HUMAN FACTORS IN SPACE TELEPRESENCE

David L. Akin
Russell D. Howard
Joseph S. Oliveira

October, 1983 SSL# 41-83

(Under NASA Contract NASW-3797)

HUMAN FACTORS IN SPACE TELEPRESENCE

David L. Akin
Russell D. Howard
Joseph S. Oliveira

Space Systems Laboratory
Massachusetts Institute of Technology
Cambridge, Massachusetts

SSL# 41-83 October, 1983
(Under NASA Contract NASW-3797)

TABLE OF CONTENTS

Human Factors in Space Telepresence

1. Introduction

This report discusses the problems of interfacing a human with a teleoperation system, for work in space. Much of the information presented here is the result of experience gained by the M.I.T. Space Systems Laboratory during the past two years of work on the ARAMIS (Automation, Robotics, and Machine Intelligence Systems) project (NASA contract #NAS8-34381).

Many factors impact the design of the man-machine interface for a teleoperator. In this paper the effects of each are described in turn. An annotated bibliography gives the key references that were used. No conclusions can be presented as a "best design," since much depends on the particular application desired, and the relevant technology is swiftly changing.

Much of the traditional work in human factors research is in the area of anthropometry. This work is mentioned in the section on Human Capabilities, but is not discussed in depth since this information is difficult to systematize, and too voluminous to enumerate here. Quite a bit of this data is required for the final design of a man-machine system, but the main issues dealt with here concern architecture-level alternatives. These depend more on some broad aspects of human behavior (which can be described concisely) than on the details of anthropometry.

The term telepresence is used synonymously with the word teleoperation here; it is used because it conveys a greater emphasis on "accommodating the human" into the system. Telepresence is a term used to describe all types of operations which involve a mechanical manipulator controlled by a human at some remote site. Strictly speaking, this definition could be construed to include even a human using a long wrench. In fact, telepresence systems are simply a class of tools which form a continuum from basic hand tools through powered tools, mechanical exoskeletons, direct-link master-slave manipulators, all the way to the most sophisticated semi-autonomous robot under the loose control of a human supervisor. Any of these tools is capable of performing useful work in space. From an economic and procedural standpoint, however, the key difference for space applications is between those systems which allow the controlling human to be a large distance from the worksite (without a direct physical connection) and those which require his proximity. The former systems are those considered here as candidates for space telepresence, since they allow the human operator to be on the ground or in a low-orbit space station, avoiding substantial transportation and life-support costs. All such systems can be broken down, as in Figure 1, into four basic elements: the task, the manipulator mechanical components, the manipulator control system, and the human in charge. This paper concentrates on the interface between the human and the machine, with the intent of summarizing the problem involved and the work done to date.

Figure 1:

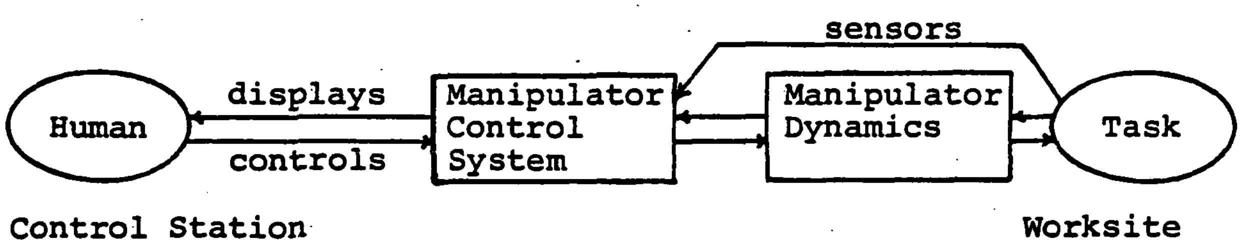

The nature of the particular tasks to be accomplished will determine the criteria to be used in designing and evaluating telepresence systems. The choice of the most effective man-machine interface depends on these criteria, and is subject to practical constraints on mechanical and control system design. The proposed tasks, for instance, will set requirements on end-effector design, size and shape of manipulator working envelope, number and type of degrees of freedom, levels of strength and positioning accuracy, as well as determining what sensors will be useful. The mechanical constraints limit what is realizable in terms of geometry and dynamics. In addition, the relatively fixed (but not entirely understood) characteristics of the human in control are important design parameters. The effects of each of these constraints on space telepresence will be discussed in turn before the results of particular experiments are summarized.

2. Description of Tasks

Human factors research in telepresence usually proceeds by testing the performance of a man-in-the-loop manipulator system on a given set of tasks. The tasks are chosen to simulate an expected application of the manipulator system. Hence, the results of such research must be interpreted in the context of these anticipated tasks. Most of the literature published to date falls into a few broad areas, distinguished by the type of tasks assumed. In this section, these underlying assumptions will be described and the differences pointed out between terrestrial and space applications.

Much of the work done in telepresence has been concerned with the problems encountered by the nuclear industry in handling radioactive materials. The tasks here include reactor maintenance: component disassembly, reassembly, and parts transportation; reactor operations such as handling and packaging of fuel and wastes; and laboratory operations such as radioactive chemical processing. Other nuclear applications occur in high-energy experimental physics, where accelerator targets must be exchanged and serviced rapidly in a radioactive environment. Some of these tasks involve large forces and extensive work spaces, requiring overhead traveling bridge cranes or rail-mounted material carriers, while others (such as chemical processing in hotlabs) require precise handling of small objects in a confined area.

The working environment is hot (thermally as well as radioactively), and is particularly hostile to solid-state electronics. Corrosion, oxidation and weakening of metals is also a problem. Maintenance is very expensive, involving elaborate decontamination procedures. The size and weight of the telepresence equipment is not of critical importance, so systems tend to be of conservative design with large safety factors. The physical separation between worksite and control station is often small, sometimes allowing direct mechanical linkage between controls and manipulator and direct vision (through shielding windows) of the worksite by the operator. Television cameras are supplied in other circumstances. The nuclear industry has gained a good deal of experience in the field of telepresence over the last forty years, although their current equipment is based on the technology of the 1960's.

Another major area of telepresence research has been for underwater applications. The U.S. Navy has been involved in many operations at depths which are not attainable by divers. Such tasks include finding and retrieving sunken objects, cutting away and removing obstructions, etc. A specific example would be to locate the wreckage of an airplane and remove a particular piece of equipment, such as a flight data recorder. Recently, commercial interests have found other applications for undersea teleoperators. Offshore oil wells, underwater cables and pipelines require regular inspection and maintenance. While divers can perform much of this work, it is often cheaper (both in terms of money and human life) to use mechanical systems.

A typical commercial task would be to clean off an area of structure with a high-velocity water jet and inspect the welds for cracks with a television camera. The worksite environment involves low temperatures and high pressures. Lighting must be provided, and vision is often obscured by sediment and debris, particularly during cleaning operations. Constantly shifting forces due to currents tend to disturb the relationship between the manipulator and its target. For shallow-water applications a human may be present at the work-site, either as a diver or within a submersible to which the telepresence system is attached. In these cases the human in control may use direct visual sensing and direct mechanical control of the manipulator. In deep water, however, the systems are connected to a control station (in a surface vessel) by a cable or an acoustic link. These links have limited information bandwidth capability and the acoustic links introduce a time delay on the order of a few seconds. Each dive may last several hours, and the equipment can be maintained and refurbished on the surface between dives, so the reliability requirements are much different from those of the nuclear industry.

The bulk of the human factors research on telepresence has been motivated by the requirements of the underwater and nuclear industries. A smaller number of contributions concern biomedical applications of telepresence, specifically orthotics and prosthetics. Some work has also been done on the topic of interest here: space telepresence.

The requirements for space telepresence differ significantly from those of the other applications discussed. Several NASA studies have identified the types of tasks which are candidates to be accomplished by remote control in the near future. The manipulator system considered would be attached to a free-flying propulsion module, and could be space-based (at a space station, for instance) or ground-based (delivered to orbit by Shuttle). The control station is usually assumed to be on the ground, with communication through TDRSS. However, control from the Shuttle or a space station is also possible.

The most basic task for space telepresence is the orbital boost or reboost of a satellite, using the propulsion module. Examples include include delivery of communications satellites to geosynchronous orbit, or astronomical observatory satellites to orbits out of the Shuttle's reach. When delivery is completed, the teleoperator may also observe and assist in the deployment of antennas or solar arrays needed to place the satellite in its operational configuration.

A potentially very profitable use of space telepresence is the maintenance and repair of satellites in orbit. With such a capability, satellites can be designed for in-space resupply of consumables such as fuel and batteries, extending their service life to previously unattainable levels. This is particularly useful for NASA's planned orbiting observatories, such as Space Telescope and AXAF (Advanced X-ray Astrophysics Facility). These are intended to be semipermanent facilities, with consumables and modular systems replaceable in orbit. Another benefit of space telepresence is the ability to repair a malfunctioning satellite. The Solar Maximum Mission and Landsat satellites illustrate the difficulties which can result from minor hardware problems, which could be fixed with an on-orbit repair system.

