaced with the suppressing signal. Output signals can be inhibited. A module can also be reset to state NIL.

architecture. Lower levels which have been well debugged continue to run when higher levels are added. Since a higher level can only suppress the outputs of lower levels by actively interfering with replacement data, in the cases that it can not produce results in a timely fashion the lower levels will still produce results which are sensible, albeit at a lower level of competence.

- **Extensibility.** An obvious way to handle extensibility is to make each new layer run on its own processor. We will see below that this is practical as there are in general fairly low bandwidth requirements on communication channels between layers. In addition we will see that the individual layers can easily be spread over many loosely coupled processors.

3.3 The Structure of Layers

But what about building each individual layer? Don't we need to decompose a single layer in the traditional manner? This is true to some extent, but the key difference is that we don't need to account for all desired perceptions, processing and generated behaviors in a single decomposition. We are free to use different decompositions for different sensor-set task-set pairs.

We have chosen to build layers from a set of small processors which send messages to each other. A prototypical such machine is shown schematically in figure 4.

Each processor is a finite state machine with the ability to hold some data structures. Processors send messages over connecting "wires". There is no handshaking or acknowledgment of messages. The processors run completely asynchronously, monitoring their input wires, and sending messages on their output wires. It is possible for messages to get lost—it

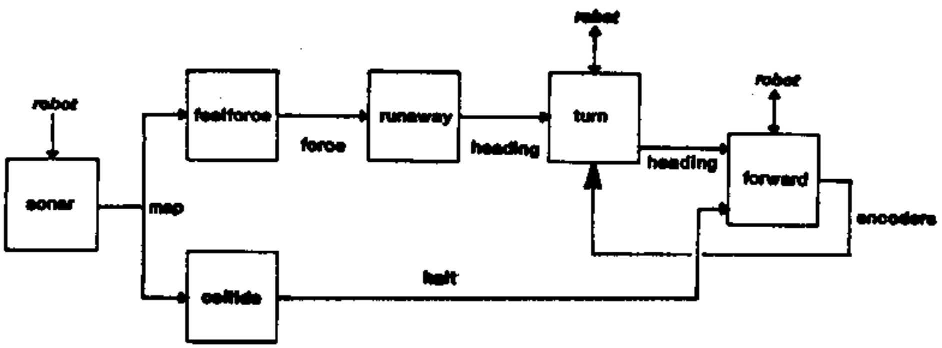

Figure 5. The level *0* control system.

actually happens quite often. There is no other form of communication between processors, in particular there is no shared global memory.

All processors (which we refer to as modules) are created equal in the sense that within a layer there is no central control. Each module merely does its thing as best it can.

Inputs to modules can be suppressed and outputs can be inhibited by wires terminating from other modules. This is the mechanism by which higher level layers subsume the role of lower levels.

4. A Robot Control System Instance

We have implemented a mobile robot control system to achieve levels *0* and *1* and *2* competence as defined above.

Zeroth Level

The lowest level layer of control makes sure that the robot does not come into contact with other objects. It thus achieves level *0* competence. See figure 5. If something approaches the robot it will move away. If in the course of moving itself it is about to collide with an object it will halt. Together these two tactics are sufficient for the robot to flee from moving obstacles, perhaps requiring many motions, without colliding with stationary obstacles. The combination of the tactics allows the robot to operate with with very coarsely calibrated sonars and a wide range of repulsive force functions. Theoretically, the robot is not invincible of course, and a sufficiently fast moving object, or a very cluttered environment might result in a collision. Over the course of a number of hours of autonomous operation, our physical robot (see section 5.2) has not collided with either a moving or fixed obstacle. The moving obstacles have, however, been careful to move slowly.

Figure 5 gives a complete description of how a set of modules are connected together.

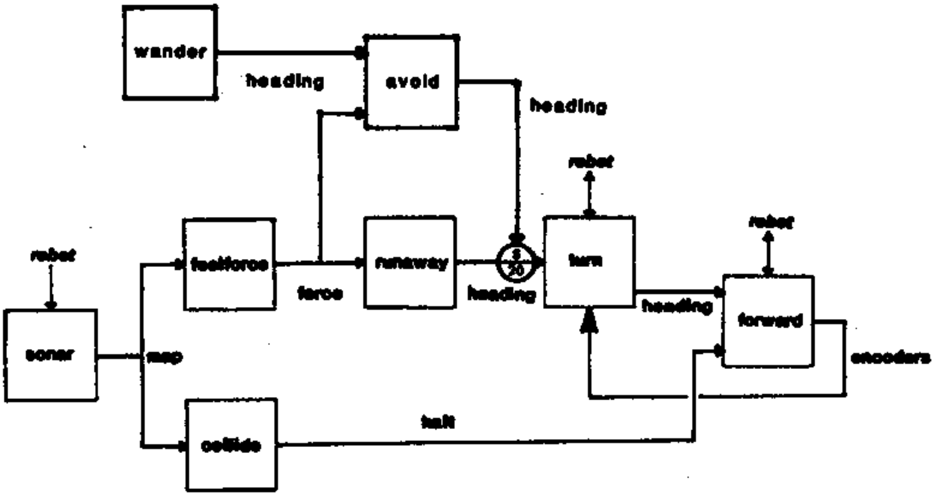

Figure 6.The level *0* control system augmented with the level *1* system.

First Level

The first level layer of control, when combined with the zeroth, imbues the robot with the ability to wander around without hitting obstacles. This was defined earlier as level *1* competence. This control level relies in a large degree on the zeroth level's aversion to hitting obstacles. In addition it uses a simple heuristic to plan ahead a little in order to avoid potential collisions which would need to be handled by the zeroth level.

Figure 6 gives a complete description of how the modules are connected together. Note that it is simply Figure 5 with some more modules and wires added.

Second Level

Level two is meant to add an exploratory mode of behavior to the robot, using visual observations to select interesting places to visit. A vision module finds corridors of free space. Additional modules provide a means of position servoing the robot to along the corridor despite the presence of local obstacles on its path (as detected with the sonar sensing system). The wiring diagram is shown in figure 7. Note that it is simply figure 6 with some more modules and wires added. The zeroth and first layers still play an active role during normal operation of the second layer. (In practice we have so far only reused the sonar data for the corridor finder, rather than use stereo vision.)

5. Performance

The control system described here has been used extensively to control both a simulated robot and an actual physical robot wandering around a cluttered laboratory and a machine

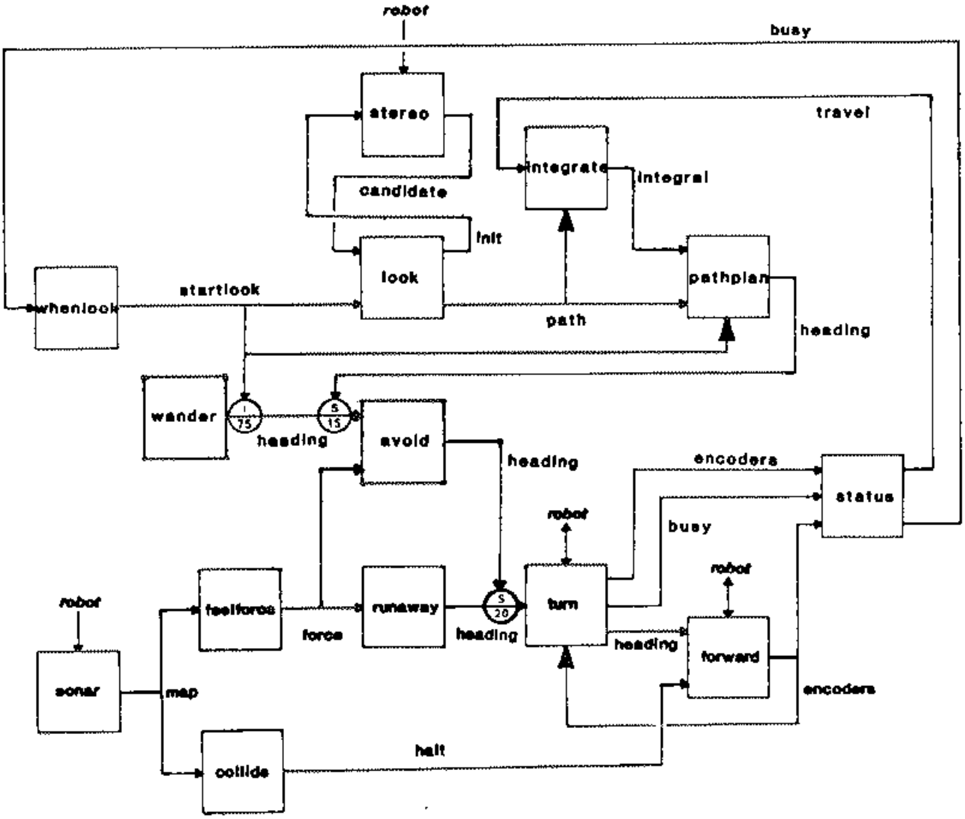

Figure 7. The level *0* and *1* control systems augmented with the level *2* system.

room. We have constructed a mobile robot shown in Figure 8.

The robot has spent a few hours wandering around a laboratory and a machine room.

Under level *0* control the robot finds a large empty space and then sits there contented until a moving obstacle approaches. Two people together can successfully herd the robot just about anywhere-through doors or between rows of disk drives, for instance.

When level *1* control is added the robot is no longer content to sit in an open space. After a few seconds it heads off in a random direction. Our uncalibrated sonars and obstacle repulsion functions make it overshoot a little to locations where the **runaway** module reacts. It would be interesting to make this the basis of adaption of certain parameters.

Under level *2* a sonar-based corridor finder usually finds the most distant point in the room. The robot heads off in that direction. People walking in front of the robot cause it to detour, but still get to the initially desired goal even when it involves squeezing between closely spaced obstacles.

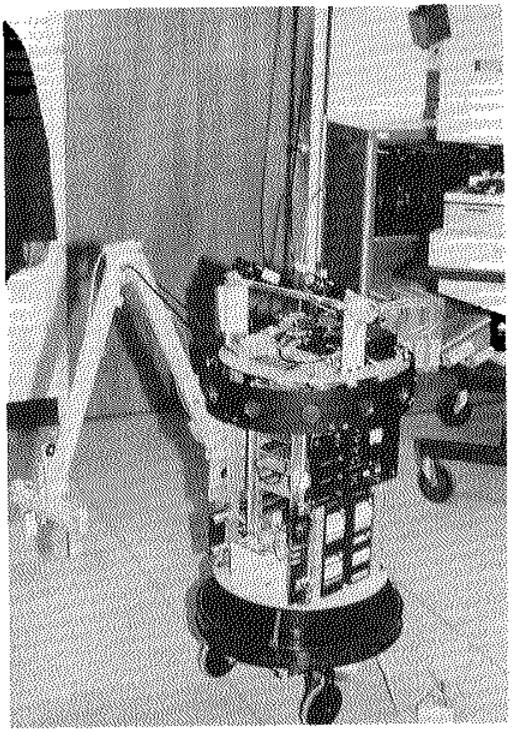

Figure 8. The MIT AI Lab mobile robot.

6. Key Ideas

The key ideas in this paper are:

- The mobile robot control problem can be decomposed in terms of behaviors rather than in terms of functional modules.
- It provides a way to incrementally build and test a complex mobile robot control system.
- Useful parallel computation can be performed on a low bandwidth loosely coupled network of asynchronous simple processors. The topology of that network is relatively fixed.
- There is no need for a central control module of a mobile robot. The control system can be viewed as a system of agents each busy with their own solipsist world.

Besides leading to a different implementation strategy it is also interesting to note the way the decomposition affected the capabilities of the robot control system we have built. In particular our control system deals with moving objects in the environment at the very lowest level, and has a specific module (**runaway**) for that purpose. Traditionally mobile robot projects have delayed handling moving objects in the environment beyond the scientific life of the project.

7. Applications

Current day underwater mobile robots are teleoperated, requiring continual attention from a human operator in a support vessel. This produces economic pressure that (1) the robot is only put in the water to do a task when it is necessary that it be done, and (2) that the robot do the task quickly. Such pressures shape the way in which the robot is designed for a particular application.

Consider the problem of cleaning off encrustations from an underwater structure in order to inspect the welds at the joints. Under current practice this is only done occasionally, when the joints need to be inspected, and so the encrustations are quite thick and hard to remove. Suppose we could dedicate a person, or a robot, to cleaning off the structure *every day*. The encrustations would be almost non-existent and so the task would be much simpler. Inspections could be scheduled with much less consideration of the cost of each and could lead to earlier detection of structural failures.

The expense of dedicating a teleoperated robot to the task would be prohibitive. However if a completely autonomous robot could live on the structure for months at a time with no human intervention perhaps it would be a practical approach. There are some technical difficulties remaining of course; the robot needs a supply of energy, and it needs to be mechanically reliable. This paper, however, has presented some ideas which lay the foundations for solving one of the major subproblems of building such a robot: how to build a reflexive robust autonomous control system which enables the robot to carry out some useful task over a long period of time.

Acknowledgments

Support for this work was provided in part by an IBM Faculty Development Award, in part by a grant from the Systems Development Foundation, and in part by the Advanced Research Projects Agency under Office of Naval Research contracts N00014-80-C-0505 and N00014-82-K-0334.

References

[Brooks 84] "Aspects of Mobile Robot Visual Map Making", Rodney A. Brooks, *Robotics Research 2*, Hanafusa and Inoue eds, MIT Press, 1984, 369–375.

[Brooks 85] "Visual Map Making for a Mobile Robot", Rodney A. Brooks, *IEEE Conference on Robotics and Automation*, St Louis, March, 1985, 824–829.

[Crowley 85] "Navigation for an Intelligent Mobile Robot", James L. Crowley, *IEEE Journal of Robotics and Automation*, RA-1, March 1985, 31–41.

[Elfes and Talukdar 83] "A Distributed Control System for the CMU Rover", Alberto Elfes and Sarosh N. Talukdar, *Proceedings IJCAI*, Karlsruhe, West Germany, August 1983, 830–833.

[Flynn 85] "Redundant Sensors for Mobile Robot Navigation", Anita Flynn, *MS Thesis, Department of Electrical Engineering and Computer Science*, MIT, July 1985.

[Giralt et al 83] "An Integrated Navigation and Motion Control System for Autonomous Multisensory Mobile Robots", Georges Giralt, Raja Chatila, and Marc Vaisset, *Robotics Research 1*, Brady and Paul eds, MIT Press, 1983, 191–214.

[Grimson 85] "Computational Experiments with a Feature Based Stereo Algorithm", W. Eric L. Grimson, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-7, January 1985, 17–34.

[Kanayama 83] "Concurrent Programming of Intelligent Robots", Yutaka Kanayama, *Proceedings IJCAI*, Karlsruhe, West Germany, August 1983, 834–838.

[Khatib 83] "Dynamic Control of Manipulators in Operational Space", Oussama Khatib, *Sixth IFTOMM Congress on Theory of Machines and Mechanisms*, New Delhi, December 1983.

[Kreithen 83] "Orientational Strategies in Birds: A Tribute to W. T. Keeton", Melvin L. Kreithen, in *Behavioral Energetics: The Cost of Survival in Vertebrates*, Ohio State University Press, 1983, 3–28.

[Moravec 83] "The Stanford Cart and the CMU Rover", Hans P. Moravec, *Proceedings of the IEEE*, 71, July 1983, 872–884.

[Nilsson 84] "Shakey the Robot", Nils J. Nilsson, *SRI AI Center Technical Note 323*, April 1984.

[Simon 69] "Sciences of the Artificial", Herbert A. Simon, MIT Press, 1969.

[Tsuji 84] "Monitoring of a Building Environment by a Mobile Robot", Saburo Tsuji, *Robotics Research 2*, Hanafusa and Inoue eds, MIT Press, 1985, 349–356.

Recognition and Localization from Noisy Sensor Data

W. Eric L. Grimson
Assistant Professor
Artificial Intelligence Laboratory
Massachusetts Institute of Technology

1. Introduction

The problem of object recognition and localization is a pervasive one in machine vision. In simplest terms, the problem is to identify a known object, by matching a set of data elements to corresponding model elements. The problem is made interesting by the fact that, in general, much of the sensor data is spurious, coming from objects other than the object of interest; by the fact that the object is frequently occluded, so that not all of it is visible to the sensor; and by the fact that the data is usually inaccurate. Once the object has been identified, by matching sensory data to model fragments, one can determine the location of the object, by solving for the six degree of freedom transformation that transforms the sensor data into the corresponding model fragments. Because of the simplicity of the approach we take to this problem, the same technique can be applied in other areas, in particular the localization of a sensor relative to a known map of its environment. For example, given a rough, possible incomplete, map of the environment, and given noisy sonar data about the current environment, a mobile robot can use the same method to determine its location relative to the map.

In this paper, we briefly sketch an approach to object recognition that has been successfully applied to a wide range of visual and laser data. We then suggest how the same method can be applied to the problem of mobile robot localization. We support this suggestion by presenting the results of a preliminary study conducted by Michael Drumheller [Drumheller 1984].

2. The problem of object recognition

Recently, we have developed a particular approach [Grimson and Lozano-Pérez 1984, 1985] to the problem of how to locate a known object from sensory data. The object may be occluded by other unknown objects, so that much of the sensory data does not arise from the object of interest. Examples of the performance of the method are shown in Figures 1 – 2.

2.1. The data and the model

We are interested in a recognition method that is applicable to a wide range of sensor types, therefore, we make very few assumptions about the character of the sense data. We assume only that the sensory data can be processed to obtain the position and surface orientation of planar patches on the object. The measured positions are assumed to be within a known error volume and the measured surface orientations to be within a known error cone.

When the objects have only three degrees of positional freedom relative to the sensor (two translational and one rotational), the positions and surface normals need only be two–dimensional. When the objects have more than three degrees of positional freedom (up to three translational and three rotational), the position and orientation data must be three–dimensional.

We assume that the objects can be modeled as sets of planar faces. Only the individual plane equations and a polygon embedded in each face is required. The model faces do not have to be connected and the model does not have to be complete.

2.2. Our approach to localisation

We approach the localization problem as a search for a consistent matching between the measured surface patches and the surfaces of the known object model. The search proceeds in two steps:

1. *Generate Feasible Interpretations:* Interpretations consist of pairings of sensed patches with some surface on the object model. Interpretations in which the sensed data are inconsistent with local geometric constraints derived from the model are discarded.
2. *Model Test:* The feasible interpretations are tested for consistency with surface equations obtained from the object models. An interpretation is legal if it is possible to solve for a rotation and translation that would place each sensed patch on an object surface. The sensed patch must lie *inside* the object face, not just on the surface defined by the face equation.

We structure the search for consistent matches as the generation and exploration of an *interpretation tree* (IT) (see Figure 3). That is, starting at a root node, we construct a tree in a depth first fashion, assigning measured patches to model faces. At the first level of the tree, we consider assigning the first measured patch to all possible faces, at the next level, we assign the second measured patch to all possible faces, and so on. The number of possible interpretations in this tree, given s sensed patches and n surfaces, is

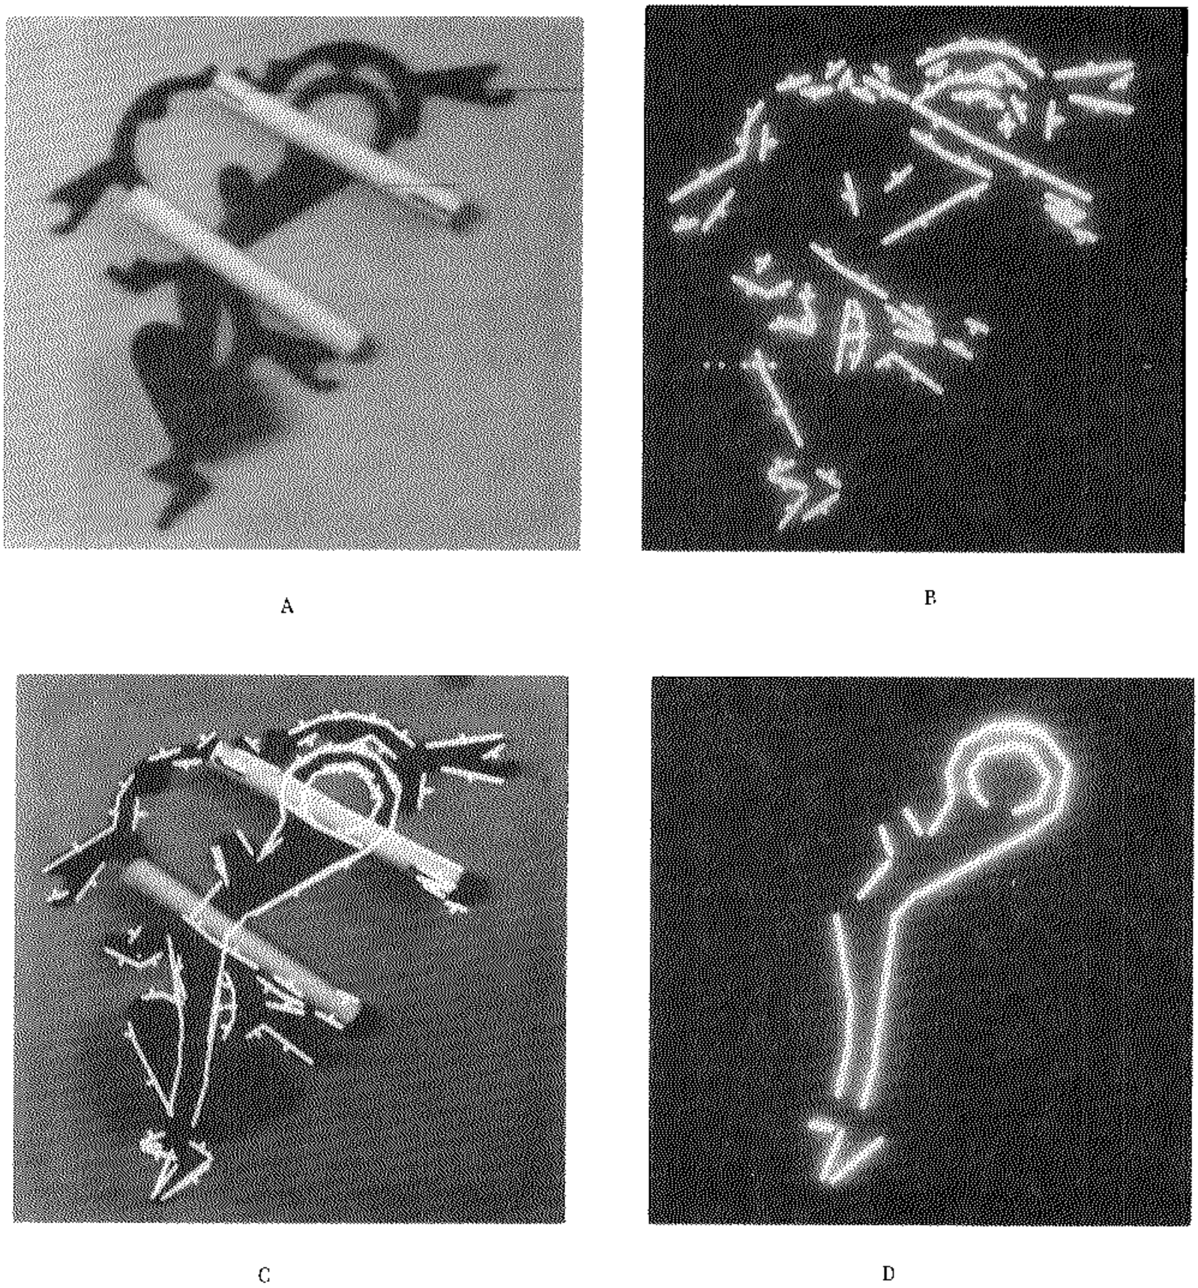

Figure 1. Two dimensional edge data. (a) Grey level images, (b) edge fragments, (c) located object in image, and (d) located object.

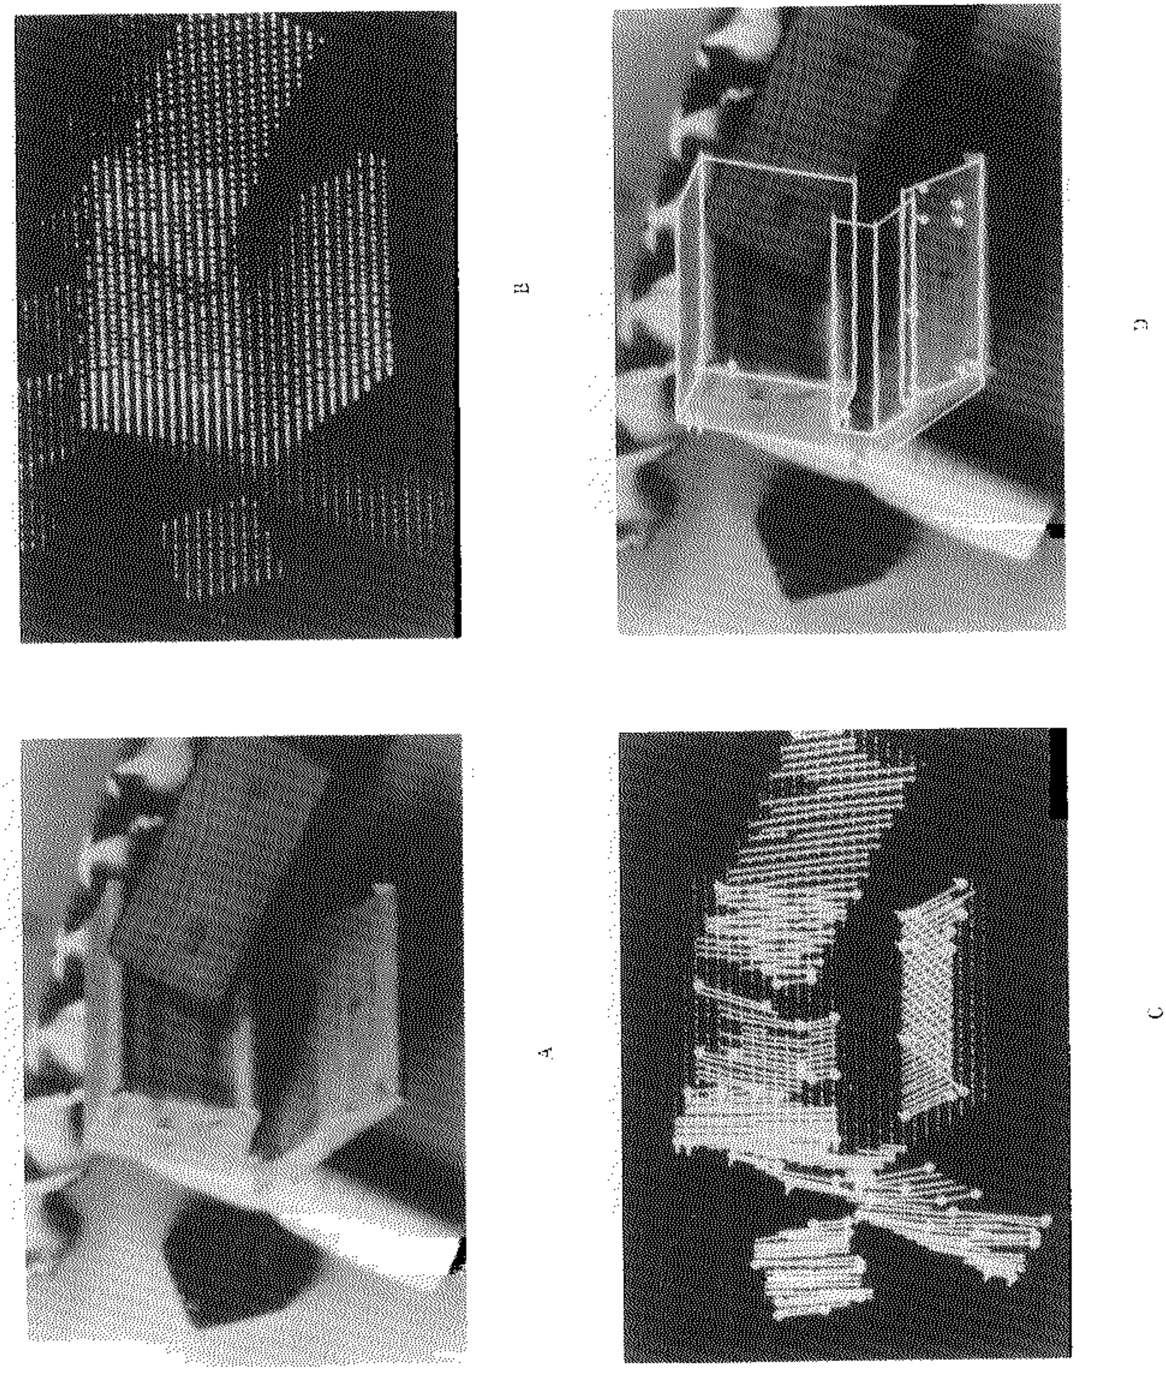

Figure 2. Three dimensional range data. (a) Original scene (b) range data where brightness encodes height, (c) planar patches with four representative points per patch, and (d) located object superimposed on original scene (the data points accounted for are indicated).

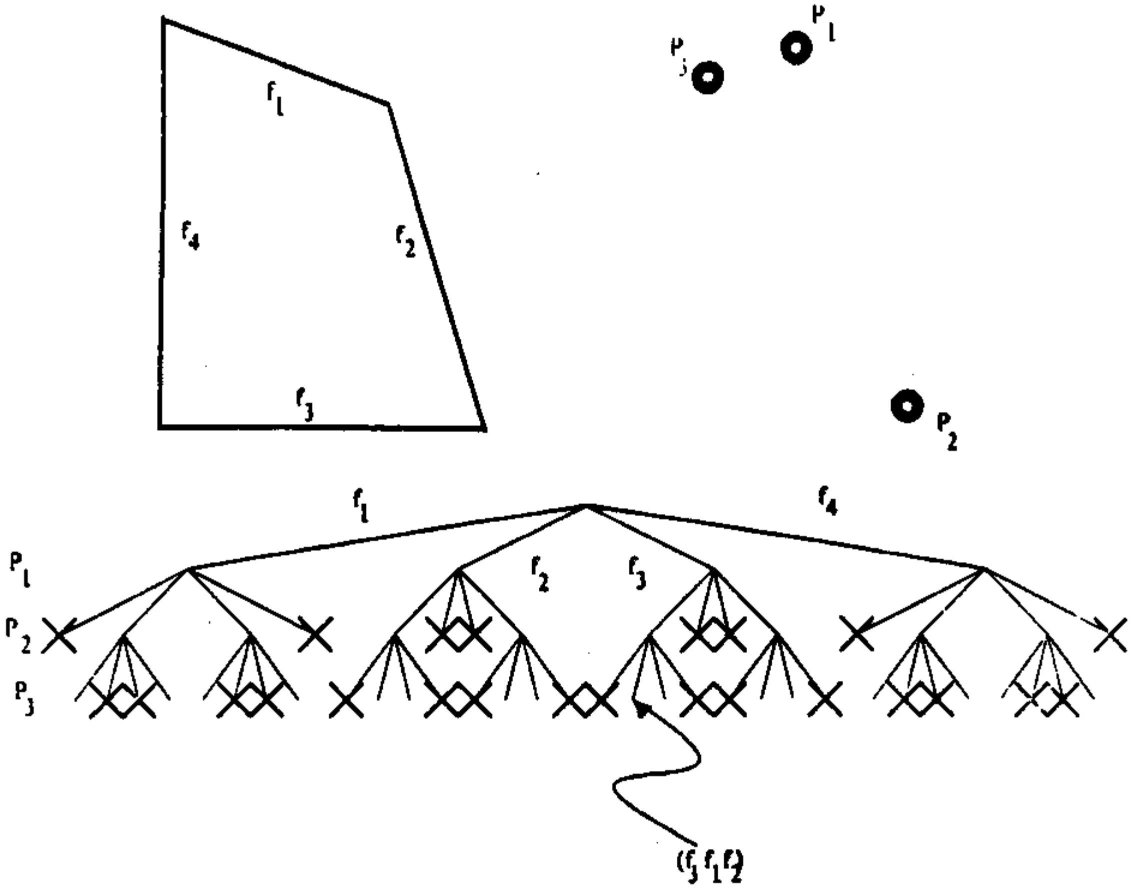

Figure 3. A simple example of constrained search. We want to find consistent matchings of the three data points to the edges of the indicated quadrilateral. If we use distance between data points as the measurement, then the table of possible ranges between the edges of the object is given by

	1	2	3	4
1	[0, 1.5]	[0, 3.25]	[2, 3.25]	[0, 2.5]
2	[0, 3.25]	[0, 2]	[0, 2.5]	[1.4, 3.25]
3	[2, 3.25]	[0, 2.5]	[0, 2]	[0, 3.25]
4	[0, 2.5]	[1.4, 3.25]	[0, 3.25]	[0, 2.5]

The tree indicates the set of possible assignments of data points to object edges, give distance as a constraint. One can see that only 16 out of 64 possible interpretations are consistent with this constraint.

n^s. Therefore, it is not feasible to explore the entire search space in order to apply a model test to all interpretations.

Our algorithm exploits local geometric constraints to remove entire subtrees from consideration. In our case, we require that the distances and angles between all pairs of data elements be consistent with the distances and angles possible between their assigned model elements. In general, the constraints must be coordinate–frame independent, that is, the constraints should embody restrictions due to object shape and not to sensing geometry.

2.3. Relation to previous work

The literature on object recognition stretches over a period of twenty years. An extensive (70 page) review of much of this literature can be found in [Besl and Jain 1985]. In terms of the approach to be described here, a number of authors have taken a similar view to ours that recognition can be structured as an explicit search for a match between data elements and model elements [Ayache and Faugeras 86, Baird 85, Bolles and Cain 82, Bolles, Horaud and Hannah 83, Faugeras and Hebert 83, Goad 83, Lowe 86, Stockman and Esteva 84]. Of these, the work of Bolles and his colleagues, Faugeras and his colleagues, and that of Baird are closest to the approach presented here.

The interpretation tree approach is an instance of the consistent labeling problem that has been studied extensively in computer vision and artificial intelligence [Waltz 75, Montanari 74, Mackworth 77, Freuder 78, 82, Haralick and Shapiro 79, Haralick and Elliott 80, Mackworth and Freuder 85]. This paper can be viewed as suggesting a particular consistency relation (the constraints on distances and angles) and exploring its performance in a wide variety of circumstances. An alternative approach to the solution of consistent labeling problems is the use of relaxation. A number of authors have investigated this approach to object recognition [Davis 79, Bhanu and Faugeras 84, Ayache and Faugeras 82]. These techniques are more suitable for implementation on parallel machines.

3. The Constraints

The key to making the search process practical is in finding constraints on the measurements that can quickly reduce the combinatorial explosion inherent in depth first search. Below, we briefly review one such set of constraints, a more detailed exposition can be found in [Grimson and Lozano-Pérez 84, 85].

3.1. The Decoupled Constraints

First construct a local coordinate frame relative to the sensed data; we use both unit normals as basis vectors. In two dimensions, these define a local system, except in the degenerate case of the unit normals being (anti–)parallel. In three dimensions, the third component of the local coordinate frame can be taken as the unit vector in the direction of the cross product of the normal vectors. In this frame, one set of coordinate–frame–

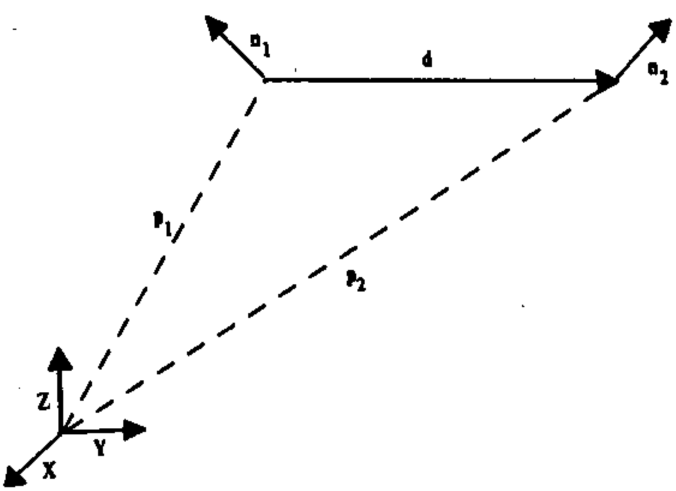

Figure 4. The constraints between pairs of measured surface patches. A given pair of sensory points P_1, P_2 can be characterised by the components of the vector $\mathbf{d}$ between them, in the direction of each of the surface normals $\mathbf{n}_1, \mathbf{n}_2$ and in the direction of their cross product, $\mathbf{n}_1 \times \mathbf{n}_2$, and by the angle between the two normals $\mathbf{n}_1 \cdot \mathbf{n}_2$.

independent measurements is: the components of the vector $\mathbf{d}$ along each of the basis directions and the angle between the two measured normals (see Figure 4). More formally,

$$\begin{gathered} \mathbf{n}_1 \cdot \mathbf{n}_2 \\ \mathbf{d} \cdot \mathbf{n}_1 \\ \mathbf{d} \cdot \mathbf{n}_2 \\ \mathbf{d} \cdot \mathbf{u} \end{gathered}$$

where $\mathbf{u}$ is a unit vector in the direction of $\mathbf{n}_1 \times \mathbf{n}_2$.

These measurements are equivalent, but not identical to the set used in [Grimson and Lozano-Pérez 84]. In the earlier paper, we used the magnitude of $\mathbf{d}$ and two of its components; this is equivalent, up to a possible sign ambiguity, to using the three components of the vector. This possible ambiguity in the earlier set of measurements was resolved using a triple product constraint.

For these measurements to constrain the search process, we must relate them to corresponding model measurements. Consider the first measurement, $\mathbf{n}_1 \cdot \mathbf{n}_2$. If this is to correspond to a measurement between two faces in the model, then the dot product of the model normals must agree with this measurement. If they do not agree, then no interpretation that assigns those patches to these model faces need be considered. In the interpretation tree, this corresponds to pruning the entire subtree below the node corresponding to that assignment. The test can be implemented efficiently by precomputing the dot product between all pairs of faces in the models. Of course, for the case of exact measurements, the dot product of the measured normals must be identical to that of the associated model normals. In practice, exact measurements are not possible, and we must take possible sensor errors into account. Given bounds on the error in a sensory measurement, we compute a range of possible values associated with the dot product of two sensed normals (see [Grimson and Lozano–Pérez 84] for details).

Similar constraints can be derived for the components of the separation vector in the directions of the unit normals. Each pair of model faces defines an infinite set of possible separation vectors, each one having its head on one face and its tail in the other. We can compute bounds on the components of this set of vectors in the direction of each of the face normals. Again, for an assignment of sensed patches to model faces to be consistent, the measured value must agree with the precomputed model values. Here also we can use the error bounds to compute a range of possible values for the components of the sensed vectors; this range must be consistent with the associated model range.

While these constraints are not guaranteed to reject all impossible interpretations, they are extremely efficient. In [Grimson and Lozano-Pérez 85], we show that more tightly coupled constraints are also possible.

Because the constraints are decoupled, the fact that all pairs of patch–surface assignments are consistent does not imply that the global assignment is consistent. To determine global consistency, we solve for a transformation from model coordinates to sensor coordinates that maps each of the sensed patches to the interior of the appropriate face. There are many methods to solve for the transformation, one is described in [Grimson and Lozano–Pérez 84] another can be found in [Faugeras and Hebert 83]. This model test is applied to interpretations surviving pruning so as to guarantee that all the available geometric constraint is satisfied. As a side effect, the model test also provides a solution to the localization problem.

3.2 Handling spurious data

While simple geometric constraints such as those given above are effective in eliminating large portions of the IT when all the data originates from a single object, when the data originates from several objects, we need to modify the approach slightly.

Our approach to handling extraneous data from unknown objects is to add one more branch to each node of the interpretation tree, IT. This branch represents the possibility of discarding the sensed patch as extraneous. Call this branch the *null face*. The search proceeds, as before, to explore the IT depth first. As each new assignment of a data patch to a model face is considered, the new interpretation thus formed is tested to see whether it satisfies the geometric constraints. In these tests, the null face behaves as a "wild card"; assigning a patch to the null face will never cause the failure of an interpretation.

Clearly, if an interpretation is legal, all subsets of this interpretation are leaves of the expanded IT. This is true since every combination of legal assignments of the null face to the sensed patches will still produce a valid interpretation. Rather than generating all of these subsets, we want to generate the "best" interpretation. The problem then arises of choosing the quality measure. [Faugeras and Hebert 83] have explored the use of a measure based on how well the computed model transformation maps measured patches into model faces. We have chosen instead to search for interpretations where the data patches have the largest combined area. The reason for our choice is that this measure is simple and fairly insensitive to measurement error. The following simple search method guarantees that we find only the most complete interpretations.

The IT is explored in a depth-first fashion, with the null face considered last when expanding a node. Now, assume an external variable, call it MAX, that keeps track of the best (largest area) valid interpretation found so far. For a node at level i in the tree, let M denote the area of the data patches assigned to non-null faces in the partial match associated with that node. Let R be the area of the data patches below this level of the tree: $P_{i+1}, \ldots, P_n$. It is only worth assigning a null face to patch P_i, if $M + R \geq MAX$. Otherwise, the area of the interpretations at all the leaves below this node will be less than that of the best interpretation already found. If we initialize MAX to some non-zero value, then only interpretations with area greater than this threshold will be found. As better interpretations are found, the value of MAX is incremented, thus ensuring that we find the most complete interpretation of the data. Note that if an interpretation of maximal area (no null-face assignments) is found, then no further null-face assignments will be considered after that point.

The search process described above can be continued until all the nodes have either been examined or discarded. This can take a very long time for realistic cases. We observed that the search located the correct interpretation fairly early on, but then spent a tremendous amount of time attempting to improve on it. This phenomenon can be avoided by the use of an area threshold (as a percentage of the model's area). The search is discontinued when an interpretation that exceeds that threshold passes the model test. We have found that this search cutoff drastically improves the execution time without adversely affecting the failure rate.

4. Simulations and Tests with Live Data

We have tested several variations of the algorithm with a wide variety of simulated data and live data. These tests are reported elsewhere [Grimson and Lozano-Pérez 84, 85]. The tests demonstrate that the method can reliably locate occluded objects from noisy data, where the data source can be grey-level visual data, laser ranging data, sonar ranging data or tactile sensing data. Examples are shown in Figures 1 – 2.

5. Using the Recognition Method to Determine a Mobile Robot's Position

Michael Drumheller [Drumheller 84] has developed a modified version of our algorithm and applied it to range data obtained from an unmodified Polaroid ultrasonic range sensor. The intended application is navigation of mobile robots. In this section, we briefly describe how our parts recognition system can be modified to handle this problem.

In the system developed by Drumheller, it was assumed that the robot's environment was a room or area inside a building. Further, it was assumed that a geometric model of the rooms was provided to the system. The rangefinding device used in his experiments was a Polaroid Ultrasonic Rangefinder, which was swept through a 360-degree path.

Given such a model of the building and data from the sonar, the matching algorithm proceeds in a manner similar to that described above. Drumheller uses four stages:

- Extract straight line segments from the sonar data.
- Generate feasible interpretations. This proceeds identically to the object recognition case, using constraints to prune a depth first tree search. Interpretations consist of pairings of data segments to model edges.
- Global model test. For each feasible interpretation, we solve for a consistent coordinate frame transformation that takes each data segment into the corresponding model segment.
- Sonar barrier test. An interpretation may represent a geometrically feasible configuration for the sonar segments alone, but an impossible configuration for the entire sonar contour. In particular, each interpretation that survives the model test must also pass the *sonar barrier test,* namely: an admissible robot configuration must not imply that any sonar ray penetrates a known solid object.

The end result of the algorithm is a determination of the position and orientation of the robot with respect to the world map.

Drumheller tested his system of a series of live experiments. In these trials, the robot has a map of the walls of the room, but much of the data obtained arises from objects on the walls, such as bookshelves, or between the robot and the walls, such as columns. Drumheller assumed an upper bound of 1.3 feet in position estimates and 10 degrees in orientation estimates for his sensor. His system was able to determine the position and orientation of the mobile robot with an average error of less than 6 inches and 1 degree. Examples are shown in Figures 5 and 6.

6. Possible Extensions

The system used by Drumheller demonstrates the feasibility of using object recognition systems to locate mobile robots relative to a known environment. Several comments are in order, however, concerning the limitations of the demonstration, and the possibility of using such a system in other domains, such as underwater.

First, it should be noted that the sonar data used in his experiments are so noisy and of such low resolution as to nearly constitute a "worst case scenario" for range data. It is likely that the performance would be significantly enhanced by higher-resolution data.

Second, the examples presented by Drumheller are all inherently two-dimensional in nature, that is, the robot has only three degrees of positional freedom, two translational and one rotational. This is fine for a land based mobile robot, which is assumed to have a stable platform. What happens when we consider the problem of an underwater robot with up to six degrees of positional freedom? First, we know from our work on recognizing

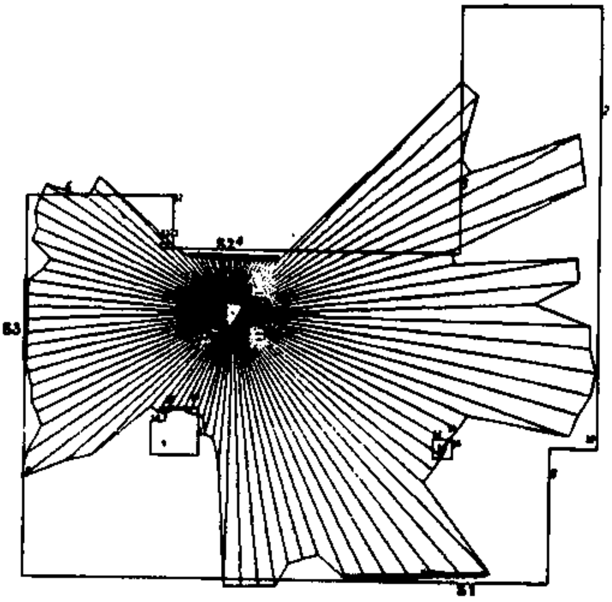

```
[(S1 W1)(S2 W4)(S3 W3)];  x = -28.9 ft.  y = 23.9 ft.  168 deg.
Actual configuration was:  x = -29.3 ft.  y = 24.6 ft.  170 deg.
```

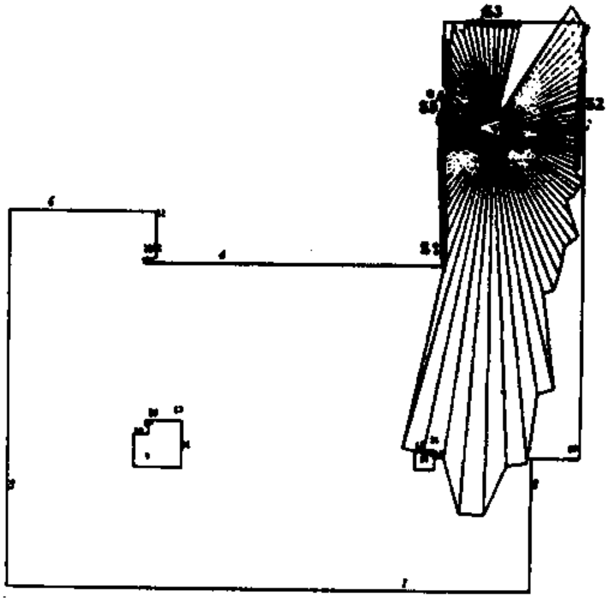

```
[(S1 W5)(S2 W2)(S3 W7)(S4 W5)(S5 W6)];  x = -3.9 ft.  y = 42.2 ft.  89 deg.
               Actual configuration was:  x = -3.5 ft.  y = 42.1 ft.  90 deg.
```

Figure 5. Examples of localization of a mobile robot. The sonar scan is superimposed on the model of the room. The printed data list the computed position and the actual position. The figure is reprinted from [Drumheller 1984].