Since satellite maintenance is likely to be the most effective use of space telepresence in the near future (15 years), the tasks involved will be examined here in some detail. Projects on the drawing boards now which incorporate orbital maintainability are designed for servicing by humans in EVA, since that technology is currently available (for orbits within reach of the Shuttle). Thus, the basic levels of dexterity, reach, and strength required to perform the designed maintenance tasks for these satellites (e.g. Space Telescope) are those of a space-suited human. Once an operational telepresence system has been demonstrated, satellite designs will begin to reflect the specific capabilities of mechanical manipulators for maintenance. This may relax some constraints on satellite design, as some human limitations do not apply to mechanical manipulators. However, in many cases it will still be desirable to allow maintenance by human in EVA, as a backup alternative.

Maintenance operations for satellites fall into three categories: scheduled, unscheduled, and contingency. Scheduled operations are designed for and take place at planned times. Unscheduled operations are designed for, but take place when required. Contingency operations only take place in the event of an unplanned component failure. A space telepresence system would be designed to handle the scheduled and unscheduled tasks, and advanced systems will be flexible enough to perform many contingency tasks as well. There will always be some classes of contingency repairs which require more dexterity than any given mechanical system can provide, but with modular design the likelihood of such a contingency is minimized. For example, an entire damaged module can be replaced if internal repairs are impossible.

The scheduled and unscheduled tasks for the Space Telescope (ST) project are well-defined at this point (the satellite is planned for a launch on STS-25 in 1985). Orbital maintenance is possible for a total of 23 orbital replacement units (ORU's) aboard ST. These consist of 5 Scientific Instruments (SI's), 3 Fine Guidance Sensors (FGS's), the Science Instrument Control and Data Handling Unit (SI C&DH), 3 Rate Sensor Units (RSU's), 3 Rate Gyro Electronics Units (RGE's), 3 Fine Guidance Electronics Units (FGE's), and 5 Batteries. Certain other malfunctions (such as faulty solar array deployment) can be handled on a contingency basis.

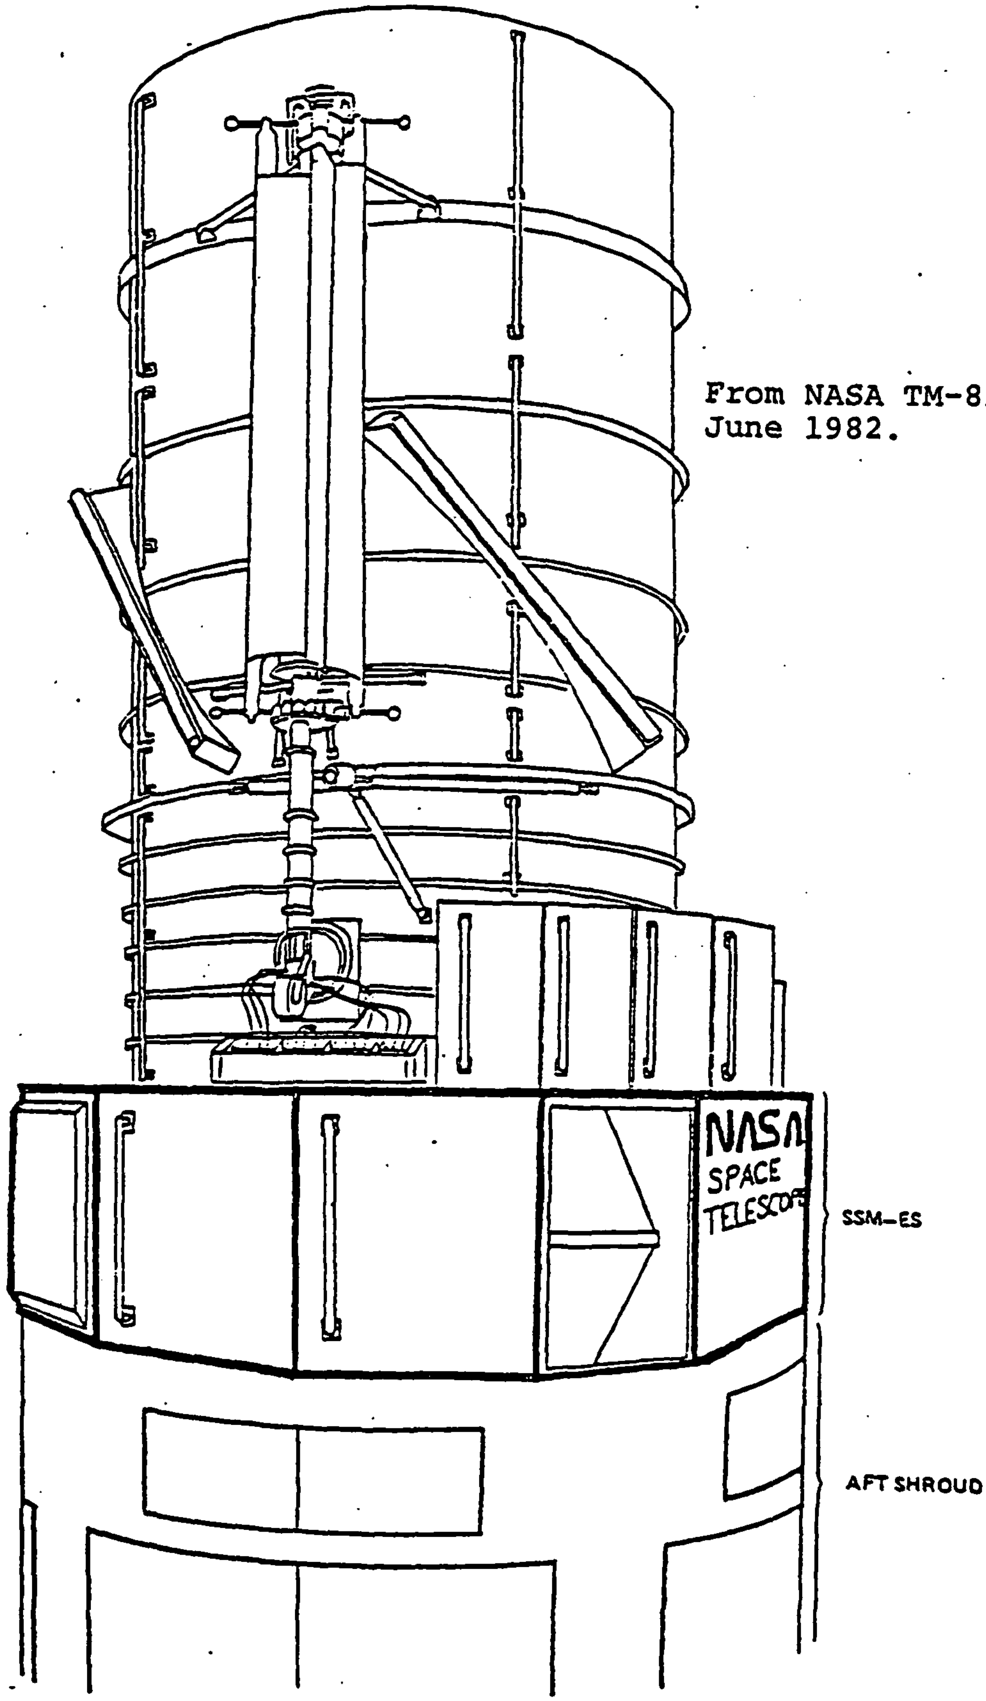

From NASA TM-82485, June 1982.

FIGURE 2: ST SUPPORT SYSTEMS MODULE EQUIPMENT SECTION AND AFT SHROUD.