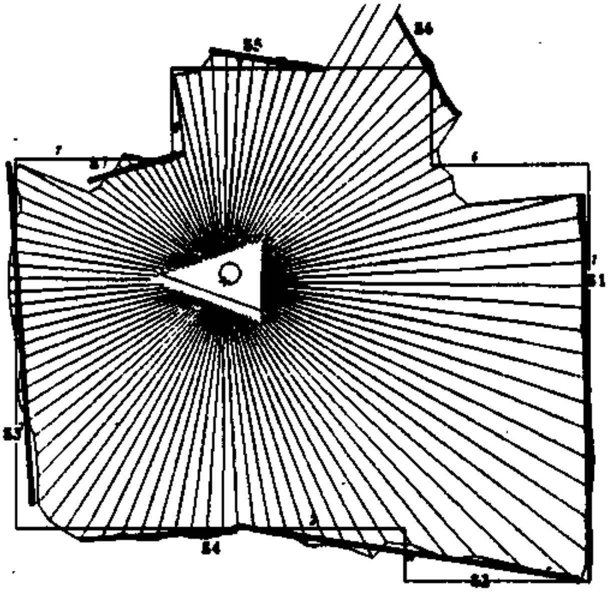

```
[(S1 W1)(S2 W*)(S3 W3)(S4 W2)(S5 W4)(S6 W*)(S7 W*)];  x = -7.4 ft.  y = 6.8 ft.  89 deg.
                        Actual configuration was:  x = -7.2 ft.  y = 6.2 ft.  90 deg.
```

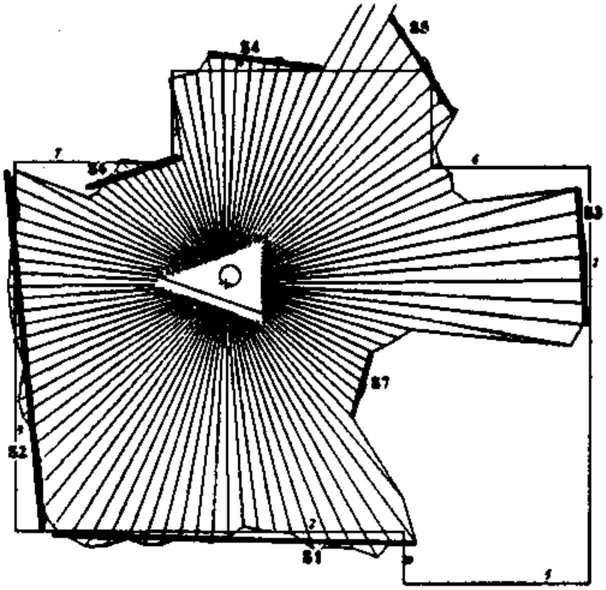

```
[(S1 W2)(S2 W3)(S3 W1)(S4 W4)(S5 W*)(S6 W*)(S7 W*)];  x = -7.4 ft.  y = 6.8 ft.  91 deg.
                        Actual configuration was:  x = -7.2 ft.  y = 6.2 ft.  90 deg.
```

Figure 6. Examples of localization of a mobile robot. The sonar scan is superimposed on the model of the room. The printed data list the computed position and the actual position. The figure is reprinted from [Drumheller 1984].

objects with six degrees of freedom that our method extends in a straightforward manner. If we have three dimensional models of the environment available to the system, the same type of robot localization should work, since the sonar data returned is also three dimensional.

There may be other solutions, however. One example would be to use bathymetric data. There are several possibilities. For example, suppose the underwater vehicle can be placed in a stable orientation, with the plane of the sonar scan perpendicular to the gravity vector, but without knowing the actual depth of the vehicle. In this case, one can simply match the sonar return against each of a set of environment models, one for each different depth contour contained in the bathymetric data. If the vehicle can also tell its depth, then the problem can be simplified to only consider the corresponding set of depth contours. If the orientation of the platform cannot be controlled, then a full three-dimensional model of the environment can be constructed, and the sonar data can be matched against the surface patches of this model.

Third, it should be noted that the techniques described here can be used not only to localize an underwater mobile robot, but also to locate nearby objects. For situations in which a robot arm is used in conjunction with the platform, it is useful to have the ability to automatically position the arm in order to pick up and manipulate objects in the environment.

9. Summary

In this paper, we have presented a brief review of a method for recognizing and locating objects from noisy, occluded data. We have suggested that this method is also applicable to the problem of positioning a mobile robot platform relative to a known map of its environment. Early experiments with such a system on real data were reported.

Acknowledgments

The image processing shown in Figure 1 was done on a hardware/software environment developed by Keith Nishihara and Noble Larson. We thank Philippe Brou for kindly providing the laser ranging system with which we obtained the data reported in Figure 2. We thank Mike Drumheller for permission to reproduce Figures5 and 6.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the Laboratory's Artificial Intelligence research is provided in part by a grant from the System Development Foundation, in part by the Office of Naval Research under contract N00014-82-K-0494, and in part by the Advanced Research Projects Agency under Office of Naval Research contracts N00014-80-C-0505 and N00014-82-K-0334.

References

Ayache, N. J. and Faugeras, O. D. 1982. Recognition of partially visible planar shapes. *Proc. 6th Intl. Conf. Pattern Recognition*, Munich.

Ayache, N. J. and Faugeras, O. D. 1986. HYPER: A new approach for the recognition and positioning of two-dimensional objects. *IEEE Trans. Pattern Anal. Machine Intell.* PAMI-8(1):44–54.

Baird, H. 1986. *Model-Based Image Matching Using Location*. Cambridge: MIT Press.

Besl, P. J. and Jain, R. C. 1985. Three-Dimensional Object Recognition. *ACM Computing Surveys* 17(1): 75–145.

Bhanu, B. and Faugeras, O. D. 1984. Shape matching of two-dimensional objects. *IEEE Trans. Pattern Anal. Machine Intell.* PAMI-6(3).

Bolles, R. C., and Cain, R. A. 1982. Recognizing and locating partially visible objects: The Local-Feature-Focus method. *Int. J. Robotics Res.* 1(3):57–82.

Bolles, R. C., Horaud, P., and Hannah, M. J. 1983. 3DPO: A three-dimensional part orientation system. Paper delivered at First International Symposium of Robotics Research, Bretton Woods, N.H. (Also in Robotics Research: The First International Symposium, edited by M. Brady and R. Paul, MIT Press, 1984, pp. 413–424.)

Davis, L. 1979. Shape matching using relaxation techniques. *IEEE Trans. Pattern Anal. Machine Intell.* PAMI-1(1):60–72.

Drumheller, M. 1984. Robot Localization Using Range Data. S. B. Thesis, Dept. of Mechanical Engineering, MIT. (see also MIT AI Lab Memo 826, Mobile Robot Localization Using Sonar.)

Faugeras, O. D. and Hebert, M. 1983 (Aug., Karlsruhe, W. Germany). A 3-D recognition and positioning algorithm using geometrical matching between primitive surfaces. *Proc. Eighth Int. Joint Conf. Artificial Intell.* Los Altos: William Kaufmann, pp. 996–1002.

Freuder, E. C. 1978. Synthesizing constraint expressions. *Comm. of the ACM*, 21(11), pp. 958–966.

Freuder, E. C. 1982. A sufficient condition for backtrack-free search. *J. ACM*, 29(1), pp. 24–32.

Goad, C. 1983. Special purpose automatic programming for 3d model-based vision. in *Proceedings of DARPA Image Understanding Workshop*.

Grimson, W. E. L., and Lozano-Pérez, T. 1984. Model-based recognition and localization from sparse range or tactile data. *Int. J. Robotics Res.* 3(3):3–35.

Grimson, W. E. L., and Lozano-Pérez, T. 1985. Recognition and localization of overlapping parts from sparse data. *MIT Artificial Intelligence Laboratory Memo 841* June, 1985.

Haralick, R. M. and Elliott, G. 1980. Increasing tree search efficiency for constraint satisfaction problems. *Artificial Intelligence*, Vol. 14, pp. 263–313.

Haralick, R. M. and Shapiro, L. G. 1979. *IEEE Trans. Pattern Anal. Machine Intell.* PAMI-1(4):173–184.

Lowe, D. G. 1986. Three-dimensional object recognition from single two-dimensional images. Courant Institute Robotics Report, No. 62.

Mackworth, A. K. 1977. Consistency in networks of constraints. *Artificial Intelligence*, Vol. 8, pp. 99–118.

Mackworth, A. K. and Freuder, E. C. 1985. The complexity of some polynomial network consistency algorithms for constraint satisfaction problems. *Artificial Intelligence*, Vol. 25, pp. 65–74.

Montanari, U. 1974. Networks of constraints: Fundamental properties and applications to picture processing. *Inform. Sci.*, Vol. 7, pp 95–132.

Stockman, G., and Esteva, J. C. 1984. Use of geometrical constraints and clustering to determine 3D object pose. TR84-002. East Lansing, Mich.:Michigan State University Department of Computer Science.

Waltz, D. 1975. Understanding line drawings of scenes with shadows. in *The Psychology of Computer Vision*, P. Winston, Ed. New York:Mc Graw Hill, pp 19 – 91.

Mobile Robots and General Intelligence

Hans P. Moravec
Senior Research Scientist
The Robotics Institute
Carnegie-Mellon University

Introduction

The significance of mobile robot research may be much greater than the sum of its applications. There is a parallel between the evolution of intelligent living organisms and the development of robots. Many of the real-world constraints that shaped life by favoring one kind of change over another in the contest for survival also affect the viability of robot characteristics. To a large extent the incremental paths of development pioneered by living things are being followed by their technological imitators. Given this, there are lessons to be learned from the diversity of life. One is that mobile organisms tend to evolve in the direction of general intelligence, immobile ones do not. The plants are an example of the latter case, vertebrates an example of the former. An especially dramatic contrast is provided in an invertebrate phylum, the molluscs. Most are shellfish like clams and oysters that move little and have tiny nervous systems and behaviors more like plants than like animals. Yet they have relatives, the cephalopods, like octopus and squid, that are mobile and have independently developed many of the characteristics of vertebrates, including imaging eyes, large nervous systems and very interesting behaviour, including major problem solving abilities.

Mobility and Intelligence in Nature

Two billion years ago our unicelled ancestors parted genetic company with the plants. By accident of energetics and heritage, large plants now live their lives fixed in place. Awesomely effective in their own right, the plants have no apparent inclinations towards intelligence; a piece of negative evidence that supports my thesis that mobility is a parent of this trait.

Animals bolster the argument on the positive side, except for the immobile minority like sponges and clams that support it on the negative.

A billion years ago, before brains or eyes were invented, when the most complicated animals were something like hydras, double layers of cells with a primitive nerve net, our progenitors split with the invertebrates. Now both clans have intelligent members. Cephalopods are the most intellectual invertebrates. Most mollusks are sessile shellfish, but octopus and squid are highly mobile, with big brains and excellent eyes. Evolved independently of us, they are different. The optic nerve connects to the back of the retina, so there is no blind spot. The brain is annular, a ring around the esophagus. The green blood is circulated by a systemic heart oxygenating the tissues and two gill hearts moving depleted blood. Hemocyanin, a copper doped protein related to hemoglobin and chlorophyll, carries the oxygen.

Octopus and their relatives are swimming light shows, their surfaces covered by a million individually controlled color changing cells. A cuttlefish placed on a checkerboard can imitate the pattern, a fleeing octopus can make deceiving seaweed shapes coruscate backward along its body. Photophores of deep sea squid, some with irises and lenses, generate bright multicolored light. Since they also have good vision, there is a potential for rich communication. Martin Moynihan identifies several dozen distinct symbolic displays, many correlated with clear emotions in **Communication and Noncommunication by Cephalopods.**

Their behavior is mammal like. Octopus are reclusive and shy, squid are occasionally very aggressive. Small octopus can learn to solve problems like how to open a container of food. Giant squid, with large nervous systems, have hardly ever been observed except as corpses. They might be as clever as whales.

Birds are vertebrates, related to us through a 300 million year old, probably not very bright, early reptile. Size-limited by the dynamics of flying, some are intellectually comparable to the highest mammals.

The intuitive number sense of crows and ravens extends to seven, compared to three or four for us. Birds outperform all mammals except higher primates and the whales in "learning set" tasks, where the idea is to generalize from specific instances. In mammals generalization depends on cerebral cortex size. In birds forebrain regions called the Wulst and the hyperstriatum are critical, while the cortex is small and unimportant.

Our last common ancestor with the whales was a primitive rat-like mammal alive 100 million years ago. Some dolphin species have body and brain masses identical to ours, and have had them for more generations. They are as good as us at many kinds of problem solving, and can grasp and communicate complex ideas. Killer whales have brains five times human size, and their ability to formulate plans is better than the dolphins', who they occasionally eat. Sperm whales, though not the largest animals, have the world's largest brains. Intelligence may be an important part of their struggle with large squid, their main food. Elephant brains are three times human size. Elephants form matriarchal tribal societies and exhibit complex behavior. Indian domestic elephants

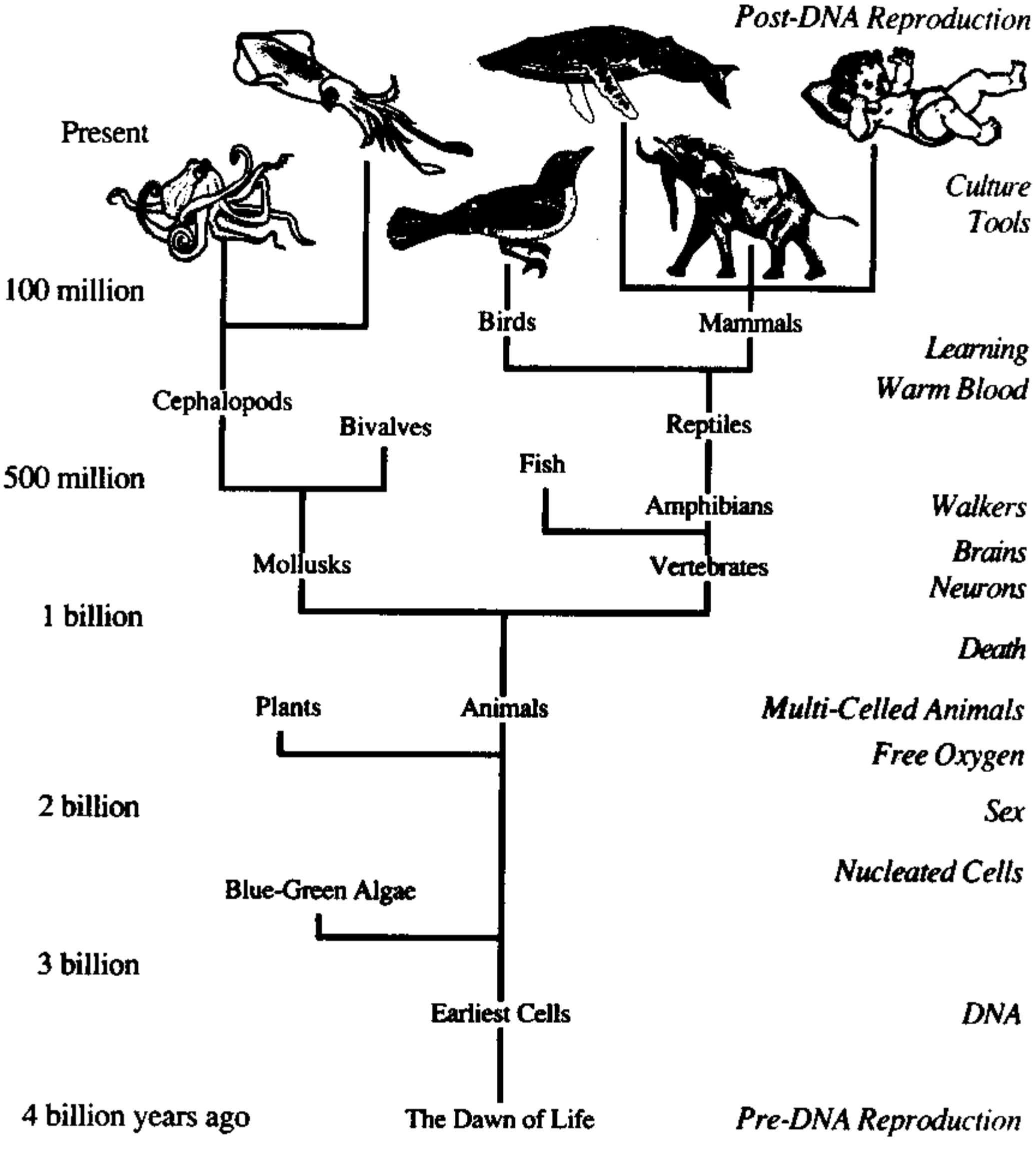

Figure 1: **Intelligence on Earth** - The diagram gives timing, family relationships and significant innovations in the development of terrestrial intelligence. It is likely that very early evolution occurred in an information carrier other than DNA. With the advent of learned behavior in mammals and birds, DNA lost a significant part of its job. More than half of what makes a modern human being is passed culturally. A self reproducing robot economy could end the DNA era altogether - our culture will have freed itself of its roots.

learn over 500 commands, and form voluntary mutual benefit relationships with their trainers, exchanging labor for baths. They can solve problems such as how to sneak into a plantation at night to steal bananas, after having been belled (answer: stuff mud into the bells). And they do have long memories.

Apes are our 10 million year cousins. Chimps and gorillas can learn to use tools and to communicate in human sign languages at a retarded level. Chimps have one third, and gorillas one half, human brainsize.

Animals exhibiting near-human behavior have hundred billion neuron nervous systems. Imaging vision alone requires a billion. The smartest insects have a million brain cells, while slugs and worms make do with a thousand, and sessile animals with a hundred. The portions of nervous systems for which tentative wiring diagrams have been obtained, including nearly all of the large neuroned sea slug, Aplysia, the flight controller of the locust and the early stages of vertebrate vision, reveal neurons configured into efficient, clever, assemblies.

Mobility and Intelligence around the Lab

The twenty year old modern robotics effort can hardly hope to rival the billion year history of large life on earth in richness of example or profundity of result. Nevertheless the evolutionary pressures that shaped life are already palpable in the robotics labs. The following is a thought experiment that we hope soon to make into a physical one.

We desire robots able to execute general tasks such as "go down the hall to the third door, go in, look for a cup and bring it back". This desire has created a pressing need - a computer language in which to concisely specify complex tasks for a rover, and a hardware and software system to embody it. Sequential control languages successfully used with industrial manipulators might seem a good starting point.

Paper attempts at defining the structures and primitives required for the mobile application revealed that the linear control structure of these state-of-the-art arm languages was inadequate for a rover. The essential difference is that a rover, in its wanderings, is regularly "surprised" by events it cannot anticipate, but with which it must deal. This requires that contingency routines be activated in arbitrary order, and run concurrently. One answer is a structure where a number of specialist programs communicating via a common data structure called a blackboard are active at the same time, some operating sensors, some controlling effectors, some integrating the results of other modules, and some providing overall direction. As conditions change the priority of the various modules changes, and control may be passed from one to another.

Character from Motion

Suppose we ask our future robot, equipped with a controller based on the blackboard system mentioned in the last section, to, in fact, go down the hall to the third door, go in, look for a cup and bring it back. This will be implemented as a process that looks very much like a program written for the arm control languages (that in turn look very much like Algol, or Basic), except that the door recognizer routine would probably be activated separately. Consider the following caricature of such a program.

```
module GO-FETCH-CUP
        wake up DOOR-RECOGNIZER with instructions
                ( on FINDING-DOOR add 1 to DOOR-NUMBER
                  record DOOR-LOCATION )

record START-LOCATION
set DOOR-NUMBER to 0
while DOOR-NUMBER < 3 WALL-FOLLOW
FACE-DOOR
if DOOR-OPEN then GO-THROUGH-OPENING
        else OPEN-DOOR-AND-GO-THROUGH
set CUP-LOCATION to result of LOOK-FOR-CUP
TRAVEL to CUP-LOCATION
PICKUP-CUP at CUP-LOCATION
TRAVEL to DOOR-LOCATION
FACE-DOOR
if DOOR-OPEN then GO-THROUGH-OPENING
        else OPEN-DOOR-AND-GO-THROUGH
TRAVEL to START-LOCATION
end
```

So far so good. We activate our program and the robot obediently begins to trundle down the hall counting doors. It correctly recognizes the first one. The second door, unfortunately, is decorated with some garish posters, and the lighting in that part of the corridor is poor, and our experimental door recognizer fails to detect it. The wall follower, however, continues to operate properly and the robot continues on down the hall, its door count short by one. It recognizes door 3, the one we had asked it to go through, but thinks it is only the second, so continues. The next door is recognized correctly, and is open. The program, thinking it is the third one, faces it and proceeds to go through. This fourth door, sadly, leads to the stairwell, and the poor robot, unequipped to travel on stairs, is in mortal danger.

Fortunately there is a process in our concurrent programming system called DETECT-CLIFF that is always running and that checks ground position data posted on the blackboard by the vision processes and also requests sonar and infrared proximity checks on the ground. It combines these, perhaps with an a-priori expectation of finding a cliff set high when operating in dangerous areas, to produce a number that indicates the likelihood there is a drop-off in the neighborhood.

A companion process DEAL-WITH-CLIFF also running continuously, but with low priority, regularly checks this number, and adjusts its own priority on the basis of it. When the cliff probability variable becomes high enough the priority of DEAL-WITH-CLIFF will exceed the priority of the current process in control, GO-FETCH-CUP in our example, and DEAL-WITH-CLIFF takes over control of the robot. A properly written DEAL-WITH-CLIFF will then proceed to stop or greatly slow down the movement of the robot, to increase the frequency of sensor measurements of the cliff, and to slowly back away from it when it has been reliably identified and located.

Now there's a curious thing about this sequence of actions. A person seeing them, not knowing about the internal mechanisms of the robot might offer the interpretation "First the robot was determined to go through the door, but then it noticed the stairs and became so frightened and preoccupied it forgot all about what it had been doing". Knowing what we do about what really happened in the robot we might be tempted to berate this poor person for using such sloppy anthropomorphic concepts as determinination, fear, preoccupation and forgetfulness in describing the actions of a machine. We could berate the person, but it would be wrong.

The robot came by the emotions and foibles indicated as honestly as any living animal - the observed behavior is the correct course of action for a being operating with uncertain data in a dangerous and uncertain world. An octopus in pursuit of a meal can be diverted by hints of danger in just the way the robot was. An octopus also happens to have a nervous system that evolved entirely independently of our own vertebrate version. Yet most of us feel no qualms about ascribing concepts like passion, pleasure, fear and pain to the actions of the animal.

We have in the behavior of the vertebrate, the mollusc and the robot a case of convergent evolution. The needs of the mobile way of life have conspired in all three instances to create an entity that has modes of operation for different circumstances, and that changes quickly from mode to mode on the basis of uncertain and noisy data prone to misinterpretation. As the complexity of the mobile robots increases their similarity to animals and humans will become even greater.

Deeper

Hold on a minute, you say. There may be some resemblance between the robot's reaction to a dangerous situation and an animal's, but surely there are differences. Isn't the robot more like a startled spider, or even a bacterium, than like a frightened human being? Wouldn't it react over and over again in exactly the same way, even if the situation turned out not to be dangerous? You've caught me.

I think the spider's nervous system is an excellent match for robot programs possible today. We passed the bacterial stage in the 1950s with light seeking electronic turtles. This does not mean that concepts like thinking and consciousness are ruled out. In the book **Animal Thinking**, animal ethologist D. G. Griffiths reviews evidence that much animal behavior, including that of insects, can be explained economically in terms of consciousness: an internal model of the self and surroundings, that, however crudely, allows consideration of alternative actions. But there are differences of degree.

Pleasure and Pain

Even single sensory neurons have been shown to habituate to over or under stimulation. Small networks of neurons can adapt in more elaborate ways, for instance by learning to associate one stimulus with another. Such mechanisms tune a nervous system to the body it inhabits, and to its environment. Vertebrates owe much of their potential to an elaboration of this arrangement.

The vertebrate brain has centralized loci for pleasure and pain. Stimulation of a pleasure center acts to encourage future expressions of the preceding behavior, while a pain stimulus discourages it. The archetypical demonstration involves an electrode implanted in the pleasure center of a rat. Allowed to operate a lever that energizes the electrode, the rat ignores food and water to rapidly and repeatedly press the lever, interrupted only by total exhaustion.

Centralization of conditioning sites probably increases long term adaptability. A new need or danger can be accomodated through a small change in the neural wiring, by connecting a detector for the condition to a pleasure or pain site. The standard learning mechanism will then insure that the animal begins to seek the conditions that meet the need, or to avoid the danger, even if the required behavior is complex.

Love and Hate

We are deep in the realm of speculation now, but the general pleasure/pain learning mechanism may provide an explanation for abstract emotions.

Let's suppose that altruism, for instance of a mother towards her offspring, can enhance the long term survival of the altruist's genes even though it has a negative effect on the individual altruist. Feeding the young may leave the mother exhausted and hungry, and defending them may involve her in risk or injury. Shouldn't the conditioning mechanisms we've just described gradually suppress this kind of behavior?

As with more immediate concerns, activities that are multi-generationally beneficial can be encouraged, and ultimately harmful behaviors suppressed, if detectors for them are wired strongly to pleasure and pain centers. For instance, mother love is encouraged if the sight, feel, sound or smell of the offspring triggers pleasure, and if absence of the young is painful.

To the extent that pain or pleasure has a subjective manifestation such non-immediate causes are likely to *feel* different from more obvious ones like skin pain or hunger. Most of the immediate causes are associated with some part of the body, and can be usefully subjectively mapped there. Multi-generational imperatives, on the other hand, cannot be so simply related to the physically apparent world. This may help explain the etherial or spiritual quality that is often associated with such transcendent motivations. Certainly they deserve respect, being the distillation of perhaps tens of millions of years of life or death trials.

What If?

Elaboration of the internal world model in higher animals made possible another twist. A rich world model allows its possessor to examine in detail alternative situations, past, future or merely hypothetical. Dangers avoided can yet be brooded over, and what *might* have happened can be imagined. If the mental simulation is accurate enough, such brooding can produce useful warnings, or point out missed opportunities. These lessons of the imagination are most effective if their consequences are tied to the conditioning mechanism, just as with real events. Such a connection is particularly easy to explain if, as we elaborate below, the most powerful aspects of reason are due to world knowledge powerfully encoded in the sensory and motor systems. The same wiring that conditions on real situations would be activated for imaginary ones.

The ability to imagine must be a key component in communication among higher animals (not to mention between you and me). Messages trigger mental scenarios that then provide conditioning (i.e. learning).

Communication that fails to engage the emotions is not very educational in this sense (yes, yes, we're talking about people now), and a waste of time. Imagine detectors for time well spent and time wasted, themselves wired to the conditioning centers. It's not too far fetched to think that these correspond to the subjective emotions of "interesting" and "boring". Humans seem to have cross wiring that allows elaborate imagining, for instance about future rewards, to deem interesting activities that might normally be boring. How else could there be intellectuals! Indeed, the conventional view of intelligence, and the bulk of work in artificial intelligence, centers on this final twist. While I believe that it is important, it is only a tiny part of the whole story, and often overrated.

Coming Soon

Some facets of the above have been explored, somewhat haphazardly, in machines. British psychologist W. Gray Walter built electronic turtles that demonstrated learning by association, represented as charges on a matrix of capacitors. Arthur Samuel (then of IBM, now Stanford) wrote a checker playing program that adjusted evaluation parameters to improve its play, and was able to learn simply by playing game after game against itself overnight. Frank Rosenblatt of Cornell invented networks of elements that resembled neurons that could be trained to do simple tasks by properly timed punish and reward signals that adjusted the thresholds of synapses between "neurons" that had recently fired. These approaches of the 1950s and 60s fell out of fashion in the last decade, but modern variations are again in vogue.

Among the natural traits in the immediate roving robot horizon is parameter adjustment learning. A precision mechanical arm in a rigid environment can usually have its kinematic self-model and its dynamic control parameters adjusted once, permanently. A mobile robot bouncing around in the muddy world is likely to continuously suffer insults like dirt buildup, tire wear, frame bends and small mounting bracket slips that mess up accurate a-priori models. Our obstacle course software, for instance, has a camera calibration phase. The robot is parked precisely in front of an painted grid of spots. A program notes how the camera images the spots and figures a correction for camera distortions, so that later programs can make precise visual angle measurements. The present code is very sensitive to mis-calibrations, and we are working on a method that will continuously calibrate the cameras just from the images perceived on normal trips through clutter. With such a procedure in place, a bump that slightly shifts one of the robot's cameras will no longer cause systematic errors in its navigation. Animals seem to tune most of their nervous systems with processes of this kind, and such accomodation may be a precursor to more general kinds of learning.

Perhaps more controversially, the begininnings of self awareness can be seen in the robots. All of the control programs of the more advanced mobile robots have internal representations, at varying levels of abstraction and precision, of the world around the robot, and of the robot's position within that world. The motion planners work with these world models in considering alternative future actions for the robot. If the programs had verbal interfaces one could ask questions that receive answers such as "I turned right because I didn't think I could fit through the opening on the left ". As it is the same information is often presented in the form of pictures drawn by the programs.

When?

How does computer speed compare with human thought? The answer has been changing.

The first electronic computers were constructed in the mid 1940s to solve problems too large for unaided humans. *Colossus*, one of a series of ultrasecret British machines, broke the German *Enigma* code, greatly influencing the course of the European war, by scanning through code keys tens of thousands of times faster than humanly possible. In the US *Eniac* computed antiaircraft artillery tables for the Army, and later did calculations for the atomic bomb, at similar speeds. Such feats earned the early machines the popular appellation *Giant Brains*.

In the mid 1950s computers more than ten times faster than Eniac appeared in many larger Universities. They did numerical scientific calculations nearly a million times faster than humans. A few visionaries took the Giant Brains metaphor seriously and began to write programs for them to solve intellectual problems going beyond mere calculation. The first such programs were encouragingly successful. Computers were soon solving logic problems, proving theorems in Euclidean geometry, playing checkers, even doing well in IQ test analogy problems. The performance level and the speed in each of these narrow areas was roughly equivalent to that of a college freshman who had recently learned the subject. The automation of thought had made a great leap, but paradoxically the term "Giant Brains" seemed less appropriate.

In the mid 1960s a few centers working in this new area of *Artificial Intelligence* added another twist: mechanical eyes, hands and ears to provide real world interaction for the thinking programs. By then computers were a thousand times faster than Eniac, but programs to do even simple things like clearing white blocks from a black tabletop turned out to be very difficult to write, and performed hundreds of times more slowly, and much less reliably, than a human. Slightly more complicated tasks took much longer, and many seemingly trivial things, like identifying a few simple objects in a jumble, still cannot be done acceptably at all twenty years later, even given hours of computer time. Forty years of research and a millionfold increase in computer power has reduced the image of computers from Giant Brains to mental midgets. Is this silly, or what?

Easy and Hard

The human evolutionary record provides a clue to the paradox. While our sensory and muscle control systems have been in development for almost a billion years, and common sense reasoning has been honed for perhaps a million, really high level, deep, thinking is little more than a parlor trick, culturally developed over a few thousand years, which a few humans, operating largely against their natures, can learn. As with Samuel Johnson's dancing dog, what is amazing is not how well it is done, but that it is done at all.

Computers can challenge humans in intellectual areas, where humans are evolutionary novices, because they can be programmed to carry on much less wastefully. Arithmetic is an extreme example, a function learned by humans with great difficulty, but instinctive to computers. A 1987 home computer can add a million large numbers in a second, astronomically faster than a person, and with no errors. Yet the 100 billion neurons in a human brain, if reorganized by a mad neurosurgeon into adders using switching logic design principles, could sum one hundred thousand times faster than the computer.

Computers do not challenge humans in perceptual and control areas because these ancient functions are carried out by large fractions of the nervous system wired for those jobs as cleanly as the hypothetical neuron adder above. Present day computers, however efficiently programmed, are simply too puny to keep up. Evidence comes from the classic program of reverse engineering of some of the visual system of vertebrates initiated by David Hubel and Thorsten Weisel in the 1960s. They elucidated some of its operation by microscopically examining the structure of the retina and the vision-related parts of the brains of cats and monkeys and using electrodes to monitor signals there as test patterns were presented to the eyes.

The vertebrate retina consists of a highly organized, ten layered, structure of densely packed neurons fed by about one hundred million light sensors (figure 2). The sensors are merged into clusters giving an effective resolution of one million picture elements, or *pixels*, to use the computer jargon. The other neurons combine their outputs in various ways to detect such things as edges, corners, curvature and motion. Each of these simple *operators* employs about 10 to 100 neurons per pixel. Thus processed, the image goes via the optic nerve to the much bigger visual cortex in the brain.

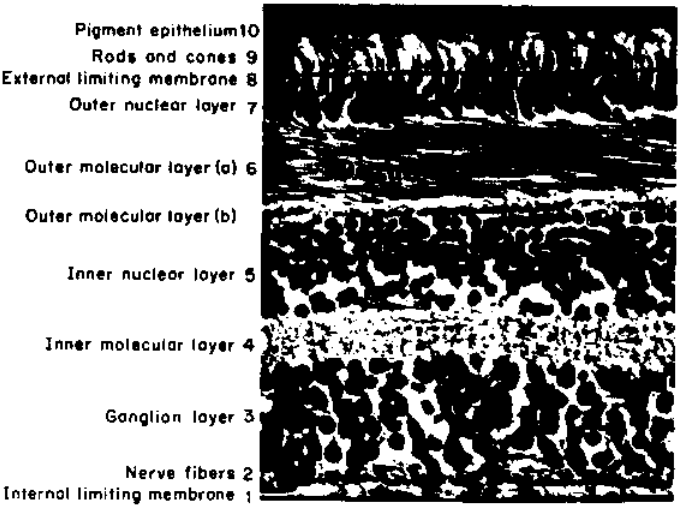

Figure 2: **Human Retina** - This cross section shows a tiny portion of each of the ten layers of neurons that form the retina. There are about 400 million efficiently organized neurons in all. (From page 185 of Photoprocesses, Photoreceptors and Evolution by Jerome J. Wolken, Academic Press, 1975.)

Assuming the visual cortex does as much computing for its size as the retina (perhaps an overestimate - the retina is a small, old and highly optimized structure; larger and more recent regions may use neurons less efficiently), we can estimate the total capability of the system. The visual cortex has about 10 billion neurons, a thousand times the number in a modest retinal operation. The eye can process ten images a second, so the cortex may do the computational equivalent of 10,000 small retinal operations a second.

Operations similar to the retinal ones have been found very useful for robot vision. An efficient program running on a typical (1 million instruction per second) computer can do the equivalent work of one small retinal operation in about 10 seconds. Thus, seeing programs on present day computers seem to be 100,000 times slower than human vision.

The whole brain is about ten times larger than the visual system, so it should be possible to write real-time human equivalent programs for a machine one million times more powerful than today's medium sized computer. In 1987 the largest supercomputers are about 1000 times slower than this desiratum.

Intellectual Voyages

Interesting computation and thought requires a processing engine of sufficient computational *power* and *capacity*. Roughly, power is the speed of the machine, and capacity is its memory size.

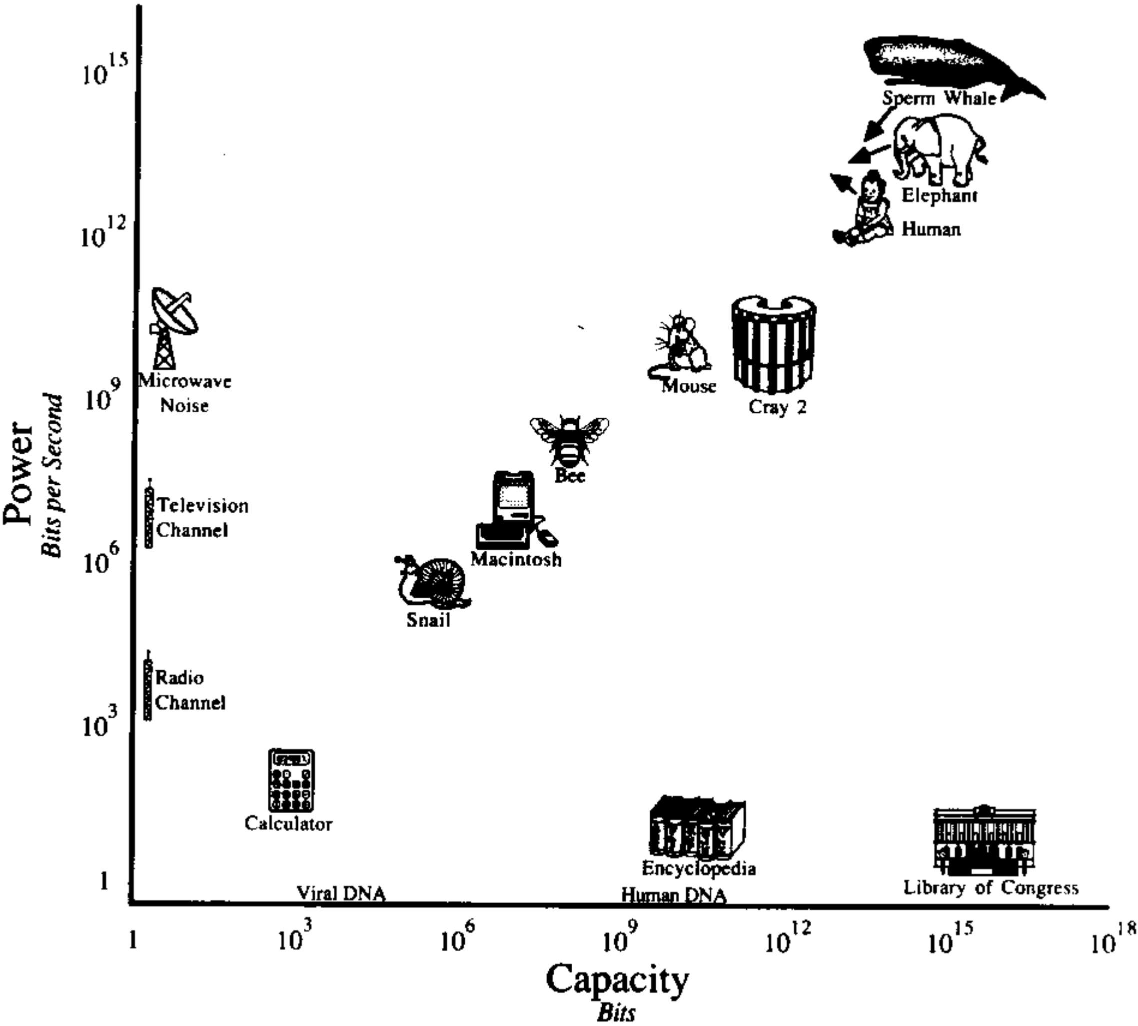

Figure 3: **Think Power -** Computing speed and memory of some animals and machines. The animal figures are for the nervous system only, calculated at 100 bits per second and 100 bits of storage per neuron. These are speculative estimates, but note that a factor of 100 one way or the other would change the appearance of the graph only slightly.

Here's a helpful metaphor. Computing is like a sea voyage in a motorboat. How fast a given journey can be completed depends on the power of the boat's engine. The maximum length of any journey is limited by the capacity of its fuel tank. The effective speed is decreased, in general, if the course of the boat is constrained, for instance to major compass directions.

Some computations are like a trip to a known location on a distant shore, others resemble a mapless search for a lost island. Parallel computing is like having a fleet of small boats - it helps in searches, and in reaching multiple goals, but not very much in problems that require a distant sprint. Special purpose machines trade a larger engine for less rudder control.

Attaching disks and tapes to a computer is like adding secondary fuel tanks to the boat. The capacity, and thus the range, is increased, but if the connecting plumbing is too thin, it will limit the fuel flow rate and thus the effective power of the engine.

Extending the metaphor, input/output devices are like boat sails. They capture power and capacity in the environment. Outside information is a source of variability, and thus power, by our definition. More concretely, it may contain answers that would otherwise have to be computed. The external medium can also function as extra memory, increasing capacity.

Figure 3 shows the power and capacity of some interesting natural and artificial thinking engines. At its best, a computer instruction has a few tens of bits of information, and a million instruction per second computer represents a few tens of millions of bits/second of power. The power ratio between nervous systems and computers is as calculated in the last section: a million instructions per second is worth about a hundred thousand neurons. I also assume that a neuron represents about 100 bits of storage, suggested by recent evidence of synaptic learning in simple neurvous systems by Eric Kandel and others. Note that change of a factor of ten or even one hundred in these ratios would hardly change the graph qualitatively. (My forthcoming book **Mind Children**, from which this paper is drawn, offers more detailed technical justifications for these numbers).

The figure shows that current laboratory computers are equal in power approximately to the nervous systems of insects. It is these machines that support essentially all the research in artificial intelligence. No wonder the results to date are so sparse! The largest supercomputers of the mid 1980s are a match for the 1 gram brain of a mouse, but at ten million dollars or more apiece they are reserved for serious work.

The Growth of Processing Power

How long before the research medium is rich enough for full intelligence?

Although a number of mechanical digital calculators were devised and built during the seventeenth and eighteenth centuries, only with the mechanical advances of the industrial revolution did they become reliable and inexpensive enough to routinely rival manual calculation. By the late nineteenth century their edge was clear, and the continuing

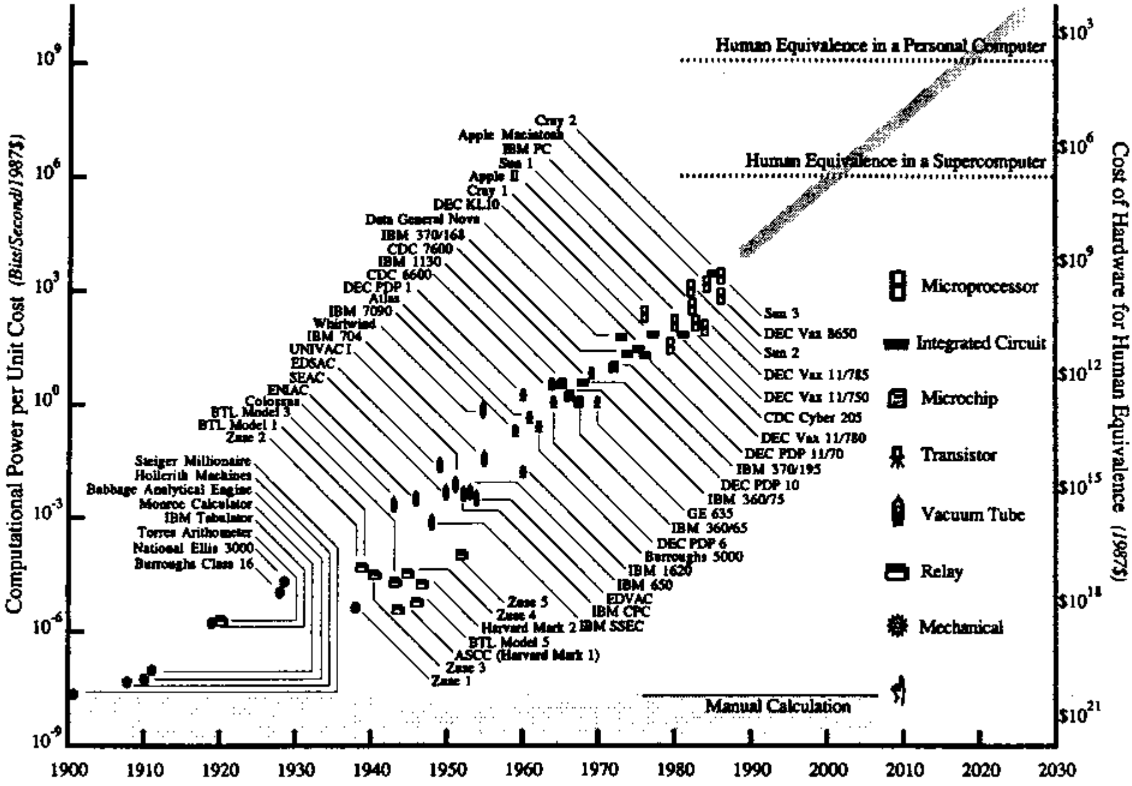

Figure 4: **A Century of Computing** - The cost of calculation has dropped a thousandfold every twenty years (or halved every two years) since the late nineteenth century. Before then mechanical calculation was an unreliable and expensive novelty with no particular edge over hand calculation. The graph shows a mind boggling *trillionfold* decrease in the cost since then. The pace has actually picked up a little since the beginning of the century. It once took 30 years to accumulate a thousandfold improvement; in recent decades it takes only 19. Human equivalence should be affordable very early in the 21st century.

progress dramatic. Since then the cost of computing has dropped a thousandfold every twenty years (figure 4).

The early improvements in speed and reliability came with advances in mechanics - precision mass produced gears and cams, for instance, improved springs and lubricants, as well as increasing design experience and competition among the calculator manufacturers. Powering calculators by electric motors provided a boost in both speed and automation in the 1920s, as did incorporating electromagnets and special switches in the innards in the 1930s. Telephone relay methods were used to make fully automatic computers during World War II, but these were quickly eclipsed by electronic tube computers using radio, and ultrafast radar, techniques. By the 1950s computers were an industry that itself spurred further major component improvements.

The curve in figure 4 is not leveling off, and the technological pipeline is full of developments that can sustain the pace for the foreseeable future. Success in this enterprise, as in others, breeds success. Not only is an increasing fraction of the best human talent engaged in the research, but the ever more powerful computers themselves feed the process. Electronics is riding this curve so quickly that it is likely to be the main occupation of the human race by the end of the century.