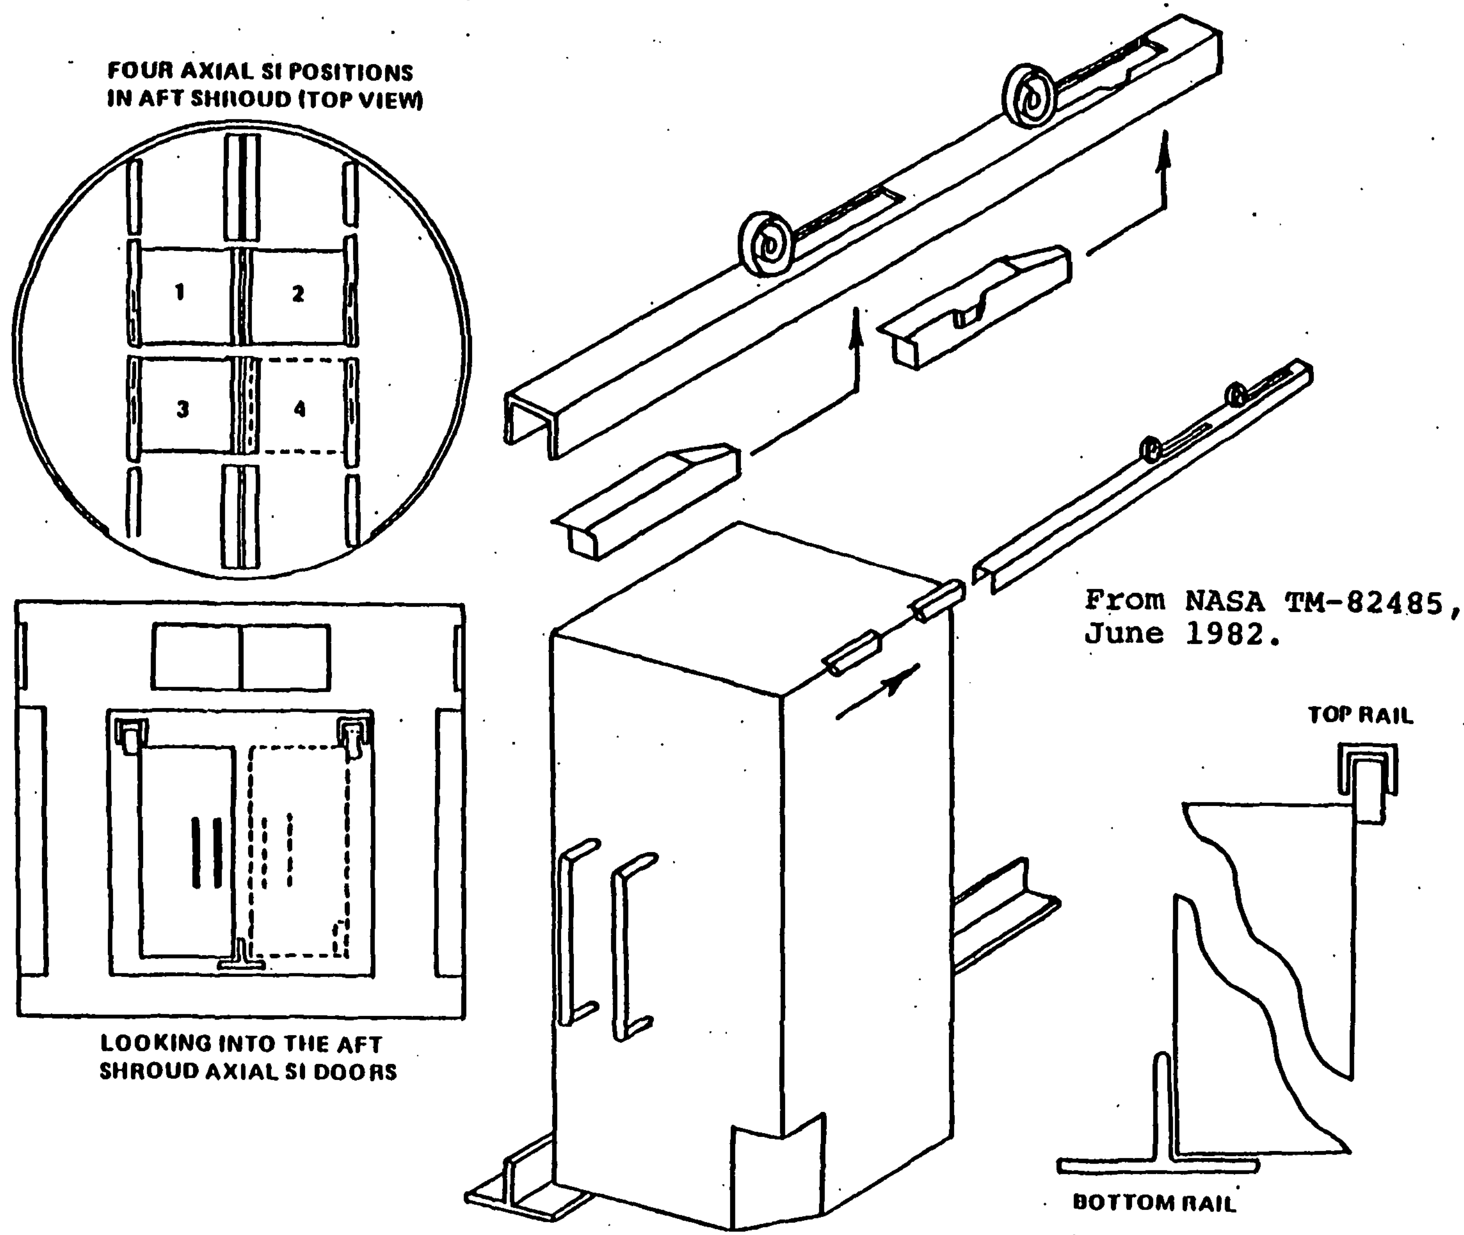

From NASA TM-82485,
June 1982.

FIGURE 3: AXIAL SI RAILS SHOWING MOVEMENT.

The support module and aft shroud are depicted in Figure 2. Figure 3 shows the location and mounting hardware for the axial SI's in the aft shroud. The location of the fine guidance sensors is shown in Figure 4. The latching mechanisms used are typified by the J-hook fasteners used on the light shields of the RSU's and on the batteries (Figure 5). Electrical connections are shown in Figure 6 and Figure 7, which shows the mounting of a typical ORU such as the SI C & DH. The maintenance tasks consist of locating the defective unit, gaining access to it, disconnecting the electrical contacts (if any) and removing the unit. Replacement is performed in reverse order.

For a near-term application of space telepresence, Space Telescope maintenance tasks are typical. Later applications will include structural assembly in space, which will possess its own vocabulary of tasks. Most large assembly projects proposed involve the connection of beams into tetrahedrons as basic structural elements, and a typical connector design (from MIT) is shown in Figure 8. More complicated versions will also be required to connect fluid- and power-transfer utilities and data lines. A more complete description of these tasks must await detailed project designs.

Other missions which a space telepresence system may be called upon to perform include rescue and exploration (lunar, asteroidal or planetary). Rescue tasks are not well-defined in advance because a large variety of situations could become

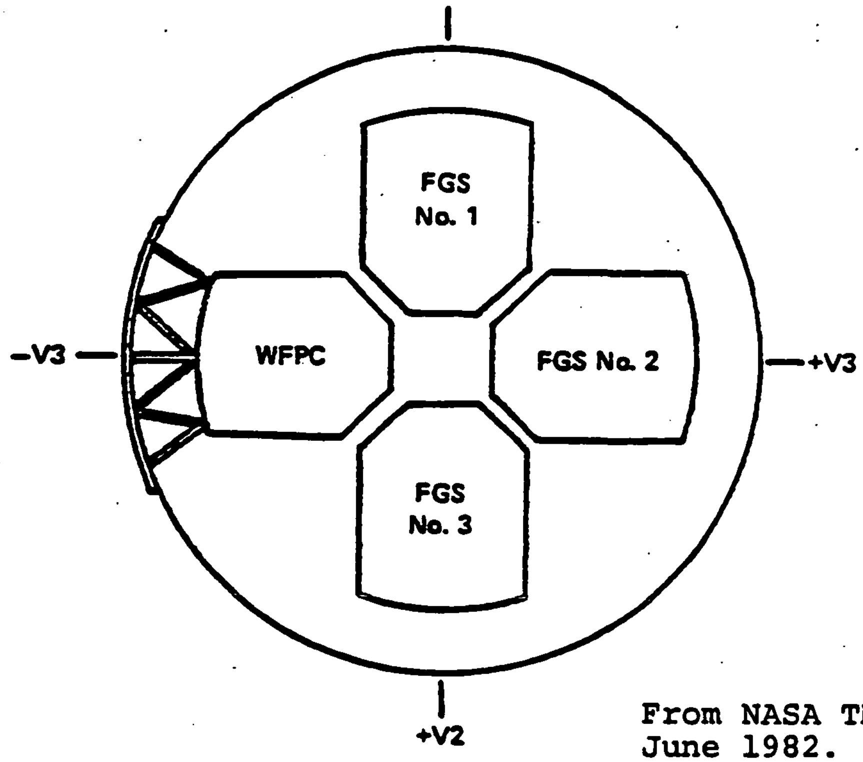

From NASA TM-82485,
June 1982.

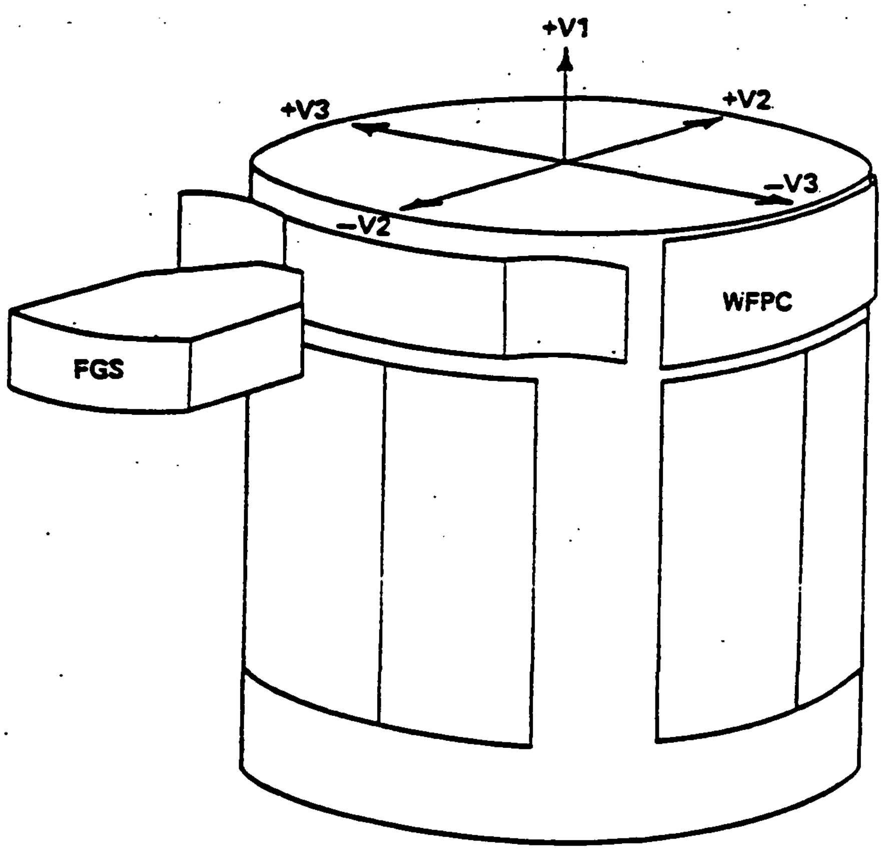

FIGURE 4: FGS ORIENTATION IN AS.