The price decline is fueled by miniaturization, which supplies a double whammy. Small components both cost less and operate more quickly. Charles Babbage, who in 1834 was the first person to conceive the idea of an automatic computer, realized this. He wrote that the speed of his design, which called for hundreds of thousands of mechanical components, could be increased in proportion if "as the mechanical art achieved higher states of perfection" his palm sized gears could be reduced to the scale of clockwork, or further to watchwork. (I fantasize an electricityless world where the best minds continued on Babbage's course. By now there would be desk and pocket sized mechanical computers containing millions of microscopic gears, computing at thousands of revolutions per second.)

To a remarkable extent the cost per pound of machinery has remained constant as its intricacy increased. This is as true of consumer electronics as of computers (merging categories in the 1980s). The radios of the 1930s were as large and as expensive as the televisions of the 1950s, the color televisions of the 1970s, and the home computers of the 1980s. The volume required to amplify or switch a single signal dropped from the size of a fist in 1940, to that of a thumb in 1950, to a pencil eraser in 1960, to a salt grain in 1970, to a small bacterium in 1980. In the same period the basic switching speed rose a millionfold, and the cost declined by the same huge amount.

Predicting the detailed future course is impossible for many reasons. Entirely new and unexpected possibilities are encountered in the course of basic research. Even among the known, many techniques are in competition, and a promising line of development may be abandoned simply because some other approach has a slight edge. I'll content myself with a short list of some of what looks promising today.

In recent years the widths of the connections within integrated circuits have shrunk to less than one micron, perilously close to the wavelength of the light used to "print" the circuitry. The manufacturers have switched from visible light to shorter wavelength ultraviolet, but this gives them only a short respite. X-rays, with much shorter wavelengths, would serve longer, but conventional X-ray sources are so weak and diffuse that they need uneconomically long exposure times. High energy particle physicists have an answer. Speeding electrons curve in magnetic fields, and spray photons like mud from a spinning wheel. Called synchotron radiation for the class of particle accelerator where it became a nuisance, the effect can be harnessed to produce powerful beamed X-rays. The stronger the magnets, the smaller can be the synchotron. With liquid helium cooled superconductiong magnets an adequate machine can fit into a truck, otherwise it is the size of a small building. Either way, synchotrons are now an area of hot interest, and promise to shrink mass-produced circuitry into the sub-micron region. Electron and ion beams are also being used to write submicron circuits, but present systems affect only small regions at a time, and must be scanned slowly across a chip. The scanned nature makes computer controlled electron beams ideal, however, for manufacturing the "masks" that act like photographic negatives in circuit printing.

Smaller circuits have less electronic "inertia" and switch both faster and with less power. On the negative side, as the number of electrons in a signal drops it becomes more prone to thermal jostling. This effect can be countered by cooling, and indeed very fast experimental circuits can now be found in many labs running in supercold liquid nitrogen, and one supercomputer is being designed this way. Liquid nitrogen is produced in huge amounts in the manufacture of liquid oxygen from air, and it is very cheap (unlike the much colder liquid helium).

The smaller the circuit, the smaller the regions across which voltages appear, calling for lower voltages. Clumping of the substances in the crystal that make the circuit becomes more of a problem as they get smaller, so more uniform "doping" methods are being developed. As the circuits become smaller quantum effects become more pronounced, creating new problems and new opportunities. Superlattices, mutiple layers of atoms-thick regions of differently doped silicon made with molecular beams, are such an opportunity. They allow the electronic characteristics of the material to be tuned, and permit entirely new switching methods, often giving tenfold improvements.

The first transistors were made of germanium; they could not stand high temperatures and tended to be unreliable. Improved understanding of semiconductor physics and ways of growing silicon crystals made possible faster and more reliable silicon transistors and integrated circuits. New materials are now coming into their own. The most immediate is gallium arsenide. Its lattice impedes electrons less than silicon, and makes circuits up to ten times faster. The Cray 3 supercomputer due in 1989 will use gallium arsenide integrated circuits, packed into a one cubic foot volume, to top the Cray 2's speed tenfold. Other compounds like indium phosphide and silicon carbide wait in the wings. Pure carbon in diamond form is an obvious possibility - it should be as much an improvement over Gallium Arsenide as that crystal is over Silicon. Among its many superlatives, perfect diamond is the best solid conductor of heat, an important property in densely packed circuitry. The vision of an utradense three dimensional circuit in a gem quality diamond is compelling. As yet no working circuits of diamond have been reported, but excitment is mounting as reports of diamond layers up to a millimeter thick grown from hot methane come from the Soviet Union, Japan and, belatedly, the United States.

Farther off the beaten track are optical circuits that use lasers and non-linear optical effects to switch light instead of electricity. Switching times of a few picoseconds, a hundred times faster than conventional circuits, have been demonstrated, but many practical problems remain. Finely tuned laser has also been used with light sensitive crystals and organic molecules in demonstration memories that store up to a trillion bits per square centimeter.

The ultimate circuits may be superconducting quantum devices, which are not only extremely fast, but extremely efficient. Various superconducting devices have been in and out of fashion several times over the past twenty years. They've had a tough time because the liquid helium environment they require is expensive, the heating/cooling cycles are stressful, and especially because rapidly improving semiconductors have offered such tough competition.

Underlying these technical advances, and preceding them, are equally amazing advances in the methods of basic physics. One recent, unexpected and somewhat unlikely, device is the inexpensive tunnelling microscope that can reliably see, identify and soon manipulate single atoms on surfaces by scanning them with a very sharp needle. The tip is positioned by three piezoelectric crystals microscopically moved by small voltages. It maintains a gap a few atoms in size by monitoring a current that jumps across it. The trickiest part is isolating the system from vibrations. It provides our first solid toehold on the atomic scale.

A new approach to miniaturization is being pursued by enthusiasts in the laboratories of both semiconductor and biotechnology companies, and elswhere. Living organisms are clearly machines when viewed at the molecular scale. Information encoded in RNA "tapes" directs protein assembly devices called ribosomes to pluck particular sequences of amino acids from their environment and attach them to the ends of growing chains. Proteins, in turn, fold up in certain ways, depending on their sequence, to do their jobs. Some have moving parts acting like hinges, springs, latches triggered by templates. Others are primarily structural, like bricks or ropes or wires. The proteins of muscle tissue work like ratcheting pistons. Minor modifications of this existing machinery are the core of today's biotechnology industry. The visionaries see much greater possibilities.

Proteins to do specific jobs can be engineered even without a perfect model of their physics. Design guidelines, with safety margins to cover the uncertainties, can substitute. The first generation of artificial molecular machinery would be made of protein by mechanisms recruited from living cells. Early products would be simple, like tailored medicines, and experimental, like little computer circuits. Gradually a bag of tricks, and computer design aids, would accumulate to build more complicated machines. Eventually it may be possible to build tiny robot arms, and equally tiny computers to control them, able to grab molecules and hold them, thermally wriggling, in place. The protein apparatus could then be used as machine tools to build a second generation of molecular devices by assembling atoms and molecules of all kinds. For instance, carbon atoms might be laid, bricklike, into ultra strong fibers of perfect diamond. The smaller, harder, tougher machines so produced would be the second generation molecular machinery.

The book **Engines of Creation** by Eric Drexler, and a forthcoming book by Conrad Schneiker, call the entire scheme *nanotechnology*, for the nanometer scale of its parts. By contrast today's integrated circuit microtechnology has micrometer features, a thousand times bigger. Some things are easier at the nanometer scale. Atoms are perfectly uniform in size and shape, if somewhat fuzzy, and behave predictably, unlike the nicked, warped and cracked parts in larger machinery.

A Stumble

It seemed to me throughout the 1970s (I was serving an extended sentence as a graduate student at the time) that the processing power available to AI programs was not increasing very rapidly. In 1970 most of my work was done on a Digital Equipment Corp. PDP-10 serving a community of perhaps thirty people. In 1980 my computer was a DEC KL-10, five times as fast and with five times the memory of the old machine, but with twice as many users. Worse, the little remaining speedup seemed to have been absorbed in computationally expensive convenience features: fancier time sharing and high level languages, graphics, screen editors, mail systems, computer networking and other luxuries that soon became necessities.

Several effects together produced this state of affairs. Support for university science in general had wound down in the aftermath of the Apollo moon landings and politics of the Vietnam war, leaving the universities to limp along with aging equipment. The same conditions caused a recession in the technical industries - unemployed engineers opened fast food restaurants instead of designing computers (the rate of change in figure 4 does slacken slightly in the mid 1970s). The initially successful "problem solving" thrust in AI had not yet run its course, and it still seemed to many that existing machines were powerful enough - if only the right programs could be found. While spectacular progress in the research became increasingly difficult, a pleasant synergism among the growing number of information utilities on the computers created an attractive diversion for the best programmers - creating more utilities.

If the 1970s were the doldrums, the 1980s more than compensated. Several salvations had been brewing. The Japanese industrial successes focused attention worldwide on the importance of technology, particularly computers and automation, in modern economies - American industries and government responded with research dollars. The Japanese stoked the fires, under the influence of a small group of senior researchers, by boldly announcing a major initiative towards future computers, the so called "Fifth Generation" project, pushing the most promising American and European research directions. The Americans responded with more money. Besides this, integrated circuitry had evolved far enough that an entire computer could fit on a chip. Suddenly computers were affordable by individuals, and a new generation of computer customers and manufacturers came into being. The market was lucrative, the competition fierce, and the evolution swift, and by the mid 1980s the momentum lost in the previous decade had been regained, with interest. Artificial intelligence research is awash in a cornucopia of powerful new "personal" workstation computers, and there is talk of applying supercomputers to the work.

Even without supercomputers, human equivalence in a research setting should be possible by around 2010, as suggested by figure 4. Now, the smallest vertebrates, shrews and hummingbirds, get interesting behavior from nervous systems one ten thousandth the size of a human's, so I expect fair motor and perceptual competence, in about a decade.

Faster Yet?

Very specialized machines can provide up to one thousand times the effective performance for a given price in well defined tasks. Some vision and control problems may be candidates for this approach. Special purpose machines are not a good solution in the groping research stage, but may dramatically lower the costs of intelligent machines when the problems and solutions are well understood. Some principals in the Japanese Fifth Generation Computer Project have been quoted as planning "man capable" systems in ten years. I believe this more optimistic projection is unlikely, but not impossible.

As the computers become more powerful and as research in this area becomes more widespread the rate of visible progress should accelerate. I think artificial intelligence via the "bottom up" approach of technological recapitulation of the evolution of mobile animals is the surest bet because the existence of independently evolved intelligent nervous systems indicates that there is an incremental route to intelligence. It is also possible, of course, that the more traditional "top down" approach will achieve its goals, growing from the narrow problem solvers of today into the much harder areas of learning, common-sense reasoning and perceptual acquisition of knowledge as computers become large and powerful enough, and the techniques are mastered. Most likely both approaches will make enough progress that they can effectively meet somewhere in the middle, for a grand synthesis into a true artificial sentience.

This artificial person will have some interesting properties. Its high level reasoning abilities should be astonishingly better than a human's - even today's puny systems are much better in some areas - but its low level perceptual and motor abilities will be comparable to ours. Most interestingly it will be highly changeable, both on an individual basis and from one of its generations to the next. And it will quickly become cheap.

Emerging Manipulator Technology

Design of the Utah/MIT Dextrous Hand

Stephen C. Jacobsen
Professor
Center for Engineering Design
University of Utah

ABSTRACT

The Center for Engineering Design at the University of Utah, and the Artificial Intelligence Laboratory at the Massachusetts Institute of Technology have developed a robotic end effector intended to function as a general purpose research tool for the study of machine dexterity. The high performance, multi-fingered hand will provide two important capabilities. First, it will permit the experimental investigation of basic concepts in manipulation theory, control system design and tactile sensing. Second, it will expand understanding required for the future design of physical machinery and will serve as a "test bed" for the development of tactile sensing systems. The paper includes: 1) a discussion of issues important to the development of manipulation machines; 2) general comments regarding design of the Utah/M.I.T. Dextrous Hand; and, 3) a detailed discussion of specific subsystems of the hand.

INTRODUCTORY COMMENTS - MACHINES THAT PERFORM MANIPULATION FUNCTIONS

Expectations Are High

Judging by recent public interest, expectations are high that machines which execute complex manipulation functions will be a reality in the near future. This wave of optimism is encouraged by science fiction presentations depicting high performance robots; by the emergence of toys which move, interact and execute locomotion functions; and by certain industrial demonstrations of robotic manipulation systems which have been carefully orchestrated to convey images which imply that robots can do much more than is currently possible.

Apparent Markets Motivate Development Efforts

In addition to the general fascination which many people have for robots, the principal motivation for financial investment in the development of advanced manipulation machines arises from three commercial opportunities. First, it is believed that systems with machine dexterity can profitably perform a variety of tasks in the industrial sector; for example, assembly procedures, interactive inspection functions, etc. Second, it is obvious that such systems permit remote human presence in hostile or distant environments such as space and undersea. And, third, advanced manipulation machines could provide expanded capabilities for various military systems.

In all of the above areas the central objective is the replacement or augmentation of humans by robots which can execute manipulation tasks with: 1) greater economy; 2) higher performance; and, 3) reduced possibility of injury to people.

Progress Has Been Slower Than Expected

Unfortunately, the development of machine manipulation systems is emerging as a problem much more difficult than previously anticipated. General theories which permit the development of successful control systems and sensors haven't "fallen out" easily and the design and construction of machinery which can execute desired physical interactions has evolved slowly. A number of serious failures have occured both in the commercial sector and in the research and development community. Optimistic claims have aroused unrealistic expectations that progress would be rapid. Disagreements among participants in research projects have arisen as a result of failures and significant confusion now exists regarding how to proceed in the future.

It is becoming clear that progress has been slow for three basic reasons. First, the problems associated with the design and construction of machine-based manipulation systems are enormously difficult. It should be expected that some time will be required before issues are fully understood and demonstrable results are achieved. Second, although there has been much publicity regarding robot manipulation systems, intensive focus, in terms of effort and resources, has only been applied to this area for a short time. It is clear that significant resources will be required to achieve successful systems. Third, and unfortunately, work in this area has been undertaken with little comprehensive management and as a result balanced efforts have occurred in only rare circumstances. Future projects must be encouraged to evolve in a balanced way, with efforts expended simultaneously in both theoretical and applied areas. Major focus should be on understanding issues related to control systems, sensing systems, and in understanding how these complex machines can be made to function reliably.

Future Efforts Should Consider: Utility, Feasibility, Persistence, Balance, Experimentation and Demonstration

In order to provide a basis for understanding decisions made during the conduct of this work, as well as guiding future efforts to develop manipulation machines, a number of general issues should be reviewed.

Existence of a Market. In a long-term, general sense there can be little doubt of the utility (disregard cost for the moment) of devices which perform comprehensive manipulation tasks. Numerous examples exist where intelligent and dexterous machines could advantageously accomplish manufacturing functions, operate in hazardous environments, explore remote locations, provide changes in scale, function at high or low speeds, or accomplish other important functions. The economic opportunities which could be generated by such systems are indeed enormous.

Clearly, future systems must be practical, reliable and within cost guidelines. However, care must be taken to avoid excessively negative projections just because present research systems are awkward and disappointing. It should be remembered that economic success will be defined in terms of the achievement of specific and relatively limited goals, rather than generating those fanciful images seen in science fiction movies. In each situation the application and resulting expectations must be understood in terms of requirements for performance, reliability, and economy. Judgments should then be based on appropriate cost/benefit issues. Too many times a research tool is evaluated according to industrial standards or an industrial machine is considered inadequate based on research demands.

Evaluation of a system based purely on cost can also be a fuzzy process. In some applications it is permissible for a manipulation system to be quite expensive since its costs are dwarfed by the expenses of the overall program; for example, in outer space or undersea applications. In other cases, such as those in manufacturing environments, the cost of a manipulation system must be less. Finally, in the case of consumer products such as toys and home robots, the cost of individual systems must be extremely small.

Feasibility. There can be little doubt that manipulation processes can be accomplished by general purpose systems. Biology provides a stunning array of existence proofs which demonstrate that intricate manipulation is within the realm of possibility. It should be emphasized here, for those with certain presumptions, that the operation of biological machines is not based in magic. Nature produces systems which utilize real hardware that operates according to physical principles. Although it can be said that the basis for many biological functions is not yet well understood, it is undoubtedly true that, as research unfolds, many important and useful discoveries will be made in this area. Finally, the intent here is not to imply that the development of such systems will be an easy task, only that such systems can be developed.

Persistence. If machine-based manipulation systems have markets and are feasible in at least some forms, it is obvious that work should proceed with vigor in order to shorten the time before such valuable systems emerge. Of initial importance is establishing a clear and balanced understanding of the issues involved so that projects can be managed with emphasis on achievement of practical results.

Balance in Approach. The development process should be managed in order to identify goals which are realistic so that results can efficiently drive an expansion of the knowledge and technology base. Early activities should target achievable goals which expand understanding of basic issues related to: 1) task definition; 2) manipulation strategies; 3) grasping functions; 4) the use and integration of sensory information; and, 5) fundamental issues governing the design of machinery which will constitute these man-made manipulation systems.

Above all, the development process should emphasize an expansion of interdisciplinary collaboration by encouraging tolerance for multiple viewpoints and methods of approach. The ultimate goal should be understanding those principles which permit the design of successful systems. An example of current intolerance is the conflict which continually surfaces between those groups which follow traditional engineering approaches versus groups which include in their activities, the study of biological systems in order to gain insights into the manipulation process. The traditional and time-tested engineering approach is to construct a sequence of prototype machines which can be used to evaluate principles. Then, based on those principles, final systems are designed. A second approach, equally valid, but sometimes criticized, is the attempt to identify important principles through the study of biological systems. In the final analysis, both approaches, as well as other methodologies, can produce desirable results, and in all cases, information which could facilitate progress should never be ignored.

Demonstration. As always, there is a limit to the human capacity for pure visualization and conceptualization. At some point, physical machines must be constructed to: 1) provide a means by which concepts can be experimentally validated; 2) to provide indications as to the performance which will be required of physical machinery; and, 3) to permit judgments to be made regarding the value of a particular approach. Unfortunately, in the manipulation area there has been a significant lack of experimentation due primarily, not to conscious avoidance, but to the unavailability of suitable research equipment.

The Utah/M.I.T. Dextrous Hand - A High Performance Research Tool for the Study of Machine Dexterity

Just as research is conducted in a number of areas using reliable, flexible, general purpose computation systems, so work in the machine manipulation area should be conducted with well developed manipulation systems which can be used as "test beds" to explore concepts related to both manipulation theory and machine design. The existence of such tools can simplify research activities and allow investigators to proceed toward understanding issues and concepts rather than being continually sidetracked by problems with experimental devices.

It was therefore decided that an effort would be undertaken to develop a research tool with: 1) many degrees of freedom; 2) very high active and passive performance; 3) acceptable reliability; 4) the capability to serve as a vehicle for studying a broad variety of tactile sensing systems; and, 5) via modularity, the possibility of geometric reconfiguration in order to address evolving experimental objectives. In 1982 the Utah/M.I.T. Dextrous Hand Project began with its objective being the development of the hand which is schematically illustrated in Figure 1. Since that time final Version III system has been produced as shown in Figure 2. The following sections describe the Dextrous Hand and review a number of general issues important to the design of such systems.

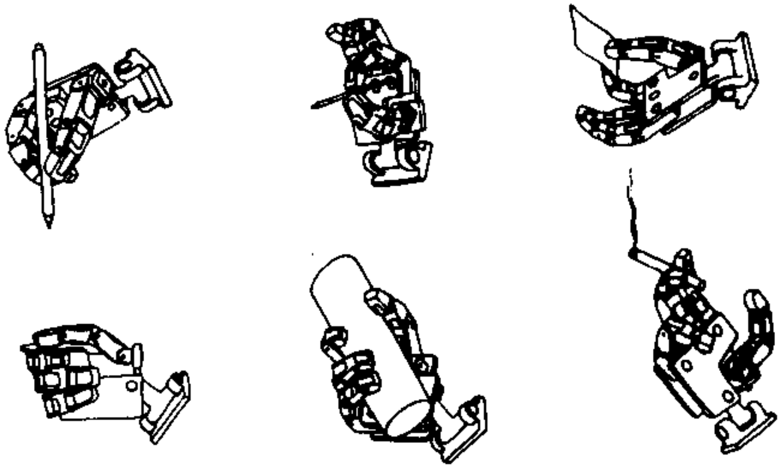

Figure 1. Originally proposed configuration for the Utah/M.I.T. Dextrous Hand. Reprinted from Reference 1.

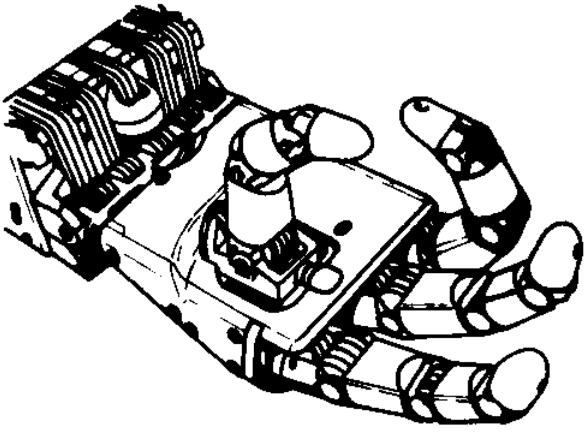

Figure 2. Line drawing of the Version III Dextrous Hand.

GENERAL COMMENTS REGARDING DESIGN OF THE UTAH/M.I.T. DEXTROUS HAND

Goals of the Project

The primary objective of the project has been the design and construction of a general purpose research tool for the study of machine-based manipulation. From the outset, the emphasis has been on producing a system with extraordinary performance (both active and passive), and as such, factors related to cost, power consumption and other practical considerations have been de-emphasized. The device is not intended for near-term application in an industrial environment, but as a research tool with sufficient functional richness to permit a broad variety of experiments to be conducted, aimed at understanding fundamental concepts and machine design issues. Goals can be reviewed in greater detail by addressing specific performance issues such as: 1) speed; 2) strength; 3) range of motion; 4) the capability for graceful behavior; 5) reliability; 6) the possibility of reconfiguration; and, 7) cost.

Speed. The speed of the digits of the Hand have been designed in anticipation of performing dynamic manipulation tasks. This requires that the fingers be capable of interacting with objects which might fall or move rapidly under the influence of gravity. Such motions require that fingers execute rapid motions with frequency components which exceed 10 Hz. At these frequencies the digits should move in such a manner that manipulated objects are not damaged by rapid finger contact. This requires that actuated masses be low so that high speed contacts produce minimal impact. Furthermore, the actuation system should be capable of tolerating large accelerations imposed by interactions with moving objects.

Strength. Since the Hand must interact with objects of "hand size," and considering that such objects will have densities ranging from Styrofoam to metals, and that lubricated surfaces will exhibit relatively low coefficients of friction, strength of the DH has been addressed with considerable emphasis. The present system is capable of producing tip forces of 7 pounds which permit very positive fixation of objects, especially when multiple fingers operate in a coordinated fashion. Even though strength of the DH is acceptable, a current objective is to identify ways to triple the strength of the hand while maintaining its speed capabilities.

Graceful Behavior. Gracefulness is a quality which can be observed in natural systems of many types. It is probably true that graceful behavior is not a feature added for aesthetic appeal, but that grace is generally a by-product of a system operating with low internal loadings on structures, and without excessive internal antagonisms which lead to inefficient operation. In other words, grace and efficiency appear to go together. The Utah/M.I.T. Hand is quite graceful in operation and this quality is primarily a result of the impedance characteristics of the pneumatic actuators. The actuators have been designed to exhibit low stiffness, low stiction, and the mass of actuated elements is low since only the tendons, pulleys and graphite pistons actually move to operate finger joints. These factors lead to a very distinctive "feel" when the DH is externally manipulated. In fact, the individual digits can be externally driven up to very high frequencies without damaging internal components.

It should also be noted here that, although compliant operation of the system is advantageous in many ways, it also presents problems for the control system. Soft systems tend to destabilize if high loop gains are used to achieve positional accuracy. It could be said here, loosely speaking, that during free motions of a manipulator, grace is a result of the "smart throwing" of segments. It is interesting that the capability to be thrown depends or their being capable of maintaining low impedance even at high speeds. Also, being thrown through space to some future rendezvous requires knowledge of the system and its task. Thus, in the case of the DH, the control system must include greater complexity in order to govern the behavior of compliant fingers involved in activities requiring rapid and precise coordination.

Reliability. The Utah/M.I.T. Dextrous Hand was not designed as a laboratory trinket for short-term investigation of some local principle. The system was designed to be a functional, reliable machine aimed at long-term operation. Subelements have been exhaustively evaluated and the design continually reviewed in order to provide information necessary to enhance the performance and reliability of future systems. Of course, subsystem reliability will become an increasingly important issue as this DH and other similar machines evolve to greater levels of complexity.

The Possibility of Reconfiguration. The Dextrous Hand was designed as a series of modules which can be interchanged for maintenance purposes, and also to permit reconfiguration of the system into alternate geometries. Even though all four digits are essentially identical, outer surfaces can be modified and base mountings repositioned in order to alter finger shape and grasp geometry. Since the system is tendon driven, the 32 actuation systems can be used in end effectors of almost any configuration by simply rerouting tendons. The low level primary control system shown in Figure 17 is also flexible in that it can, via inputs from the higher control system, execute system servo functions with dynamically varying loop gains according to the specific requirements of the end effector and its task.

Cost. Since the initial system is intended to be a research tool manufactured in small quantities, cost has been a secondary consideration. Nevertheless, system modules have been designed to permit low cost if manufactured in significant numbers. It is intended that, subsequent to a period of research, simplified versions of the Utah/M.I.T. Dextrous Hand will be produced for industrial application using less expensive modules evolved from the original system.

On the Choice of Anthropomorphic Geometry

At the beginning of the project a search was made for previous work which would indicate the most desirable configuration for a research oriented robotic end effector. A number of papers were found which reviewed the subject under a rather restricted set of assumptions which related to: 1) the number of actuated elements (fingers) and the geometry of their contact with manipulated objects; 2) the quality of contact between fingers and objects (frictional characteristics); 3) the geometric complexity of the manipulated object (planar versus three dimensional objects); 4) the type of sensory information available to the central controller (joint torques, touch information, etc.) and other issues. The analyses included only simple types of static grasp using fingers with localized contact points and no real consideration of the palm as a valuable platform against which fingers can position objects. Furthermore, papers included no discussion of hand geometry and its influence on higher control issues such as those relating to task, manipulation strategy and grasping functions. Previous work was, due to the complexity of the problem, quite inconclusive in the sense that no clear direction was inferred regarding end effector configuration.

It was therefore decided to attempt an anthropomorphic geometry for the following three reasons. First, the human hand is an existence proof that a wide variety of very complicated, manipulative tasks can be accomplished by a single system provided that the end effector includes sufficient complexity to address proposed tasks and that the control system and its sensors are adequate. It was considered that, if a dexterous hand were developed with performance capabilities similar to the natural hand, that research could progress with a primary emphasis on control and sensing issues without being hampered by marginal performance of end effector machinery. Secondly, it appears that, from an experimental standpoint, an anthropomorphic configuration is desirable since it allows the human researcher to compare operations of the robot hand with operations of his own natural hand. It seems clear that, in a situation where almost nothing is really known regarding ideal configuration versus task objectives, the advantage of beginning with a system similar to the researcher's hand is unassailable. Thirdly, an anthropomorphic configuration has potential application as a slave element in a teleoperation system. In fact, a master system has been developed and the Utah/M.I.T. Dextrous Hand has been used successfully as a research device aimed at understanding issues which influence the performance of remote manipulation systems.

Unfortunately, it was not actually possible to achieve an exact anthropomorphic configuration as a result of packaging problems relating primarily to tendon routing. The final, quasi-anthropomorphic geometry, is as shown in Figure 3. The Figure illustrates the spatial positioning of the joints and the orientation of their axes.

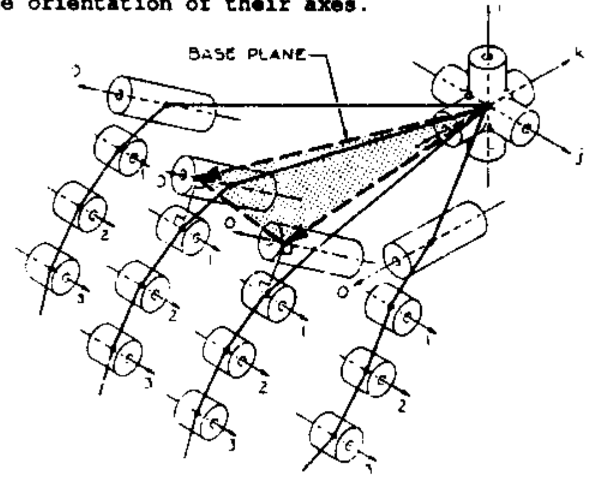

Figure 3. Configuration of the Dextrous Hand. The orientation of the axes of each joint are shown. Note that in the Version III hand, the 0 joints of the fingers are parallel to the base plane.

Specific deviations from anthropomorphic geometry include the following. First, the hand contains only three fingers and one thumb. The fourth, or "little finger," was eliminated to avoid complexity since the necessity of that finger could not be immediately shown. Second, the first two joints of each digit (the 0 and 1 joints as shown in Figure 3) were necessarily separated in order to allow tendons to be routed in a manner which would result in reliable operation. Note that for reasons of strength and fatigue life, flat belt tendons were selected and that routing such flat elements must be accomplished via a sequence of planar bends separated by axial twists (see Figure 4). If the joint included two intersecting axes, certain regions of the tendon would be subjected to undesirable two dimensional deformations.

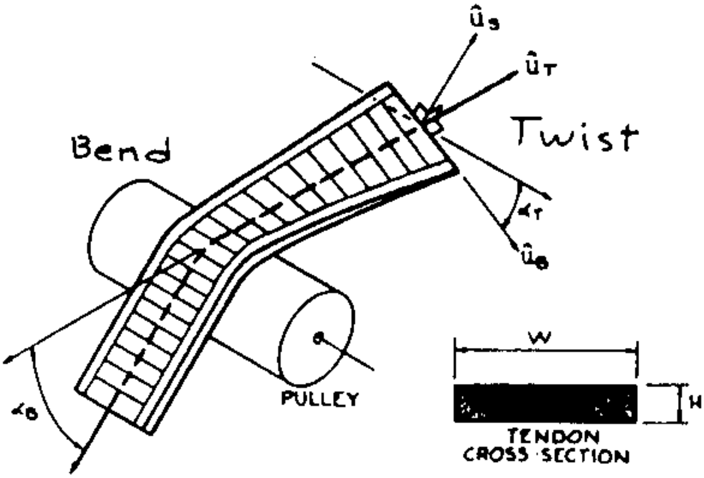

Figure 4. Tendons are routed throughout the system via a series of axial twists and bends over pulleys.

Third, the axis of the most proximal joint (the 0 joint) lies parallel to the base plane rather than tilted at 30 degrees as shown in Figure 3. Note that, in the natural joint, the 0 joint is essentially perpendicular to the base plane. Motivation for this change again relates primarily to tendon routing limitations. This alternative should not be considered negatively however, since it allows the fingers to achieve significant side-to-side excursions (0 joint mobility) when the 1 joint is flexed to the 90 degree position. These 0 joint motions actually improve mobility of fingertips when they oppose the thumb in contrast to the situation found in the natural hand where 0 joint mobility is almost null when the 1 joints are flexed to 90 degrees. Fourth, due to tendon routing difficulties, the base of the thumb was placed in the palmer section between the first and second fingers. This allowed tendons to be routed over the wrist and through the palm to the thumb in a reliable manner. Although, in a non-anthropomorphic configuration, the thumb does maintain sufficient 0 joint mobility to interact with all fingertips in a near natural manner. The existing configuration has been successfully operated in a teleoperation mode with suitable transformations made between motions of the master unit and finger positions of the DH. Fifth, and again because of tendon routing problems, the wrist joint is larger than desired. The enlargement, which causes some appearance problems, in fact provides additional space for the placement of 32 tendon tension sensors whose output is used for management of actuation systems and estimation of individual joint torques.

Comments Regarding Design of System Elements (DH)

From the outset our objective was to use the most advanced, yet practical technology so that performance could be maximized while achieving an acceptable level of reliability. The most important guiding principle used to manage the project has been the maintenance of balance. That is, understanding that design involves the simultaneous satisfaction of multiple conflicting objectives and that no specific area can be overemphasized at the expense of another. General success requires individual successes in all subareas, both conceptual and those relating to machine design. Efforts have included simultaneous work in eight areas of system design including: 1) structures; 2) internal sensors; 3) external sensors; 4) actuation systems; 5) remotizing systems; 6) covers; 7) communication networks; and, 8) computation systems.

Efforts have also focused on understanding the requirements for lower level control systems which provide system management functions. Essentially, the lower control system is responsible for producing a system which "does what it's told" with speed, strength and stability. Mid and higher control systems are now being investigated since the DH is now available for experimental activities.

In the design of all subsystems, simplicity has been a major objective. Every attempt has been made to eliminate unnecessary components and to design various system elements to be free from the need for continual calibration and maintenance. Closely associated with the avoidance of complexity has been the avoidance of precision. In all subsystems, close tolerances and the possibility of interfering components has been avoided.

In order to simplify control requirements actuator elements have been designed to exhibit desirable qualities as a result of intrinsic characteristics, rather than via attempts to modify actuator performance through the use of compensating feedback loops. In fact, the principle reason for the excellent passive performance of the DH is the intrinsic qualities of the pneumatic actuators. Finally, as previously mentioned, the DH was built in a modular fashion with a minimal number of different subsystems so that intensive efforts could be focused on understanding the behavior of a fewer number of elements rather than dealing with many different components.

Relocation of Actuators Outside the Hand. Of course the most desirable configuration for an anthropomorphic robotic hand is as shown in Figure 5. In this case all structures, actuators, sensors and covers are an integral part of one system with only simple connectors emerging for power, input commands, and to output sensor information. Unfortunately, a number of realities totally preclude such a possibility if the system is to include reasonable performance goals in relation to strength, speed, reliability, etc. The difficulty of generating a totally self-contained hand is demonstrated by natural systems which also require remotization of major muscles to the forearm area proximal to the hand.

The first retreat from the totally self-contained alternative was the relocation of actuators outside of DH to some convenient location. The relocation allows more volume within the DH to be used by structures, joints and sensors. This compromise also permits greater design flexibility for the actuators since they can be larger, with an emphasis on performance rather than restrictions in size, shape, weight and volume. Obviously, remotization of actuators also provides an opportunity for cost reduction.

The Use of Tendons. With the decision to remotize actuators, various methods for the transmission of mechanical energy from actuators to the hand were explored including hydraulic and pneumatic systems, mechanical linkages and flexible tensile tendons. Various trade-offs were exhaustively reviewed and the tendon approach was finally selected. Tendons provide a number of advantages including very low mass, and the possibility for stiff transmission of energy over complex pathways. Tendons require no bulky terminal energy transformation systems such as cylinders used in hydraulic systems and motors with transmissions used in electrical systems. It is also important to note that tendons possess an additional advantage in that their routes can be designed to augment the loading of structures as well as imposing intrinsic coupling between the motions of joints (i.e., tendonesis).

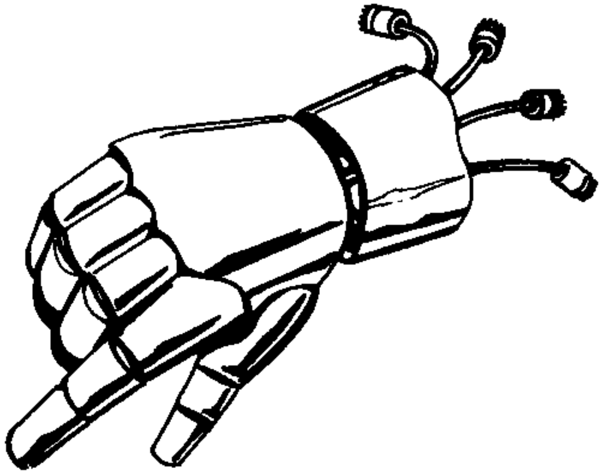

Figure 5. An ideal configuration for an anthropomorphic, dexterous, robotic hand.

Originally it was intended that high strength polymeric tendons be developed which were suitable to operate in lubricated and flexible channels similar to biological systems and to the Bowden type brake cables found in bicycles and motorcycles. Unfortunately, after exhaustive work with various tendon structures, with wet and dry lubricants, and with assorted bearing surfaces, the lubricated-channel alternative had to be disregarded in favor of pulley-based systems. Pulleys appear to be the only method which allows for both high strength and low friction operation. Pulleys unfortunately include problems with complexity and reliability, and they consume significant volumes within the DH, especially in critical regions surrounding joints.

In order that each joint could be individually controlled, subject to a minimal number of additional constraints, it was decided that two tendons would be routed over each of "n" joints (the "2n" versus the "n+1" configuration). Although this approach consumes additional volume and imposes higher levels of complexity within the DH, it was determined that the "2n" configuration was the more conservative approach and that since the system was to be a research tool, maximum levels of flexibility in operation should be maintained. If, however, it is determined at a later point that a lower number of tendons can be successfully utilized, that approach will be pursued.

It should be noted here that tendons have not been as much a problem with reliability as previously anticipated. The current tendon configuration, during conservative life tests as described in Reference 1, performed over 6 million cycles. It is our goal, with continued work, to achieve a life of 50 million cycles under the same experimental conditions.

Motivation for the Remotizing System. It was intended that, during initial experimental phases, the DH would be positioned in space by a manipulator arm configured such as shown in Figure 6. Due to the limited lifting capabilities of existing robots it was unlikely that the entire Dextrous Hand and its actuation package could be accurately and quickly moved around in space, so it was decided to construct a remotizing system to conduct tendons from the actuation package, which remains in a static position, to the DH, which is oriented by the robot arm. The remotizing system shown in Figure 6 includes 32 tendon pathways in four subsystems, each of which include series of longitudinal beams and rolling joints. The longitudinal beams support the compressive stresses imposed by tendons and the system of rolling joints permit motion of the remotizer without altering tendon path lengths. This configuration includes the side benefit that no torques are imposed on the remotizer joints as tendon tensions are varied. The remotizer is then a very passive system which allows the DH to be freely positioned in space while receiving substantial energy from the actuation package. The present configuration allows the DH to be positioned everywhere within a cube with three foot sides.

Selection of the Actuation System. The selection of the actuation approach was a difficult procedure since severe constraints exist with respect to both performance and configuration. Actuators must exhibit extraordinary performance under active control where the DH is primarily driven by commands from the central control system. The actuators must also attain performance of a different kind when passively manipulated by external loads such as those produced when the hand interacts with objects. Also important are factors relating to weight, size and geometry of individual actuators, especially considering that 32 individual actuators are required to operate the DH.

A number of alternative systems based on hydraulics, electrics and pneumatics were investigated. Unfortunately, when comprehensive sets of constraints were applied in the selection process, no commercial systems were found to be suitable for use. After a series of fundamental studies, simulations and experimental procedures, the pneumatic approach was selected for implementation. It should be emphasized here that the pneumatic approach was not selected because it represents the "easy" approach. The development of the present electro-pneumatic actuator was a complex process spanning three years, a number of complex computer simulations (up to 22nd order system models), four stages of prototyping, exhaustive experimentation and the development of a high performance, low level control system for supervision of the DH and its actuator package.

The selection of the pneumatic approach was not based on a single quantitative analysis, but: 1) a series of general performance requirements formulated during previous work; and, 2) broad knowledge of the characteristics of various real, physical systems. Attitudes regarding performance are very briefly reviewed in the following two comments:

Active Performance. Active performance relates to the operation of the Dextrous Hand in response to commands from the central controller. Since it is intended that the DH have strength approximately equivalent to the natural hand, the actuators were targeted to produce forces of 50 pounds. (The existing system actually generates only 25 pounds.) Since the DH is intended to engage in real time manipulation of objects of approximately hand size it was shown that the finger should operate at frequencies exceeding 10 Hz. In order to maintain acceptable positional accuracies during free space operation of the hand it was concluded that positional stiffnesses of the finger tip should be at least three pounds per inch without destabilizing finger operation.

Passive Performance. Developing a system with suitable passive performance has probably been the most difficult aspect of the project. Passive performance is the least well understood issue and requires the design of systems which emphasize subtlety and delicate behavior rather than brute strength. Simply put, our goal was the development of actuators which could maintain very low intrinsic output impedance at substantial frequencies (for example, above 25 Hz). To address this goal, the possibility of using pneumatic cylinders was explored since such systems can possess very low actuated mass (piston masses) while being capable of producing large forces. The compressability of the gas adds an intrinsic compliance to the system which functions at speeds in excess of those achievable via an active feedback system for behavior modification. The pneumatic cylinder approach however exhibits a primary disadvantage if the cylinder is driven by a flow control valve. With a fixed flow, the compressability of the gas allows the piston to act as a spring which oscillates against the masses of the finger. This produces a system with low damping and, because of delays, a high tendency towards unstable oscillation. It was therefore decided to develop a pneumatic valve with an integrated pressure control loop so that the driven pneumatic cylinder would operate as a force source thereby avoiding oscillation problems induced by gas compressability or by the compliance in remotizing structures and tendons. The development of a pressure control valve was a complex undertaking since the valve must possess speed and flow capacities sufficient to dominate system operation at frequencies exceeding cylinder resonant frequencies (i.e., whistle frequencies).

Ultimately the two stage jet pipe system shown in Figure 13 was selected. As shown in Figure 14 each joint is operated by two antagonist actuators which are controlled to produce desired output torques and co-contraction levels as further discussed in References 1, 3 and 4. The use of the actuators as antagonist pairs results in a very high bandwidth, low output impedance system which enables the fingers to easily interact with objects under force control with no tendency towards instability. When the system operates without contact in free space, instabilities occur if joint stiffnesses are adjusted above certain levels. In order to extend the level of stiffness achievable during non-contact motions, dampers were added to each actuator as shown in Figure 12. Also, as discussed in Reference 1, stiffness can be automatically increased upon contact with objects as indicated by tactile sensors (proximal stiffening). (Note that without dampers the actuation package is capable of operating the finger's number 3 joint at frequencies in excess of 60 Hz). The resulting system is an excellent compromise which permits free operation at speeds up to 10 Hz while simultaneously allowing for high compliance interactions with solid objects. If desired, actuator performance can be further enhanced by expanding controller function based on information from tendon tension sensors in the DH.

Joint Angle Sensing Systems. In order to obtain accurate joint angle information for control purposes, it was decided that the sensors be located at the joints within the fingers themselves. An alternate choice would be to monitor tendon deflections at the actuators and determine joint angles by back computation. This approach would be subject to excessive errors and would drift should lengths of tendons vary or their terminations slip. A number of alternative methods for angle sensing were investigated including potentiometric, capacitive, optical and magnetic. Potentiometric alternatives seemed straightforward, but introduced problems with packaging, susceptability to intrusion by contaminants and reliability. Optical and capacitive systems, based on discrete measurement techniques, were complex, had limited resolution, were fragile and could be unreliable in certain circumstances. Also, due to their digital nature, the acquisition of angular velocity information introduced significant computational time delays which are undesirable if such information is to be used within servo loops. Finally, a magnetic approach using Hall effect sensors was explored and has proven to be very desirable. This system is reliable, proportional, and compact. System elements are totally encapsulated so that intrusion of dirt and other contaminants is not possible. Noise levels produced by the system are low and signals are smooth enough for direct differentiation to provide velocity information. The present system exhibits one disadvantage in that operation in the presence of strong magnetic fields can produce errors. Consequently, present efforts are aimed at a dual Hall effect system which configures transistors in a bridge in order to desensitize the system to external magnetic fields.

Tendon Tension Sensors. For a number of reasons it appeared important that tendon tensions be monitored in order to provide information regarding the torque imposed on individual joints as well as the possibility of providing information to the controller for actuator compensation. Ideally, tendon tension sensors would be installed distally at the insertion of tendons at each joint and proximally at the output of each actuator. Such a dual system could provide valuable information for the compensation of elastic and frictional characteristics of the tendons. However, due to packaging problems and issues related to complexity, a compromise system was selected. As shown in Figure 6, 7, 10 and 32, tendon tension sensors are located in the wrist. Each sensor uses a semiconductor strain gauge bridge to monitor beam deflection which is proportional to tendon tension.

DESCRIPTION OF THE UTAH/M.I.T. DEXTROUS HAND (DH)

Specific Subsystems of the DH

Systems for Spatially Positioning the Hand. Figure 6A is a photograph depicting a typical experimental configuration for the DH. Figure 6B is a line drawing included to indicate system subelements. The Dextrous Hand (A) is holding an object for assembly. The Hand connects to the wrist (B), which attaches to the remotizer (C). The remotizer carries tendons from the actuator package (D) which is maintained in a static position by an external mounting. The hand is oriented in space by a PUMA robot (F). To generate additional experimental flexibility components in the manipulation field can be oriented with respect to the hand assembly by a powered, moving table (E). Degrees of freedom of the table are: translation X, translation Y and rotation q.

DH Structures and Joints. The Version II Hand, shown in Figure 7, is essentially a prototype system used for investigating issues in packaging, joint angle sensing, tendon routing, electrical interconnections and hand geometry.

The Version II Hand was a predecessor to the Version III system shown in Figures 2A and 2B. Both hands include three fingers and one thumb, with each digit containing four joints. All three distal joints (1, 2 and 3) of the fingers and the thumb are capable of excursions from 0 - 95 degrees. The proximal, or 0 joint, functions differently depending on whether it is the base for a thumb or a finger. The 0 joint in the finger is capable of ± 25 degree motions and the 0 joint in the thumb is capable of ± 45 degree motions. The orientation of the base, or 0 joints, in the finger are illustrated in Figure 3. All elements of the Hand are machined 7075 aluminum with dual precision ball bearings installed at each joint. The hand and wrist include 184 low friction pulleys for the purpose of tendon routing. Joint angle sensors are located at the side of each joint, with their electrical connections running along lateral slots located at the neutral axis of each finger. On the opposite side of each finger are similar slots for the purpose of routing additional communication lines for external sensors.