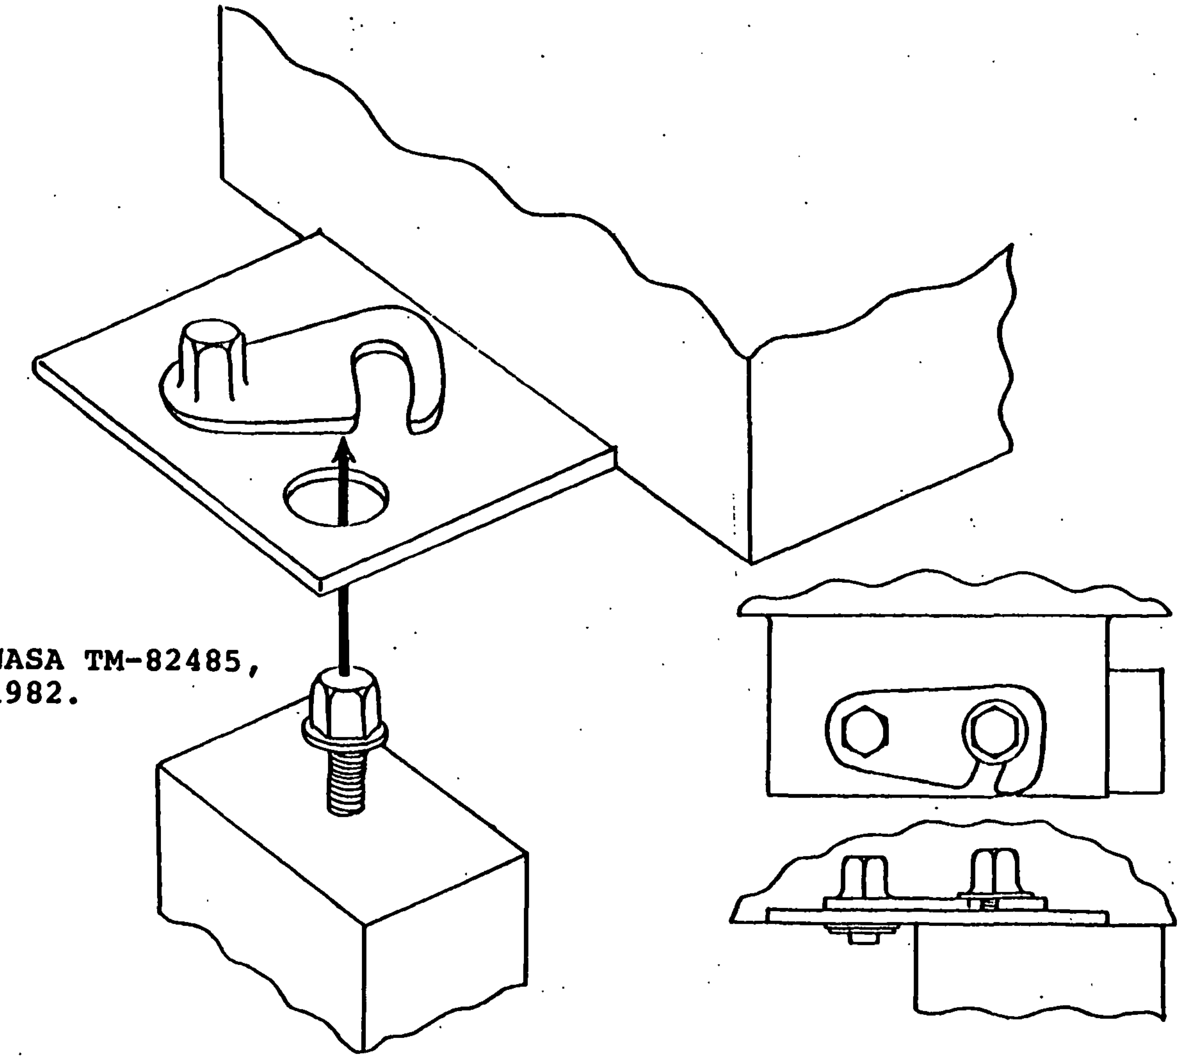

From NASA TM-82485,
June 1982.

FIGURE 5: J-HOOK LATCH.

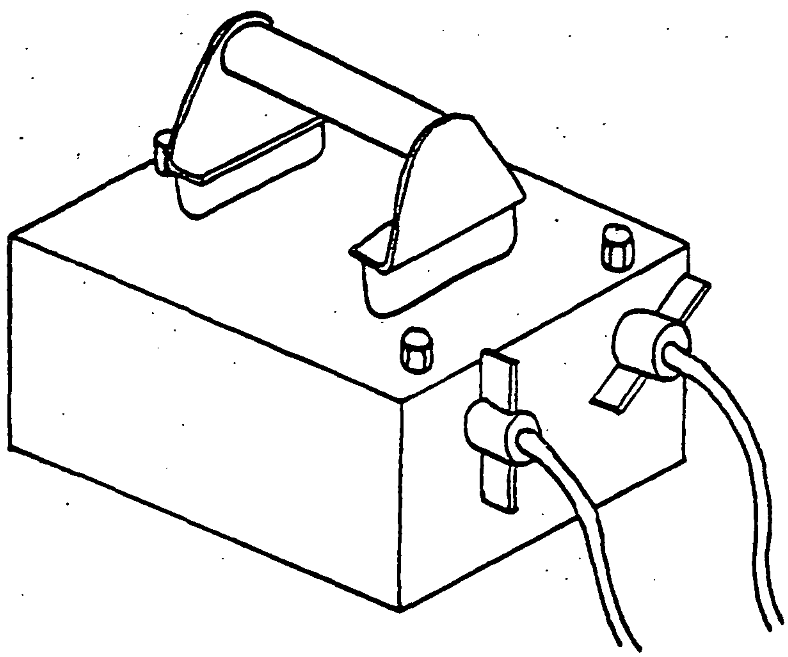

FIGURE 6: RSU ELECTRICAL WING TAB CONNECTORS.

FROM NASA TM-82485, JUNE 1982.

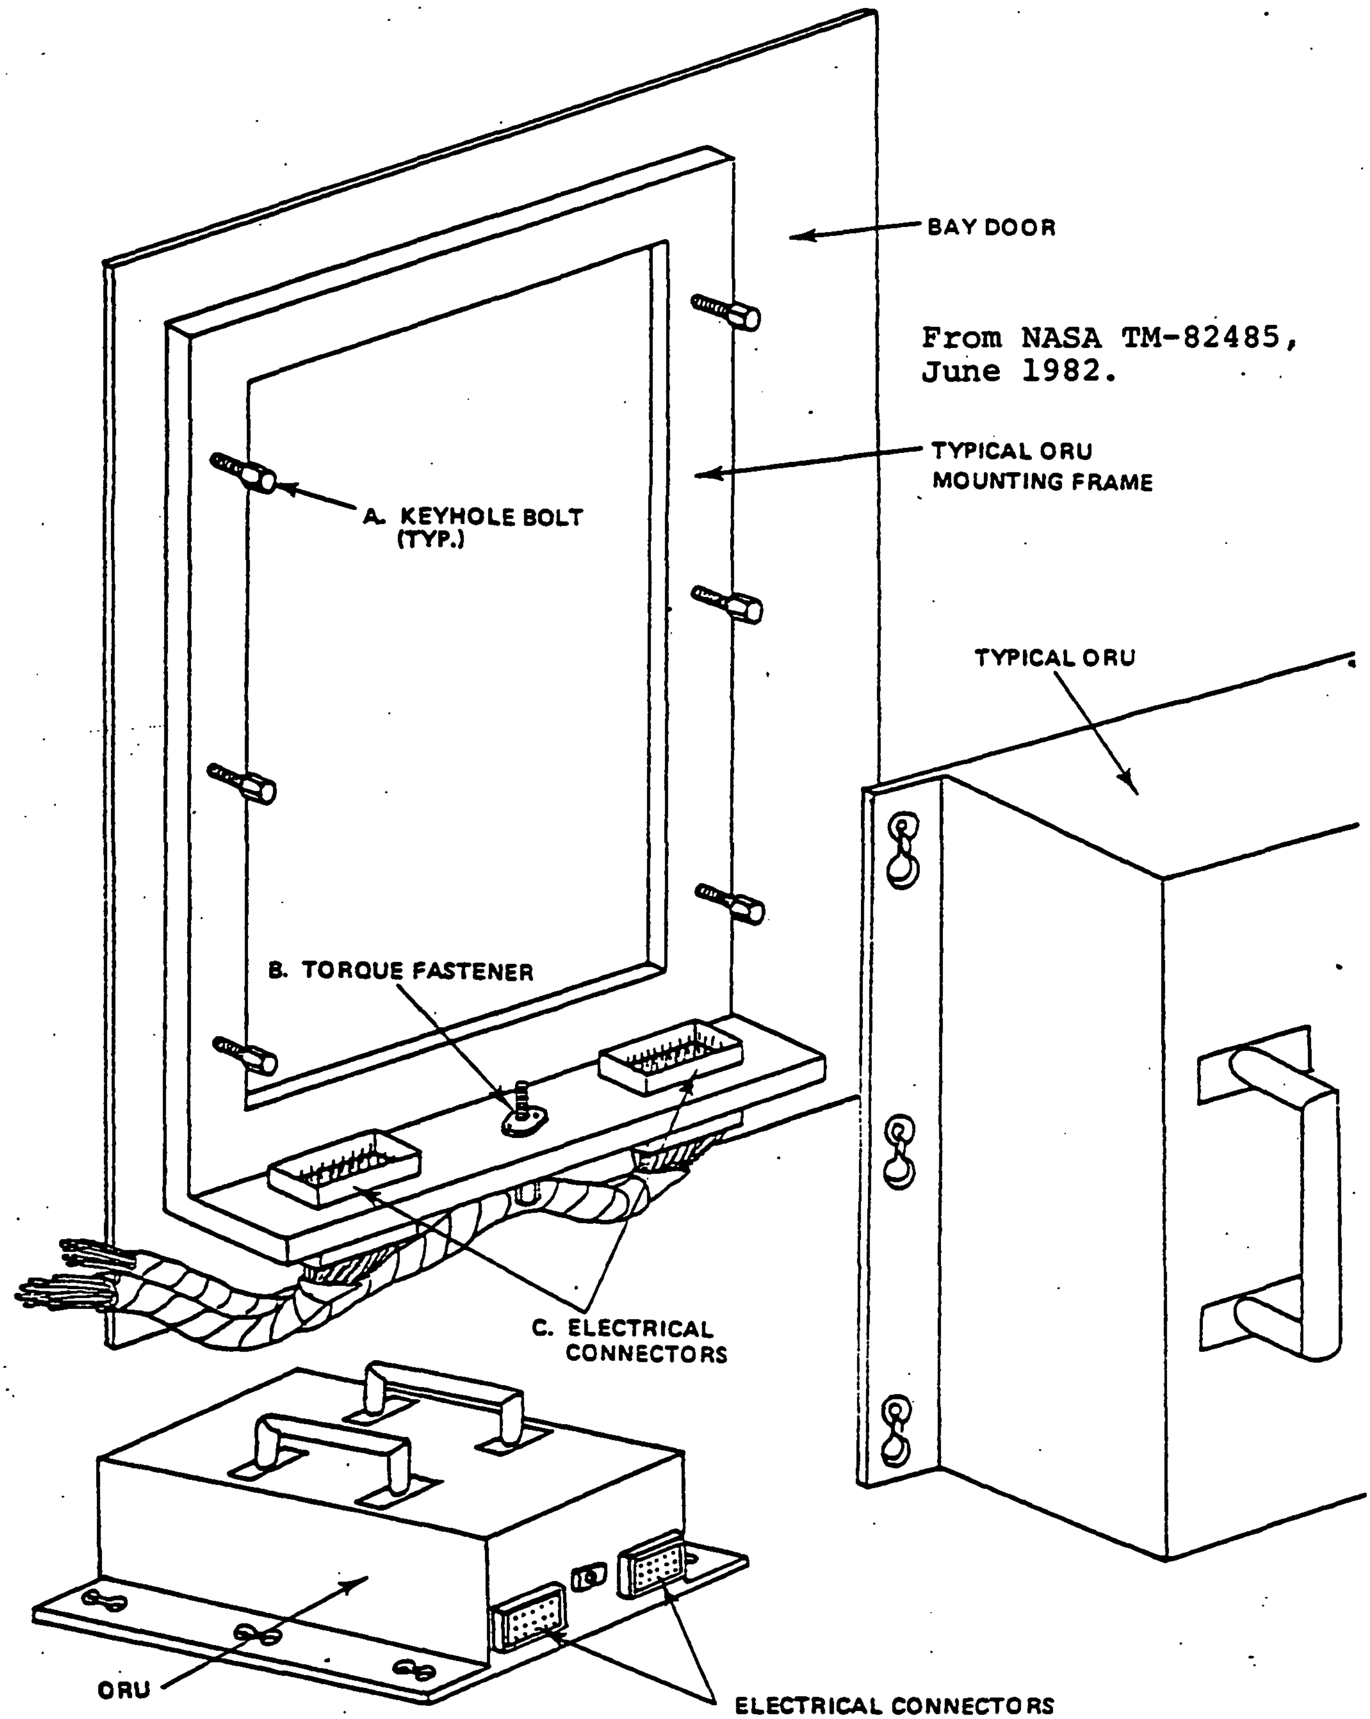

FIGURE 7: TYPICAL ORU (E.G., SI C&H) DOOR MOUNTING SYSTEM.

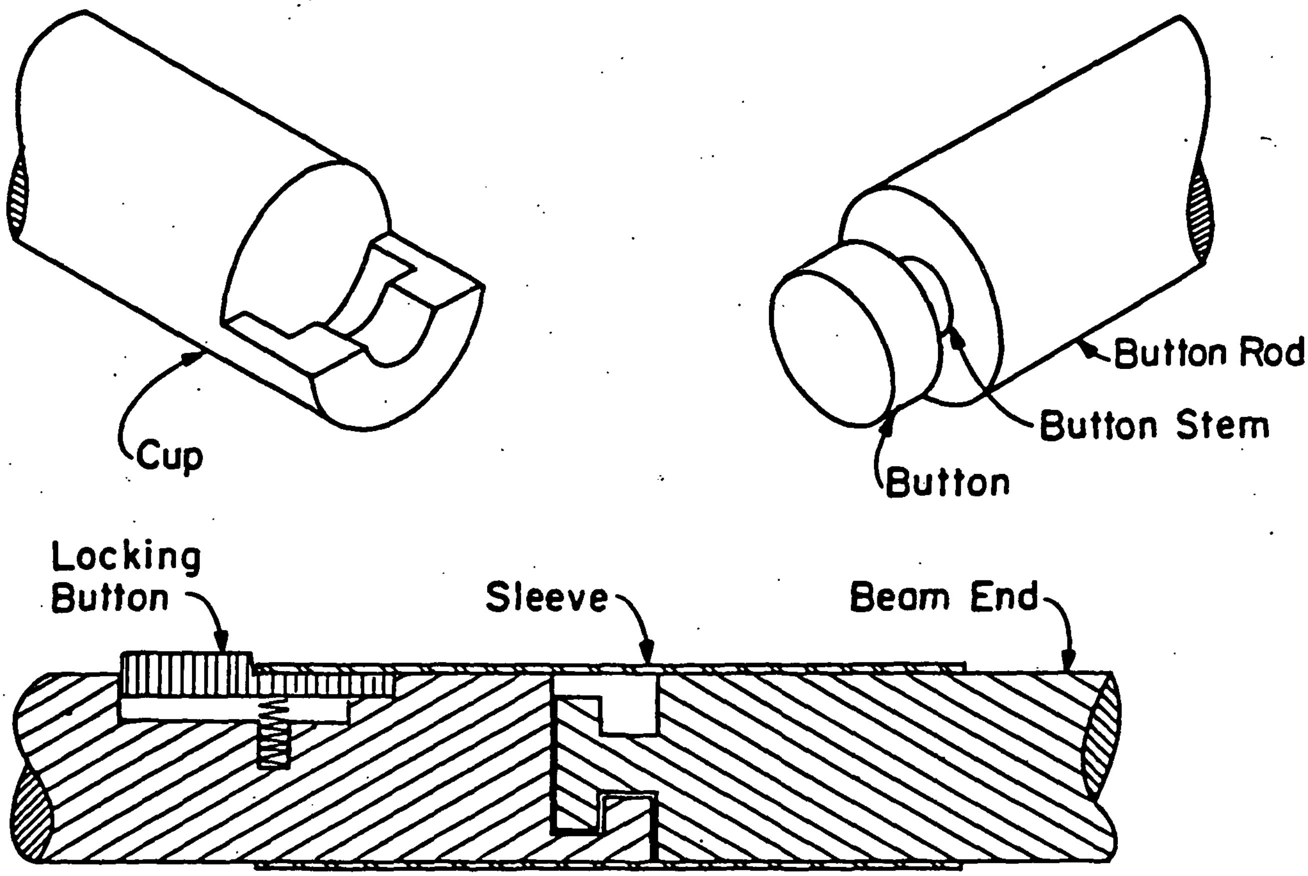

FIGURE 8: MIT CONNECTOR DESIGN

hazardous. In an emergency, any additional rescue capability provided by a telepresence system would be welcome.

Telepresence has already been used in lunar (Surveyor's shovel) and planetary (Viking's sampling arm) exploration. The controls were rather cumbersome on each and improvements can be expected, but they demonstrated the utility of even crude telepresence for the analysis and exploration of a planetary surface. For extensive explorations roving vehicles have been proposed, using telepresence techniques.