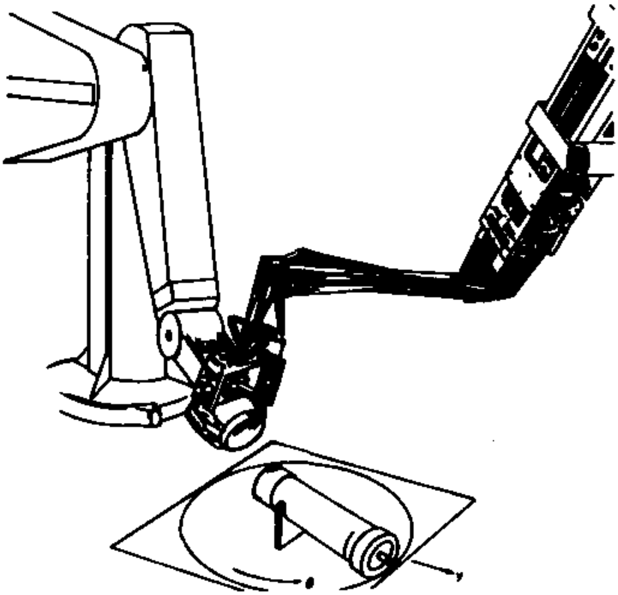

Figure 6. Line drawing illustrating components in Hand, Remotizer, Actuator Package, PUMA and positioning table.

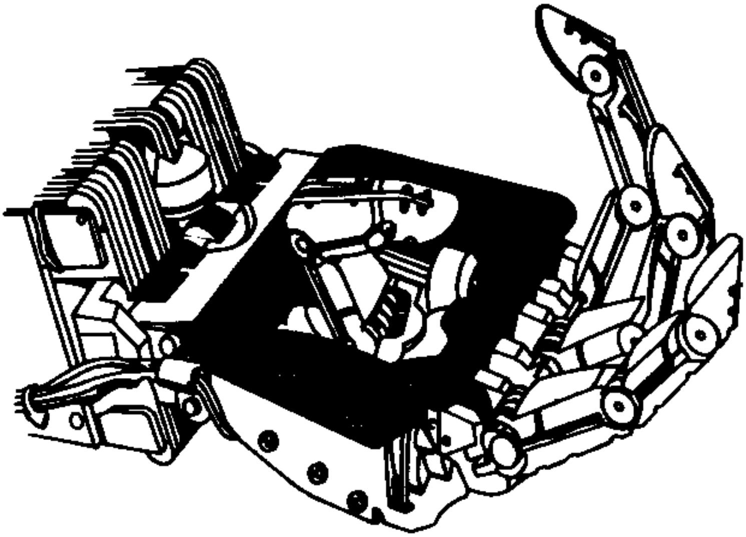

Figure 7. The Version II Utah/M.I.T. Dextrous Hand.

Figure 8 illustrates the configuration of the wrist, which includes two perpendicular axes, implemented by a crossed yoke mechanism. A third orthogonal axis is made possible by axial rotation of the remotizer compression rods.

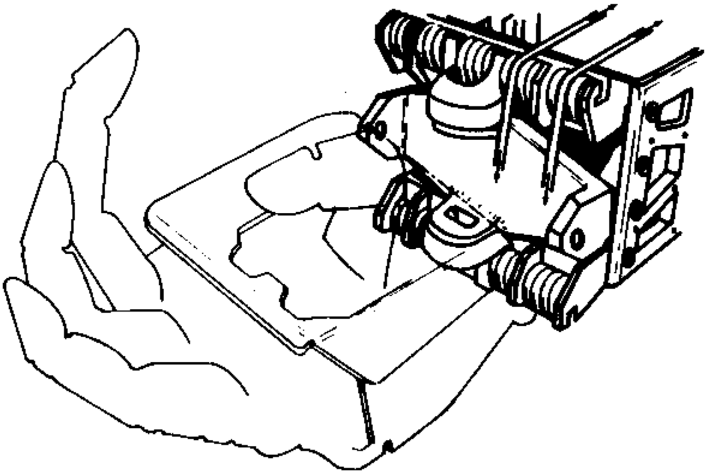

Figure 8. Configuration of the wrist in the Utah/M.I.T. Dextrous Hand.

Permissible deflections for the wrist are ± 45 degrees for wrist flexion/extension; ± 15 degrees for wrist abduction/aduction; and ± 135 degrees for wrist rotation. All three axes intersect at a central point in order to simplify inversion computations. Again, the wrist includes 32 tendon tension sensors located immediately behind the outer row of pulleys marked (A) in Figure 8.

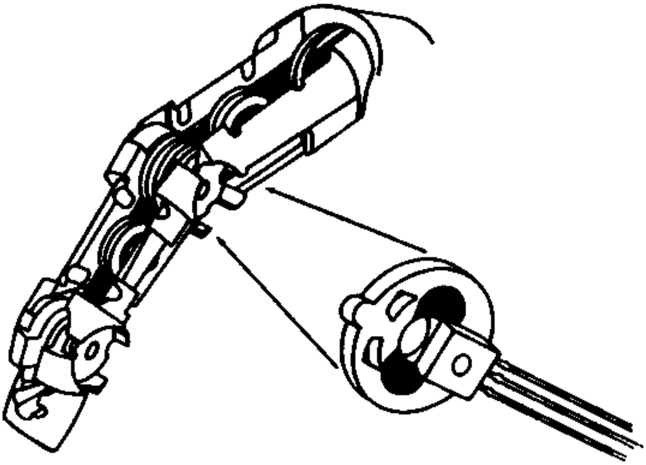

Figure 9. Configuration of the Hall effect sensor used to measure angular deflection of the joints.

<u>Internal Sensors - Joint Angle and Tendon Tension</u>. Figure 9 illustrates that each joint contains a sensor to measure angular deflection. Part A in the figure is a magnetically sensitive Hall effect device located in the proximal link, while part B, attached to the distal

link, includes two cobalt samarium magnets operating in a dipole configuration. As the Hall effect device sweeps the magnetic field it produces an output current which corresponds to angular deflections between 0 and 95 degrees with a linearity within 5%. The Hall effect sensor system includes a number of advantages such as continuous output signal, high bandwidth operation, low friction, no mechanical contact, long life, and tolerance to surrounding contaminants.

Figure 10 illustrates one of the 32 tendon tension sensing systems located in the wrist. The pulley is positioned in order to perturb the path of the tendon such that tendon tension imposes a load on the cantelevered beam. A semiconductor strain gauge bridge detects beam strain and provides a linear output for tendon tensions from 0 to 30 pounds. Supporting electronics for both the angle and tendon tension sensors are located in the low level control system (LLCS) shown in Figure 17.

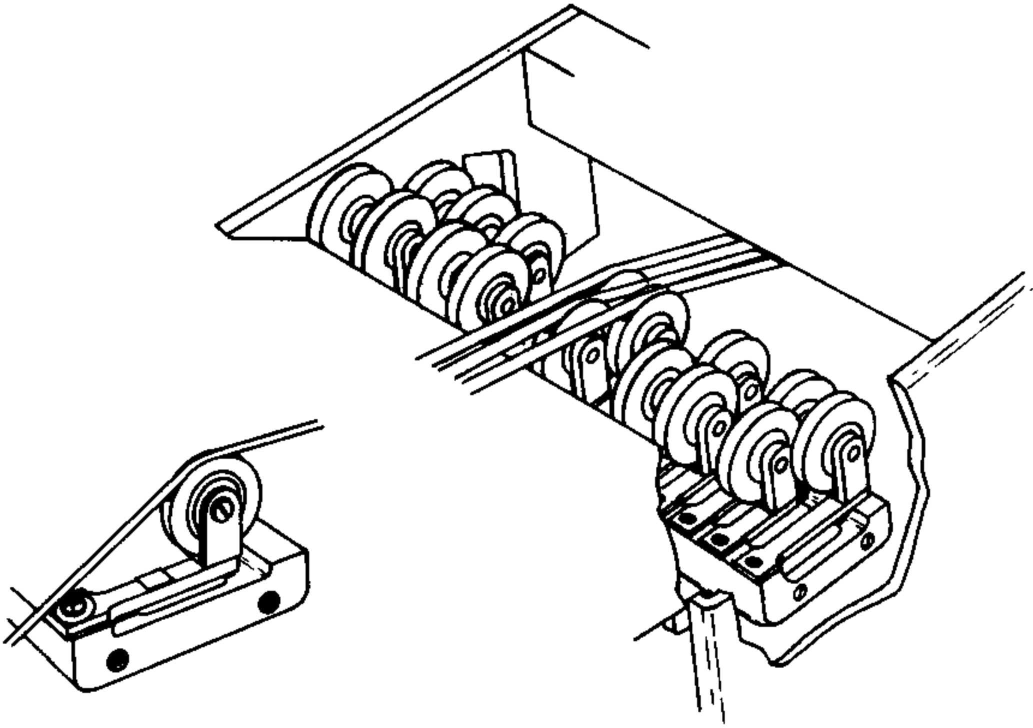

Figure 10. 32 tendon tension sensors located in the wrist of the Dextrous Hand.

External Sensors and Covering Systems. Figure 11 illustrates five removable segments which occupy void spaces in the finger structure. Each segment is approximately 1/10th of an inch thick and can be injection molded from either rigid or flexible materials. Depending on the particular experiment, selected sections of these rather disposable elements can be machined away to allow space for tactile sensing transducers; for example, detectors to sense direct contact, normal pressure, shear stress, temperature, etc. This approach maintains experimental flexibility so that various methodologies can be tried without committing to a specific geometry. Note that the entire system can be operated with the segments exposed or covered with a flexible glove to isolate internal components from unwanted environmental contaminants. As mentioned in the previous section, communication with external sensors will be provided via conduits which run along lateral slots in the fingers.

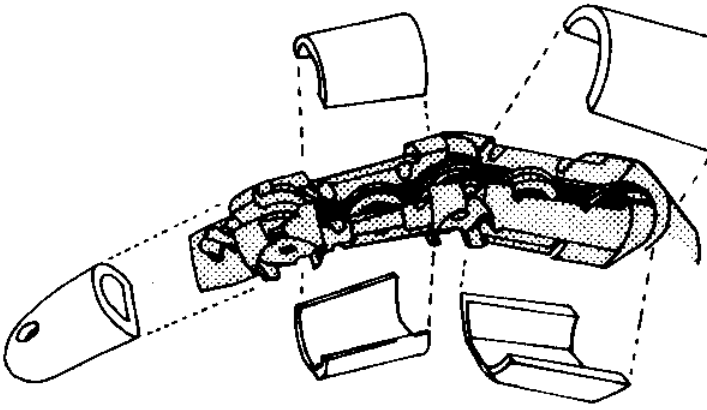

Figure 11. Five removable segments intended to house tactile sensing systems.

<u>Actuators and the Actuator Package</u>. As shown in Figure 12A, each joint receives two tendons, driven by actuators which consist of pressure controlling valves and low stiction cylinders. Each valve receives an electrical signal which commands a cylinder pressure. As described in References 1, 3, and 4, the signals are modulated so that the torque (differential tendon tension) can be varied simultaneously with co-contraction (sum of tendon tensions) of the tendons. The actual module used in the DH is shown in the line drawing of Figure 12B.

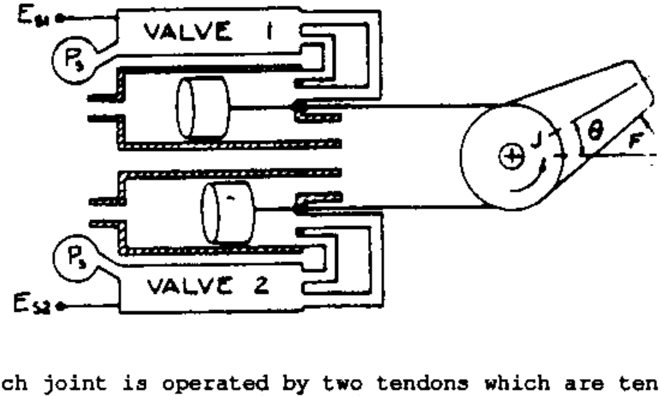

Figure 12A. Each joint is operated by two tendons which are tensioned by actuators consisting of low stiction cylinders and pressure controlling valves.

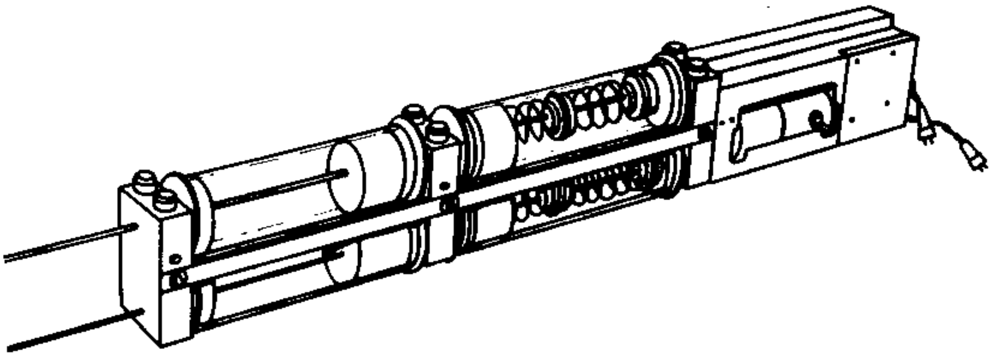

Figure 12B. A dual actuator module which includes valves, cylinders and dampers.

Each module contains two valves (A), two actuator cylinders (B), and two adjustable pneumatic dampers (C). Also, spring tensioning systems (B) are included within cylinders (B) to maintain tendon tensions when the system is unpressurized. Residual tension during inactive periods prevents misalignment of tendons which can, during operation, seriously impair tendon life. The cylinders for actuation and damping are constructed of glass and close fitting graphite pistons. This configuration, operating at low pressures (50-100 psi), permits the elimination of seals, thereby minimizing stiction effects. Dampers, which operate in a pressurized mode in order to increase damping and minimize compressibility effects, are adjustable in order to allow a variable trade-off to be made between high frequency operation and stability during free-space operation. The actuator modules are compactly placed in a 4.25 x 4.25 x 24 inch rectangular assembly containing air manifolds and weighing 20 pounds, as shown in Figure 6. Electrical servo amplifiers for the system are contained in the low level control system as shown in Figure 19.

As further described in References 1, 2, and 3 and in Figure 13 each actuator consists of a two-stage jet pipe valve. The first stage is electrically driven to provide a pressure signal to a second stage such that an increase in primary stage pressure drives the second stage upper diaphragm downward to deflect the jet pipe and increase piston cavity pressure. Piston cavity pressure is fed back to a lower diaphragm on the second stage jet pipe which antagonizes the upper diaphragm such that the valve behaves approximately as a pressure source modulated by input electrical current. This valve assembly is very fast. In fact, as shown in Reference 2 it can drive the pressure in the cylinder cavities, at maximum volume, with a flat frequency response up to 21 Hz.

Tendons, Remotizer and Disconnection System. As shown in Figure 15A, the tendons in the DH are composite structures which consist of Dacron fibers woven around multiple longitudinal elements of Kevlar. The Dacron outer sheath serves to align and protect the internal load bearing Kevlar fibers. Tendons are routed throughout the hand via sequences of pulley-induced bends and axial twists as shown in Figure 4. Pulleys used have diameters of .95, .72, .49 and .38 inches with centrally domed contact surfaces in order to provide tendons with self-aligning properties.

The fixation of the ends of tendons is critical to system reliability and a number of termination approaches have been explored. Three successful termination methodologies are: 1) the permanent clip (C) as shown in Figure 14; 2) the age-old knot; and, 3) the overwrapped friction lock (A) in Figure 14 and Figure 15. The friction lock permits non-destructive and reversible adjustment of the tendons in order to initially locate fingers with respect to actuating piston positions. After adjustment, the tendon is locked via a supplemental machine screw. The molded section (A) shown in Figures 14 and 15 allows for easy, quick disconnect between actuator rods and the tendons.

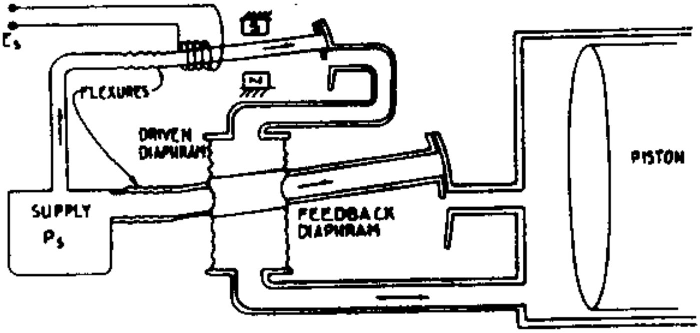

Figure 13. Schematic diagram showing the function of the two-stage electropneumatic, pressure controlling, jet pipe valve.

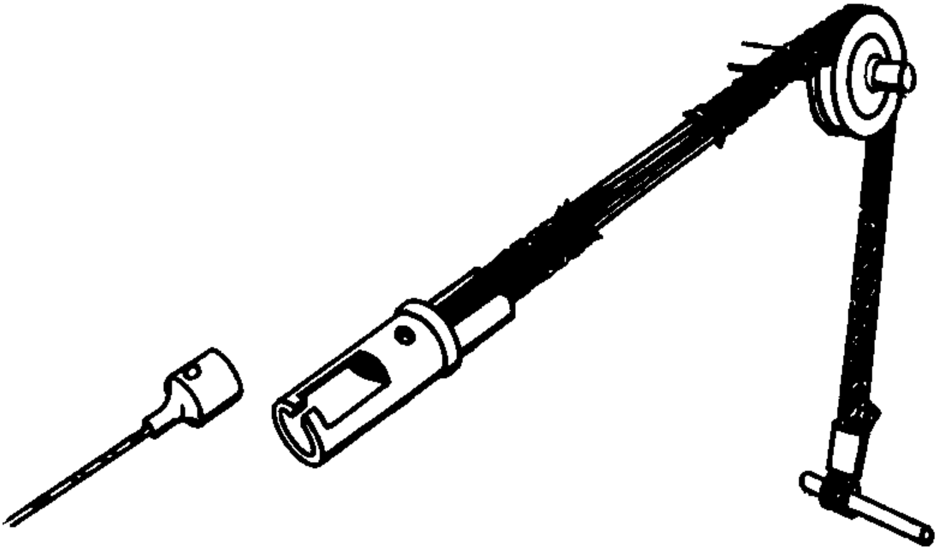

Figure 14. Tendons which are constructed of Dacron and Kevlar are routed over pulleys (B) and terminate either permanently as shown at point (C) or in a reversible manner shown at point (A).

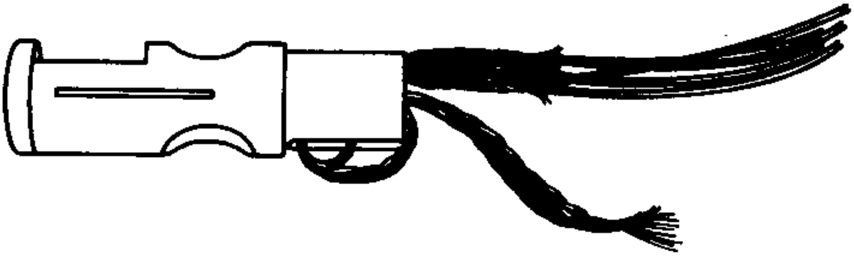

Figure 15. Injection molded component used to terminate tendons and provide quick disconnect capability.

Figure 6 illustrates the total remotizer which consists of four subsystems which each include: 1) rolling joints constructed of injection molded composite materials into which gears have been insert molded. The orientation of the two halves of the joint is maintained by coupling members and the gears which constrain the joint to roll in a manner which maintains tendon length regardless of angular deflection of the joint. Each rolling joint includes 16 pulleys for tendon routing which means that the entire remotizer utilizes 288 pulleys from the actuation package up to tendon tension sensors in the wrist; 2) compression rods which counteract tendon tensions and allow axial rotation in order to add degrees of freedom to the remotizer assembly; and, 3) eight tendons which run between the actuator package and the tendon. Note that the four subsystems must be flexibly tethered in such a way to maintain their general coordination while permitting simultaneous relative motion as the remotizing system is bent and twisted.

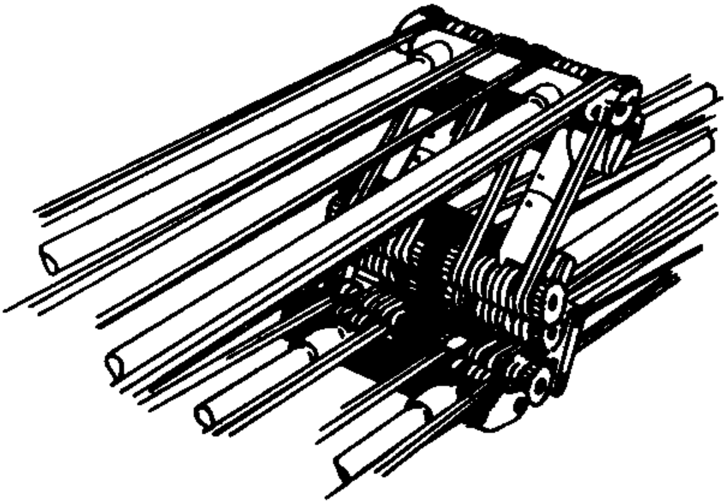

Figure 16A. Oblique view of the central joint (elbow) of the remotizer. Note four subassemblies operate in parallel to conduct 32 tendons from actuator to hand.

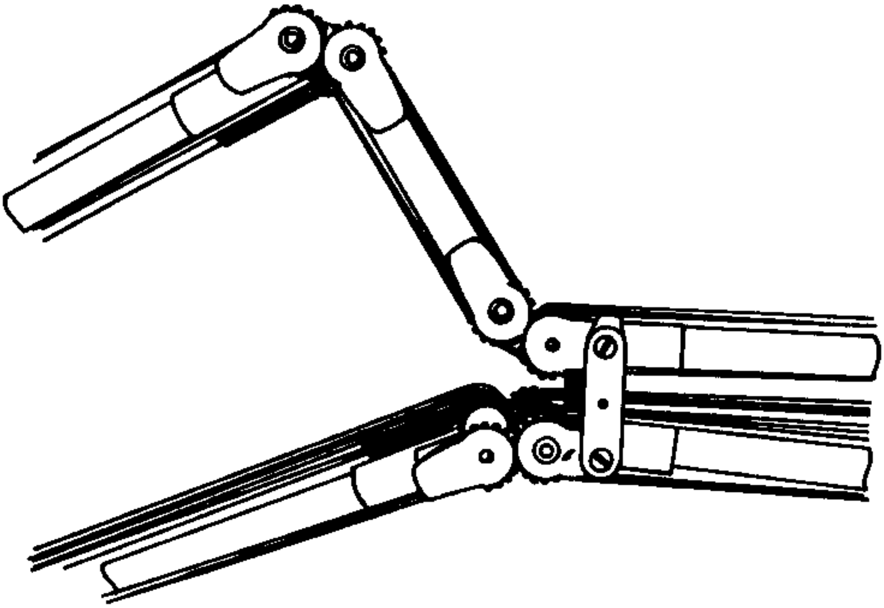

Figure 16B. Side view of the central remotizer joint showing rolling action of remotizer subsection joints.

Analog and Digital Computation System. Figure 17 shows the Low Level Control System (LLCS) which can stand alone to insure that all subsystems are functioning or receive complex analog inputs from higher digital control systems. The analog system is advantageous since it executes a of number servoing functions smoothly and with high bandwidth, thereby reducing computational requirements on the digital control system.

The LLCS includes 16 variable-loop-gain position servos to operate finger joints and 32 variable-loop-gain tension servos to modulate actuator behavior such that tendon tensions can be closely controlled. Amplification and signal conditioning circuitry is also included to: 1) provide current sources for driving pneumatic valves; 2) drive and monitor tendon tension sensors and drive and monitor joint angle sensors. For reasons of flexibility, system inputs include: 1) 16 inputs for control of angular position; 2) 32 inputs for control of desired tendon tension; 3) 16 inputs to vary position servo loop gain; and 4) 32 inputs to vary tendon tension servo loop gain. Also, a number of auxiliary inputs are available to control damping, co-contraction levels and to allow direct control of servo valve currents. The face of each of the 16 subsystems of the LLCS includes 13, proportional, multicolor light emitting diodes (LED) for the purpose of diagnostically displaying important system parameters. The console also includes 16 potentiometer inputs for the purpose of manually adjusting joint angles. The LLCS also provides analog outputs of all sensor signals generated within the hand.