The factor which most distinguishes the tasks of space telepresence from their terrestrial counterparts is the environment in which they take place. One important difference is the distance between the control station and the worksite. In the undersea or nuclear applications this distance ranges from a few meters to perhaps a kilometer, while for space systems the separation is typically thousands of kilometers. The most obvious consequence of this separation (characteristic of all space operations) is the large transportation cost involved in getting the manipulator to the worksite. Delivery to low earth orbit costs about $2000/kilogram, so there is incentive to eliminate excess weight and bulk. These costs also affect any maintenance and refurbishment needed, so a successful design would stress reliability and longevity, while remaining compact and lightweight.

Another consequence of large distances is the communication problem. Communications will probably be through the TDRS system, which has a maximum capacity of 300 million bits per second (the equivalent of 600 television frames per second) on ku-band. This limit should not constrain teleoperator performance appreciably, but there is another communications-related factor which has a critical effect on control system design: time delays. For a link from the worksite in orbit to a controller on the ground, the information is transmitted first to a TDRSS satellite in geosynchronous orbit, relayed from there to the ground station at White Sands and then transmitted over surface lines to the operator's control station. The control commands retrace this path in reverse. The time delays in the loop come primarily from information handling and reformatting, with some contribution from the finite speed of light traversing the distance. The total round-trip delay is between .5 and 2 seconds, depending on circumstances. For planetary exploration applications, the time delays can become minutes or hours.

Radiation levels in earth orbit can be as high as 10^9 electrons per square centimeter per second (at energies greater than 0.5 MeV). This is one of the reasons for using telepresence in space: humans require shielding in such an environment, particularly in high orbits (such as geosynchronous). Such levels can also affect solid-state electronic devices, and must be taken into account in teleoperator design.

There are some important differences in the visual environment between space and terrestrial applications. In space, with no intervening material between the cameras and the target, images are clear and undistorted, unlike the undersea projects in which vision is frequently obscured (often by debris stirred up by the manipulator itself). The objects viewed (satellite components) are not subject to corrosion or sedimentation which would change their appearance over time, so they are easier to recognize. Lighting is provided by the Sun, the Earth and whatever lights are carried by the tele-operator, and is completely controllable if desired.

In contrast with the undersea environment, there are no currents to continually disturb the relationship between the manipulator vehicle and its target. However, rigid docking will be required simply to take reaction loads imposed by the manipulations. The mechanical design of the manipulator for space will not need to take gravity loads into account, although some tasks may involve working on a rotation structure, possibly requiring compensation for centrifugal forces.

The tasks required of a space telepresence system, and the environment in which they take place are different enough from those of terrestrial applications that major design tradeoffs are shifted. One of these is the tradeoff between manipulator speed and accuracy. Productivity is the key for many earthbound tasks -- speed is directly related to profit. However, in space the time spent on actual manipulations represents a small fraction of the total mission cost. Far more important is the requirement that the mission be successful, i.e. the intended manipulations are accomplished without causing unintentional damage.

Another tradeoff affected is the structural compromise between rigidity and light weight. On earth there is no great penalty for conservative design, but weight is directly related to the transportation costs in space missions. It may be possible to build a light, somewhat flexible manipulator and use a more sophisticated control system to achieve the same results for a lower mission cost.

In summary, the constraints imposed on space telepresence systems are significantly different than those for terrestrial tasks. These differences can have an important effect on man-machine interface design. Human factors studies typically involve assumptions about the type of tasks to be performed and the worksite environment. The applicability of a given study's results to the space telepresence problem depends on the correspondence of these assumptions to the expected space tasks and environment, as described above.

3. Manipulator Design

Another set of constraints which are common to all telepresence systems are the state-of-the-art limitations of manipulator construction, dynamics, and control. Many of the considerations that apply to this end of the telepresence system are identical to those encountered in the design of robotic (fully autonomous) manipulators, which are now becoming common in industry. Much of the technology developed for robotics is directly applicable to telepresence. A brief overview of possible manipulator types and their properties will be given here, since the design of the man-machine control interface depends heavily on the type of information required by the machine.

The geometrical properties of a manipulator are determined by the type and number of its joints and the links which connect them. Once these are specified, the working envelope of the manipulator is determined.

Joints connect links and permit relative motion. The majority of joints in use are of two types: revolute (R) or prismatic (P). An R-type joint is simply a hinge, allowing relative rotation of two links about an axis. Such a joint can be simply constructed and is easily driven by motors, gears, pulleys, or other rotary actuators.

A P-type joint permits sliding (translation) but no rotation. These joints are often of rectangular cross-section to prevent rotation, and are easily driven by linear actuators (hydraulic, for instance). The manipulator depicted in Figure 9 possesses both R- and P-type joints.

Other types of joints are possible. Spherical ball-and-socket joints can be modelled as three independent co-located R-type joints. Cylindrical and screw-type joints can be modelled as coaxial R and P joints.

A link is depicted in Figure 10. The actual shape does not affect the manipulator kinematics, beyond the specification of two parameters: the length ℓ and the twist α. The length is the minimum distance between the axes of the joints at either end of the link, and the twist is defined as the angle between these axes in a plane perpendicular to ℓ.

Two other parameters specify the condition of the joint: the distance s between the two links connected to it, measured along the common axis, and the angle θ between the links measured in a plane normal to this axis. In R-type joints, θ varies during motion and s is fixed; for a P-type joint the reverse is true.

From the user's standpoint, the manipulator is just a means of putting the end-effector (usually a gripper or claw) where it is needed, and in the desired orientation. The configuration of links and joints supporting the end-effector is important only to the degree that it doesn't interfere with itself or other objects in the workspace. Thus, from this point of view, the important features of a manipulator are: its working envelope -- the volume composed of all

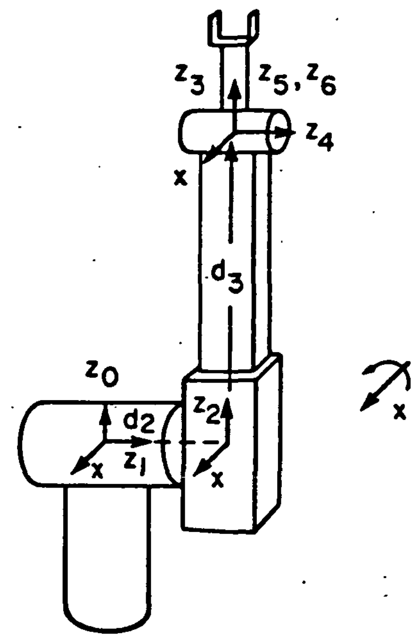

FIGURE 9: A MANIPULATOR WITH R- AND P-TYPE JOINTS

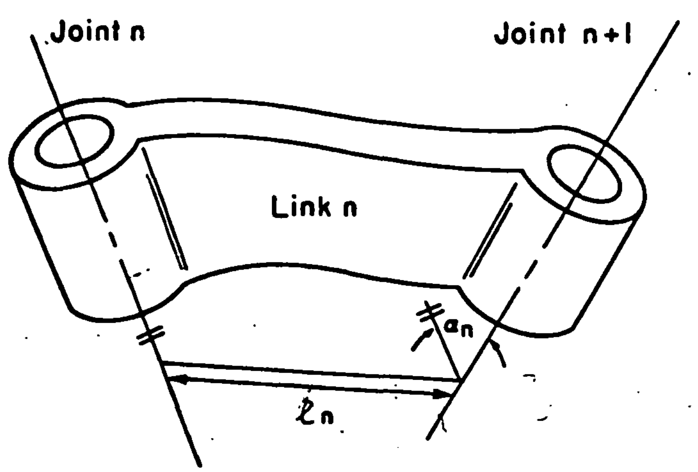

FIGURE 10: A GENERALIZED LINK

FROM PAUL, RICHARD P., "ROBOT MANIPULATORS", MIT PRESS, 1981.

attainable end-effector locations, the approach-angle characteristics (the range of orientations the end-effector can assume at each point in its working envelope), and the number of ways in which the manipulator can reach a given position and orientation.

In general, six degrees of freedom are desired at the end-effector. Three degrees of freedom allow it to be brought to any position and another three are required for orientation. This requires at least six joints, including at least three R-type joints. There is an advantage to designing all link parameters (α's and ℓ's) nonzero; the motion obtainable by a system with a zero parameter is less general than otherwise. For instance, if a link has R-type joints at either end and $\alpha = 0$ (no twist), the resulting movements are constrained to a single plane.

In addition, the more general the link parameters, the more possible ways of reaching a given position and orientation. For example, in a manipulator with six R-type joints: if $\ell_1 = \ell_3 = \ell_5 = 0$ there are at most four ways; if $\ell_3 = \ell_5 = 0$ there are at most eight ways; if $\ell_3 = 0$ there are at most sixteen ways; and it is believed that if all parameters are nonzero there are at most thirty-two different ways to reach the target position and orientation.

The type of joint used also affects the number of ways to position and orient the end-effector. In general, the use of an R-type joint instead of a P-type joint doubles the number of possible ways, increasing the ability of the system to cope with obstacles in the workspace.

Of course, not every location and orientation can be reached by the maximum number of ways referred to above. Positions outside the working envelope cannot be reached at all (by definition), and typically toward the boundaries pairs of possible ways degenerate into a single configuration.

The shape of the working envelope and the approach-angle characteristics can be calculated for any given manipulator design, but for six (or more) joints the problem is very complicated. No simple design rules have resulted from such analyses, but many have been calculated, and the results are available in the literature.

For tasks requiring dexterity, an extra (seventh) degree of freedom provided by another joint is often desirable, allowing a wide range of arm positions for any task. A figure-of-merit which is useful in evaluating systems for flexibility or dexterity is the "aspect ratio," defined as the ratio of working envelope volume to the volume of the arm itself.

Serial manipulators, in which each link depends from the previous one, typically have the highest aspect ratios. There are some disadvantages to this arrangement, however. Innaccuracies cascade through the joints; a small angular error in the "shoulder" can lead to large discrepancies at the "hand". Also, in practice, the number of control and sensor leads that are brought out from the serial arrangement can become quite large. Since most of these must twist through all the cascaded joints, lead failures may occur.

Error cascading can be reduced by arranging motions in parallel. For example, the errors in an x-y table (in x and y) are independent to first order. Parallel manipulator designs tend to "enclose" the workspace to a much greater extent than serial configurations. This makes them more suited to industrial robot applications (where the workspace is defined in advance and fixed) than for telepresence.

The actuators used can be hydraulic, pneumatic, or electric. Hydraulics are best for many applications calling for small actuators and large forces. They are used almost universally in underwater telepresence. Hydraulics are shunned by the nuclear community, primarily for historical reasons: early designs leaked in hot cells, spreading alpha contamination. Leakage and long-term degradation may make them less desirable for space applications as well. The Shuttle RMS uses electric motors to good effect.

The simplest designs place the actuators in proximity to the joints they drive. On earth, this leads to the introduction of heavy counterweights at each joint, to reduce the torque requirements due to gravity loads. Since this is not a problem in space, the counterweights can be dispensed with, reducing overall system mass considerably. Other designs trade mechanical complexity for minimum arm mass by using pulleys or tendons, allowing the actuators to be placed at the base (shoulder) of the manipulator.

Throughout manipulator research the paradigm for a mechanical arm has been the human arm, with its six-seven degrees of freedom. Perhaps because we tend to conceive of manipulation tasks in terms of our own capabilities, this design is a good compromise for general-purpose manipulation. It is likely that the best choice for space telepresence would be a six- or seven-joint serial manipulator. This type is the most popular for terrestrial telepresence, and is usually the configuration used in man-controlled manipulator research.

4. Manipulator Control

The control of a manipulator involves many decisions of varying degrees of complexity. A useful way to look at the problem is to construct a hierarchy of decision levels. Each control level deals with wider aspects of overall systems behavior than the lower levels. The upper levels deal with the system aspects that vary more slowly.

A common division is into a hierarchy of four levels, in which the highest recognizes the obstacles in the operating space and the conditions under which a task is being performed, and plans how it is to be accomplished. The next (strategic) level divides the operation into elementary movements. The tactical level performs the distribution of an elementary movement to the individual degrees of freedom, and the executive level drives the actuators on the joints.

In a telepresence system the higher levels of control are performed by man. Depending on the sophistication of the machine, this may mean that direct human control is required all the way down to the tactical level (since we are only considering systems with large controller-worksite separations, the human cannot directly perform the executive level of control), or, in the opposite extreme, only occasional human guidance is needed at the highest level. These upper-level options will be discussed under the topic of Man-Machine Interfaces.

Since the executive level is automated in all systems of interest, the characteristics and limitations of modern control methods are reflected in their performance. Control system synthesis begins with the equations of motion of the manipulator. For six or seven degrees of freedom, the derivation of these equations can be extremely complicated, made possible only recently by the development of computer programs capable of symbolic manipulation (e.g. M.I.T.'s MACSYMA). These dynamics equations relate forces and torques to positions, velocities, and accelerations, and they typically contain many thousands of terms. The next step is simplification, in which approximations appropriate to the desired performance are made to reduce the equations to manageable size.

Traditionally, control has been implemented with separate analog servos closed around each joint, or digital simulations of this. For this type of control, the dynamics are simplified by discarding all velocity-dependent (such as Coriolis and centripetal) terms, as well as nonlinear terms and those representing coupling between joints. This is a radical simplification of the dynamics, giving values for the "effective inertias" of each joint. Since these inertias vary with the position of the manipulator, the simplest approach is to use the highest values which will be encountered as the design values, and size the actuators and feedback gains accordingly. Manipulator response is always designed to be overdamped, since an underdamped (oscillatory, with overshoot) system would often collide unintentionally with its surroundings.

These simplifications lead to errors, particularly at high speeds where the velocity-dependent terms are important. When this occurs during rough motions such as parts transfer in an uncluttered workspace, it may be of no concern. However, when fine motions are required with greater accuracy, it must be noted that in the traditional control method the actuator signals are derived from the error, so a quick motion requires a large error to provide an adequate actuator signal. For this reason, accurate motions must be performed slowly.

Much theoretical work has been done on the application of digital optimal control methods to the manipulator problem. These schemes try to take into account more of the dynamics, such as the coupling between joints. Taking advantage of recent advances in semiconductor memory capabilities, many complicated functions can be pre-computed and stored in lookup tables, saving on the amount of computation which must be done in real time.

The application of artificial intelligence techniques may solve the control problem in another way, similar to the control of the human arm. The human arm has no positional transducers. Accuracy is achieved solely by successive approximations in the arm-eye-force sensing systems (in the cerebellum). Complete adoption of such a programming scheme would eventually require only the most rudimentary accuracy capabilities to be implemented in hardware.

5. Human Capabilities

Whatever the capabilities of the manipulator system, the controls must interface with a human. The relevant physical parameters are straight-forward to define and quantify. For instance, typical data for an average, male human arm are:

upper arm length	30 cm
lower arm length	27 cm
distance from center of palm to wrist	9 cm
lifting capability hand outstretched	15 kg
best fit cube for comfortable working volume	45 cm on side

A vast amount of such information is available (see bibliography), ranging from average dimensions and weights to ranges of motion and strength. The human senses have also been thoroughly described with such parameters as frequency range and discrimination, angular resolution, etc. Some intellectual components such as memory can also be directly tested and quantified, although the underlying mechanisms are not clear.

The more complicated aspects of human performance are more difficult to characterize. The ability to use information to modify behavior, the effects of training on performance, the limitations imposed by fatigue for various tasks -- these all represent functions of a complex system that is poorly understood. Usually a simple model is proposed for a narrow range of behavior, and experiments are conducted to validate the model and determine the values of the relevant parameters. Many studies of this nature have been done to determine the performance of a human in a proposed task, to directly assist in the design of "user-friendly" man-machine interfaces.

Similar studies have been performed, for different purposes, by the artificial intelligence community and by psychologists. The biological development of man's information processing systems provides examples in which complex problem solving tasks of apparent infinite degrees of freedom are reduced to real time computations. In the domain of human problem solving the division of processing labors is distributed through a hierarchy of low-level and high-level processing operations. The evolutionary aspects of human problem solving suggest a vast amount of parallel computation with a system of self-modification: a system which learns. Biological systems employ learning as a tool, by which they reduce the complexities of problem-solving.

Biological systems are goal-directed machines capable of self-organized adaptive behavior. In the construction of smart machine extensions of ourselves it would be helpful to understand the strategies by which biological systems solve complex problems and the operational procedures which characterize the process of continuous problem reduction, interpretation and solution.

The Artificial Intelligence community does not say that machines cannot be constructed before we know how Man works; but that our machines, if they are to be true extensions of ourselves, should be built in our image. The problems of parallel computation in staged hierarchical information processing structures, continuous representation of inconsistent information in a consistent form and learning are but a few of the issues which should be addressed if we are to build true extensions of ourselves.

6. Man-Machine Interface

In the telepresence systems under discussion here, the human receives most of his information about the worksite through a television system. One of the reasons for this is that the television camera is an important tool in itself, and no space teleoperator will be without one; inspection and observation are the most fundamental of its tasks. Another reason is the human's ability to quickly derive spatial relationships from visual data. Detailed results of the evaluations of different camera and operator configurations are voluminous and available in the literature (see bibliography). The task of integrating all of this information and recommending the best system would be considerably beyond the scope of this paper. However, the combination of the discussion in the text and the appended bibliography should allow the reader to identify the issues in his field of interest, and refer him to the original sources for more detailed information.

Human manipulations depend to a large extent on hand-eye coordination, a task to which a significant fraction of the brain (the cerebellum) is devoted. For this reason, efforts are made to ensure that the TV system can be used in a natural way, to take best advantage of human experience and ability. The questions commonly addressed are the minimum required number of cameras for each type of task, and their placement; the minimum resolution needed; the relative advantages of color vs. black-and-white.

Humans are by nature adaptable to new circumstances, and an important question is just how far the telepresence design engineer must go to make the operator feel natural. With training, operators can become comfortable with and quite proficient at tasks which seem to bear little relationship to previous experience (video games are a familiar example). Three-dimensional displays, including Fresnel screens and stereo TV systems (using two cameras and monitors) have been evaluated for their effect on teleoperator performance. The Naval Ocean Systems Command has developed a system which simulates a human very closely: a pair of TV cameras at the correct interocular distance mounted on a "head", whose motion is slaved to the motion of the operator's head (the TV monitors are fixed on the operator's helmet). This system is part of a very anthropomorphic device, which also includes two manipulator arms attached to a movable "trunk".