The digital control system, which has been previously described in Reference 5, will only be briefly reviewed here. The system consists of five Motorola 68000 microprocessors, a Multibus card cage, 40 channels of digital to analog conversion and 320 channels of analog to digital conversion. All software development was accomplished on a VAX 11/750, utilizing the programming language "C." The system includes four finger controllers each doing all computations required for four joints. The system also includes a master controller which: 1) manages all system gains; 2) deploys data; 3) interprets primitive commands; and, 4) monitors the system for errors. High speed data taking and plotting have been accomplished via a PDP 11/44.

System Performance

Providing a comprehensive description of the performance of various system elements is a lengthy task and is beyond the goals of this document. Additional information regarding performance can be obtained via References 1, 2 and 3. However, in order to provide the general impression regarding the performance of an individual finger, Figures 18A and 18B have been included. Figure 18A illustrates the combined step response of joints 0, 1, 2, and 3. The finger in this circumstance executes a grabbing motion, combined with a side sweep due to the 0 joint motion. The approximate rise time can be seen to be approximately 70 milliseconds.

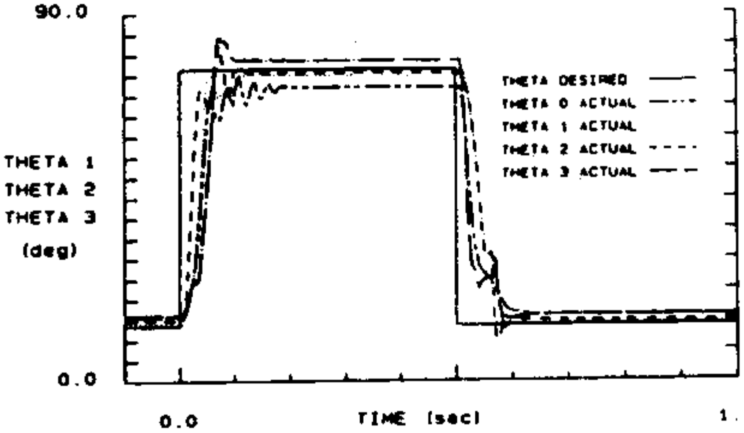

Figure 17A. A complete Version I finger as outlined in Reference 3.

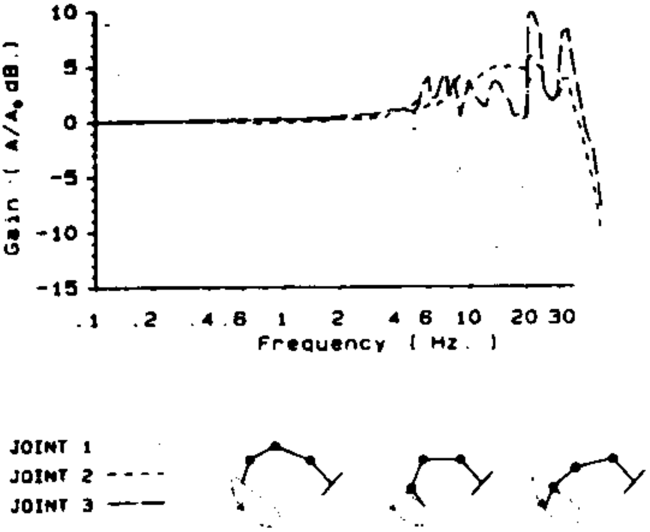

Figure 17B. Gain plot for an individual finger executing a forward running motion as described in Reference 3.

Figure 17B illustrates the gain ratio for a finger executing a motion similar to forward running. Joint 1 is driven by the function A_0 sin (wt), while joint 2 and 3 are driven by the function A_0 sin (wt - π / 2) where A_0 is 60° on joints 1, 2 and 3 and 45° on joint 0 and w is the frequency. Note that this system exhibits a relatively flat response up to frequencies of approximately 8 Hz. The output is less well behaved but substantial at frequencies up to 30 Hz. The somewhat unwieldy behavior in the range from 10 to 30 Hz is a result of coupling motions between degrees of freedom which can be reduced via compensation in the controller if desired.

CONCLUSIONS

Previous sections have reviewed the development and characteristics of the Utah/M.I.T. Dextrous Hand which promises to be an effective tool for the investigation of issues in machine manipulation. The DH, together with its wrist and remotizer includes over 25 degrees of freedom thereby providing a rich environment for experimentation. The system exhibits high electromechanical performance in both active and passive operation so that both high and low speed manipulation procedures can be explored. Finally, the system has been designed for maximum reliability in order to allow experiments of significant complexity to proceed without interruptions due to machine failures.

Now that the hand and its low level control system are in existence, a number of experiments will begin aimed at understanding higher control system issues such as task definition, manipulation strategies, grasping functions and the collection and utilization of comprehensive sensory information. With equal importance, work will also continue towards understanding issues related to machinery so that future manipulation systems can be designed to produce high performance, with substantial reliability and at a reasonable cost.

The paper also reviewed a number of important issues which impinge on the development of future systems. Of particular importance to future work will be: 1) understanding appropriate goals; 2) developing approaches for the effective management of directed research efforts; and, 3) achieving a balanced focus on both theoretical and practical issues.

ACKNOWLEDGEMENTS

This work was conducted in collaboration with the Artificial Intelligence Laboratory at the Massachussetts Institute of Technology. Work was supported via Defense Advanced Research Projects Agency, contract F33615-82-K-5125; System Development Foundation; and the Office of Naval Research, contract N00014-82-K-0367.

The authors also thank Professor John M. Hollerbach of the Massachusetts Institute of Technology for his efforts in this project. Appreciation is also expressed to M.I.T. Research Assistants, Mr. David Siegel and Mr. Sundar Narasimhan.

REFERENCES

1. S.C. Jacobsen, J.E. Wood, D.F. Knutti, K.B. Biggers and E.K. Iversen, "The Utah/M.I.T. Dextrous Hand: Work in Progress," *The International Journal of Robotics Research,* vol. 4, no. 3, pp. 21-50 (Winter 1984).

2. S.C.Jacobsen, D.F. Knutti, K.B. Biggers, E.K. Iversen and J.E. Wood, "An Electropneumatic Actuation System for the Utah/M.I.T. Dextrous

Hand," *Proceedings of the Vth CISM-IFToMM (International Federation for the Theory of Machines and Mechanisms) Symposium*, Udine, Italy, June 1984.

3. S.C. Jacobsen, J.E. Wood, D.F. Knutti, K.B. Biggers and E.K. Iversen, "The Version I Utah/M.I.T. Dextrous Hand," *Proceedings of the Second International Symposium of Robotics Research*, MIT Press, 1984.

4. K.B. Biggers, S.C. Jacobsen, G.E. Gerpheide, "Low Level Control of the Utah/MIT Dextrous Hand," accepted to the 1986 IEEE International Conference on Robotics and Automation, April 14-17, 1986 in San Francisco, California.

5. D.M. Siegel, I. Garabieta and J.M. Hollerbach, "A Capacitive Tactile Sensor," *SPIE Conference on Intelligent Robots and Computer Vision*, Cambridge, Massachusetts, September 1985.

6. D.M. Siegel, S. Narasimhan, J.M. Hollerbach, G.E. Gerpheide and D.J. Kriegman, "The Computational Architecture of the Utah-MIT Dextrous Hand," *IEEE International Conference on Robotics*, St. Louis, Missouri, March 1985.

7. J.M. Hollerbach and T. Flash, "Dynamic Interactions Between Limb Segments During Planar Arm Movements," *Biol. Cyber.* 44(1):67-77, 1982.

8. R.P. Paul and C.N. Stevenson, "Kinematics of Robot Wrists," *International Journal of Robotics Research* vol. 2, no. 1, 1983.

9. M.H. Raibert and J.J. Craig, "Hybrid Position/Force Control of Manipulators," *Transactions of ASME J. Dyn. Sys., Measurement, Contr.*, 103(2):126-133, 1981.

10. J.K. Salisbury, "Kinematic and Force Analysis of Articulated Hands," *Ph.D. Thesis, Department of Mechanical Engineering, Stanford University*, May 1982.

11. B.W. Andersen, "The Analysis and Design of Pneumatic Systems," Robert E. Kreiger Publishing Company, New York, 1976.

12. S.A. Kapanji, "The Physiology of Joints: Upper Limb," E & S Liverstone, London, 1970.

13. H. Hanafusa and H. Asada, "Stable Prehension by a Robot Hand with Elastic Fingers," *Proceedings 7th Int. Industr. Robots*, pp. 361-368, Tokyo, Japan, 1977.

14. H. Kobayashi, "Control and Geometrical Considerations for an Articulated Robot Hand," *The International Journal of Robotics Research*, vol. 4, no. 1, pp. 3-12, MIT Press, 1985.

15. M.H. Horn, P.G. Riewald and C.H. Zweben, "Strength and Durability Characteristics of Ropes and Cables from Kelvar Aramid Fibers," *Proceedings of Oceans '77 Conference*, pp. 24E-1 - 24E-10, 1977.

16. A.K. Bejczy, "Smart Sensors for Smart Hands," *AIAA/NASA Conference on Smart Sensors*, Hampton, Virginia, Nov. 14-16, 1978.

17. L.D. Harmon, "Touch-sensing Technology: A Review," *Society for Manufacturing Engineers*, MSR80-03, Dearborn, Michigan, 1980.

18. J.M. Hollerbach, "Dynamic Scaling of Manipulator Trajectories," *AIM-700*, Artificial Intelligence Laboratory, M.I.T., Cambridge, Massachusetts, 1983.

Design and Performance of a Cable Controlled Parallel Link Manipulator

Samuel E. Landsberger
Research Assistant
Department of Mechanical Engineering
Massachusetts Institute of Technology

1. Introduction

Manipulator geometries may be classified as serial or parallel, or hybrids thereof. The serial linkage is an open kinematic chain, with successive links connected by rotary or prismatic joints. End effector positioning is achieved by controlling the various joint angles/extensions via actuators normally located at the joints. The resulting mechanism, an emulation of the human arm, exhibits good freedom of movement and the ability to work around obstacles in its workspace. Its drawbacks are poor strength/stiffness/accuracy - to - weight ratios, consequent low-frequency vibrational modes, and complex kinematic and dynamic control algorithms. A parallel linkage can overcome these deficiencies, with a tradeoff in dexterity and obstacle avoidance capability. Fig. 1 illustrates the two configurations schematically.

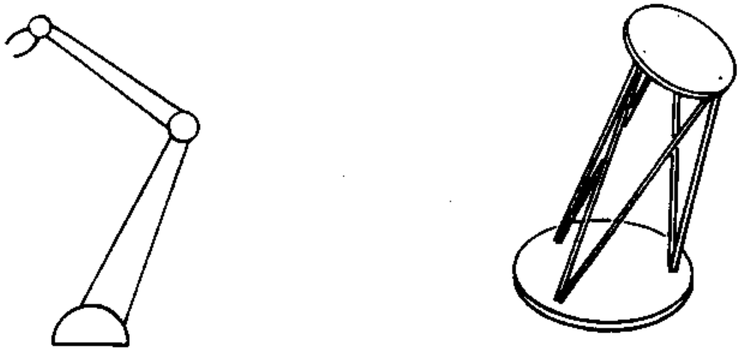

Figure 1. Serial and Parallel linkages

The parallel geometry is a closed kinematic chain wherein multiple links join end effector to base. The geometry of the manipulator which is the topic of this paper resembles that of a "Stewart Platform"[1] wherein six pivoting links of adjustable length specify the position and orientation of an upper platform. In the present design the six links are cables spooled by motors resting at the base. Preload in the system is maintained by a central telescopic compression member, the spine.

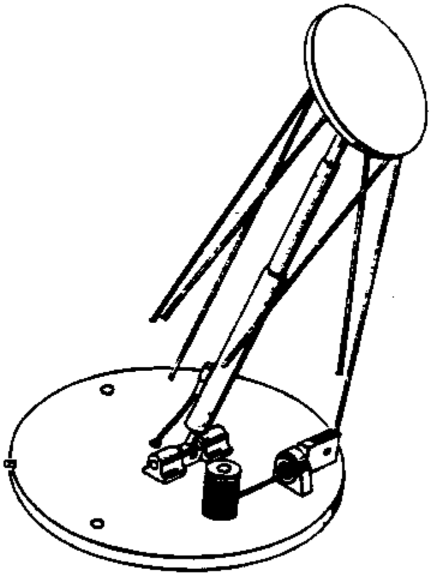

Figure 2. Six-degree-of-freedom cable manipulator

2. Advantages and Drawbacks of the Cable Linkage

2.1 Advantages

(i) Very high strength and stiffness-to-weight ratios result because actuating links bear no moment loads, but are instead simple tension-compression members. In a serial arm, by contrast, all links are subject to cantilevered loading.

(ii) Because the actuators act in parallel rather than series to position the end effector, the force and moment capacity of the manipulator is much higher than that of the individual servomotors. A serial link arm is no stronger than its weakest link/actuator, whereas the parallel arm capacity is roughly the sum of the individual actuator strengths.

(iii) Manipulator inertia is minimal in the parallel cable configuration because bulky links and motors are not being waved about in space - only the end effector. This results in economy of power and superior dynamic performance. The manipulator bandwidth may be approximated as:

$$\omega \approx \sqrt{\frac{stiffness(k)}{mass(m)}}$$

where we have noted the high stiffness of the configuration. The low dynamic inertia is of particular interest in space (and small ROV) applications where "tail wagging the dog" (manipulator inertial reactions moving the vehicle) phenomena must be considered. Furthermore, the inertial properties of the cable system are approximately invariant with respect to manipulator configuration, so that accounting for them is easy. Cable and spine mass are so low that in effect the dynamics reduce to that of a single rigid body. This is in contrast to the inertia matrices of the serial type linkage, whose highly nonlinear configuration-dependence render dynamic compensation a slow and difficult task.

(iv) Friction is minimal, and backlash-endplay zero due to the direct, preloaded connection between actuators and the end effector. Thus very high accuracy and smooth motion are realized from simple components: six motors, idler pulleys/cable conduits and a compression member. High resolution force control is enhanced by the low-friction cable drive, in which the only source of "stiction" are the motor shaft seals.

(v) The design is simple and modular, hence more reliable and cheaper to manufacture than the close tolerance joints, links, gearboxes and lightweight motors found on conventional arms. The arrangement permits static mounting of the actuators beneath the base or in a separate module, connected to the base via teflon-coated cable conduits. This results in a lightweight, better-protected manipulator. In the event of a collision or other mishap, expensive links and actuators aren't "out on a limb". Cable and spine replacement is easy and inexpensive.

(vi) The cable links, stored on spools, collapse to a small rest size while permitting great extension of the active manipulator. Currently, 60" of 2000 lb-test cable is stored on a spool 2" in diameter by 3" high, allowing a 5' reach. Different spools, upper platforms and telescopic spines may easily be substituted, without great expense, allowing radical configuration changes of the manipulator to suit a variety of tasks, from localized assembly work requiring great precision and compliance control, to pick-and-place work where speed and work envelope are paramount. Increasing the spine preload improves disturbance load rejection in much the same way as tightening one's muscles stiffens the body's response to external forces, while decreasing it "relaxes" the arm, and permits safe operation of the manipulator in fragile environments. The parallel link cable mechanism with opposing telescopic spine, wherein no rigid link exists between base and end effector, proves to be a simple, compact and versatile manipulator whose workspace, load capacity and compliance characteristics may be varied to suit a wide variety of tasks.

(vii) The inverse kinematics and dynamics computation is orders of magnitude easier for the parallel than serial geometry, allowing real-time control with an inexpensive microprocessor.

2.2 Disadvantages

(i) The parallel configuration makes obstacle avoidance more difficult than in the case of a serial-type linkage since the manipulator physically occupies a greater percentage of its workspace, and cannot "articulate" around objects in its path.

(ii) The outward normal to the upper platform cannot incline more than 90° to the direction in which the spine points at that position.

(iii) Workspace bounds are difficult to express in terms of link lengths, whereas normally the stops of a serial arm are specified directly in terms of joint rotation limit angles.

(iv) Singular points of the parallel kinematics result in the introduction of unwanted degrees of freedom (floppiness), whereas serial singularities correspond to a safer "lock up" mode where some freedom is lost.

3. Hardware

The prototype which has been built is intended for undersea use on MIT's ROV, *SEAGRANT I.* The design incorporates Planet ball-piston hydraulic motors with integral servovalves and potentiometers, outputing 240 in-lb of torque @ 3000 psi. Keyed directly to the motor shafts are ten-turn spools which accumulate the 1/8" stainless cable. These pass over swiveling idler pulleys before attaching to the upper plate via universal joints. The central compression member is gimbaled at the base and upper platform. Currently, a 9" stroke cylinder is in use, to be replaced by a 4' stroke multistage member when funds become available. This will give the manipulator a 10' diameter workspace. The 22" diameter base, 7" upper plate, cable spools and idler pulleys are made of Delrin, a strong, low water-absorbtion thermoplastic yeilding a 95 lb manipulator of which 5 lb is dynamic mass.

4. Control

The six cable servoes operate independently under simple analog proportional position control. A desired trajectory is computed in spherical-eulerian coordinates as a sequence of arc increments for which cable lengths are evaluated.

Although description of a smooth arc requires coordinated motions of all six cables, the lengths may be independently varied without interference. Spine preload is controlled via a pressure servo to modulate end effector apparent compliance independently of the cable servo gains. With low preload the manipulator's load capacity is restricted so that safe interaction with a fragile or obstacle-ridden environment is possible, whereas with a 1200 lb preload loads of that magnitude may be supported. The preload parameter also influences the dynamic response. A moderate spine preload can often effectively hide dynamic effects as well as disturbance loading from the environment: a 1 kg mass accelerating at 3 g's induces an inertial load of about 5 lb on the upper plate - noise compared with nominal cable tensions of 100 lb!

5. Kinematics

The kinematic equations relating coordinates natural to the manipulator to natural observer reference coords are of inverse complexity in the case of parallel and serial geometries. For the Stewart Platform, determination of the link lengths corresponding to a desired end effector position involves little more than six applications of the pythagorean theorem. Once the position of the upper platform is specified, the positions of the six cable tie points may be determined using a translation vector and rotation matrix (or one 4-D transformation matrix[2] ,[3]). Since the base tie points are approximately fixed, vector subtraction of a base point from its corresponding upper tie point yeilds the desired cable vector, whose magnitude is the position input to the cable motor servo. In the serial case, the inverse kinematics cannot in general be solved in closed form, and are plagued by singular points. On the other hand, the forward kinematics relating a given set of joint angles to endpoint location are straightforward for the serial arm, but difficult in the parallel case, and also plagued by multiple roots or "assembly modes".[4] Fortunately the latter are not normally required in control applications, and where necessary, as in position error estimation, the inverse Jacobian suffices.

6. Workspace

The working bounds are conveniently expressed as follows:

(a) Inner and outer radial stops. The upper platform must lie between spherical shells of radii corresponding to the minimum and maximum extension of the spine. Currently the range is from 19" to 28". The multistage spine will increase this to a range of 20" to 60".

(b) The inclination of the spine from a vertical to the base cannot exceed 60° (total angle 120°).

(c) Relative inclination of upper plate to the radial vector pointing in the direction of the spine cannot exceed 90°.

(d) The "twist" of the upper plate cannot normally exceed 35°, although this limit is a function of endpoint position.

Note that constraints (a) and (b) are principally hardware-induced limits, whereas (c) and (d) reflect constraints inherent to the cable linkage geometry, wherein cables must be maintained in tension, and the approximate octahedron formed by the planes containing { cables + base + upper platform } must remain convex. If these criteria are not met, the linkage will lose rigidity.

7. Performance

Time response and loading tests have been conducted on a three-degree-of-freedom (tripod) version of this manipulator, and results are extrapolated for the six cable arm, whose performance evaluation is in progress.

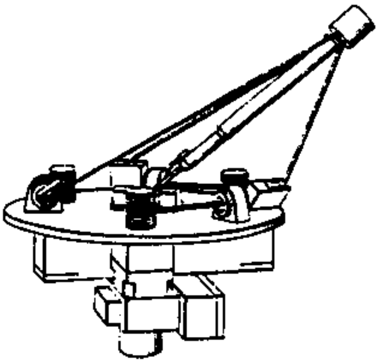

Figure 3. Tripod (3 d.o.f.) arm

Load Capacity is highly configuration dependent, ranging from a maximal 1100 lb force applied to the manipulator in its stronger region (cables of roughly equal length) to a 30 lb load on the manipulator extended 5' at an inclination of 30° to the horizontal. The moment capacity has a similar range, with a 100 ft-lb upper bound.

Accuracy: positioning repeatability of .005", as measured with a dial indicator. Absolute accuracy has not yet been measured. The smallest repeatable step increment: .0005".

Speed: maximum slew rate of 140 in/sec (corresponding to servovalve flow saturation). A minimum speed of .002 in/sec was measured, smoothly tracking a .004" amplitude, .1 Hz signal.

Dynamic Response: in response to sinusoidal input, -3db amplitude attenuation occured at 30 Hz, both for the loaded and unloaded manipulator (5 kg mass attached to tip). Step response settling times were 30 - 40 ms for the manipulator, again with little sensitivity to the presence of a 5 kg mass.

8. Future

The author's efforts are but a first iteration in a continuing design process. There is great room for improvement in the realm of both control and design parameter optimization, e.g. base and upper plate sizing. It will be fruitful to explore hybrid systems in which some degrees of freedom are controlled via parallel, some via serial linkages. For example, a tripod arm may be equipped with an end effector itself capable of three rotational degrees of freedom. A parallel linkage might serve as a moveable base for a complete 6 d.o.f. serial arm. Or a strong and dextrous serial linkage might be synthesized by staging multiple cable platforms one atop another (with flexible cable conduits providing upper stage control). Even the present relatively unsophisticated effort has produced a manipulator which, granted its workspace restrictions, has demonstrated the high performance and practical versatility of a parallel cable-actuated workstation.

REFERENCES

1. Stewart, D., "A Platform with Six Degrees of Freedom,"*Proceedings of the Institution of Mechanical Engineers,* Vol. 180, part 1, No. 15, 1965-66, pp. 371-386.

2. McCallion, H., and Truong, P.D., "The Analysis of a Six-Degree-of-Freedom Workstation for Mechanised Assembly," *Proceedings of the Fifth World Congress for the Theory of Machines and Mechanisms,* an ASME Publication, 1979, pp. 611-616.

3. Yang, D.C.H., and Lee, T.W., "Feasibility Study of a Platform Type of Robotic Manipulators from a Kinematic Viewpoint,"*Journal of Mechanisms, Transmissions, and Automation in Design,* an ASME publication, Vol. 106, June 1984, pp. 191-198.

4. Hunt, K. H., "Structural Kinematics of In-Parallel-Actuated Robot Arms," ASME Paper 82-DET-105, 1982.

4. Landsberger, S.E., Master's Thesis, Massachusetts Institute of Technology, September, 1984.

4. Landsberger, S.E. and Sheridan, T.B., "A New Design for Parallel Link Manipulators," *Proceedings of the 1985 IEEE International Conference on Systems, Man and Cybernetics,* pp.812-814.