The idea of camera control by the operator's head movements allows a single individual to control both the manipulator and the camera. Further, when the monitor is fixed to his helmet, the operator can establish a natural sense of his surroundings just by looking around. Such a display is called an environmentally-stabilized visual reference, since it appears to the operator that his body is fixed in the teleoperator's frame.

This type of display contributes to the illusion that the operator is at the worksite, and reduces the chance of disorientation. With such a system it is possible to obtain depth information from a single TV camera and monitor, by small sideways motions of the operator's head (causing a change in parallax). Time-delays in the control system may reduce the "natural" effect of this display. Other disadvantages appear if the manipulator system requires cameras in locations other than the natural "head position" (e.g. substantially off to one side of the manipulator), or large changes in the camera position for some particular task. Also, some tasks may require the operator to hold his head in an uncomfortable and fatiguing posture for long periods of time.

In practice, most terrestrial telepresence systems have a control station with facilities for a video operator as well as the manipulator controller. A typical control station is depicted in Figure 11. The video operator controls the aiming, zoom and selection of cameras for display on the monitors, in response to verbal requests from the manipulator controller. In a typical arrangement the controller has one large high-resolution (1000 lines) monitor screen and two smaller ones to use as direct references in manipulation. When the two operators are trained as a team, they can switch positions occasionally to reduce fatigue. An experienced video operator often learns to anticipate the needs of the other controller, resulting in rapid and efficient coordinated action.

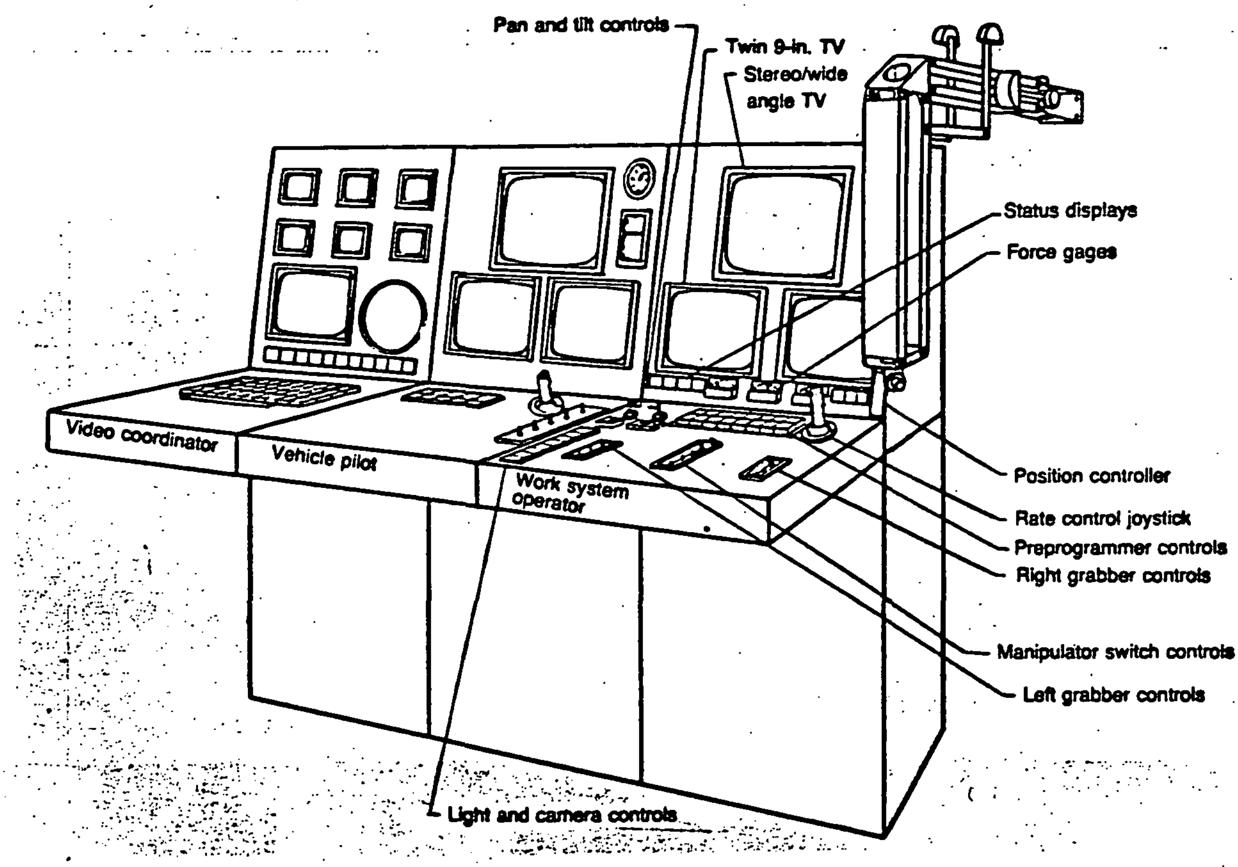

FIGURE 11: CONTROL CONSOLE CONCEPT

FROM WERNLI, R.L., "ROBOTICS UNDERSEA," MECHANICAL ENGINEERING, AUGUST 1982.

It is evident from experience that for many tasks, the sense next in importance to sight is touch. Manipulation tasks can be divided into two categories, which differ in the type of sensing and control which is most suitable. One of these categories is typified by parts transfer, and the other by assembly. The tasks identified previously for a space telepresence system include elements from both categories, but early terrestrial manipulator systems concentrated on the former, partly because of their simplicity. This led to the adoption of position control schemes, which are still the most common in telepresence and robotics. In such a scheme, the operator specifies a desired position and orientation (in some way) to the control system, which attempts to achieve that configuration of the manipulator. If an obstruction prevents the desired configuration from being attained, large and potentially destructive forces can be generated by the control system's attempts. For an industrial manipulator moving parts around this is not a great drawback, since the motions can be planned carefully, and little physical contact with the environment is required.

Assembly tasks, however, involve what is termed compliant motion. Simple examples of compliant motion are sliding along a tabletop, or pulling out a drawer. Both involve interaction with environmental constraints which are not known accurately in advance. A typical assembly motion is the insertion of a peg into a hole. This is a simple task, and a good system should accomplish it quickly and without exerting undue forces between the objects in contact.

Another example which occurs frequently in the space telepresence tasks is the tightening of a bolt. This can be done with a position-controlled manipulator: if the location and orientation of the bolt is known accurately, the circle the wrench must move along can be computed and the control system can execute the motion. To do so without exerting unnecessary forces using a rigid manipulator requires high-resolution positional transducers and fine mechanical tolerances Even so, if a human is computing the path by eye from a television image, errors are unavoidable. In a rigidly-coupled system like this a small error in position can give rise to very large reaction forces.

The alternative is simple: compliance. Compliance and force-sensing in the human arm enable a man to directly tighten a bolt with a wrench when he has only a general impression of where things are. Without force-sensing of some kind, he wouldn't even know when the bolt was tight enough.

Compliance is the ability of the manipulator to respond to forces imposed on it by the environment. It may take the form of passive compliance, mechanically built into the manipulator, or active compliance, wherein the forces are sensed and the manipulator commands are modified accordingly. Active compliance has the most general application, and, for telepresence, the loop can be closed either in the control system or in the human (by relaying the force information to him).

A system with active compliance can operate in a new control mode -- force control. In this mode the operator specifies the components of force which the manipulator should exert, and the control system produces the required motion. For many assembly-type tasks, this is more natural than position control.

Whether position or force control is appropriate for a given task depends on the constraints. If the task implies position constraints (such as the drawer example) it is not appropriate for the control system to provide conflicting position constraints in the same directions. To illustrate an extreme case, consider a manipulator whose end-effector is imbedded in a fixed object. The manipulator has no positional freedom, and position control is meaningless. Conversely, the manipulator is free to exert any force commanded. In the opposite extreme, consider a manipulator whose end is free in space and unconstrained. In this case force control is meaningless and position control is natural.

In practice, most tasks fall between these extremes, and the best solution is a hybrid of position and force control. For instance, if the constraints can be expressed as a surface, position control should be used tangent to the surface and force control normal to it.

The implementation of active compliance and force control requires force sensing and a means of closing the loop. When the loop is to be closed in the human operator, the sensory information must be relayed to him in a form he can use. Force sensing can also provide useful information not directly related to the control problem. Properties of objects in the environment can be estimated, including mass, moment of inertia, and frictional resistance. A sense of touch (essentially a more refined version of force sensing) can be used to provide information about environmental features which are not available from visual data. Various types of tactile sensors are under development for manipulators. These include "artificial skin," consisting of an array of pressure transducers imbedded in a flexible matrix, as well as more conventional pressure switches. To aid in maintaining a grip on an object, slip sensors have been devised which can detect the direction and magnitude of relative motion between the manipulator's hand and the object's surface.

To avoid damaging a delicate object, the manipulator operator will often approach it slowly, so that the "collision" occurs at a low velocity. Proximity sensors (Figure 12) have been developed to help the operator control this phase of manipulation. JPL's proximity sensors are electro-optical: the sensor contains an infrared light source which is focused on the target area (a few centimeters in front of the manipulator

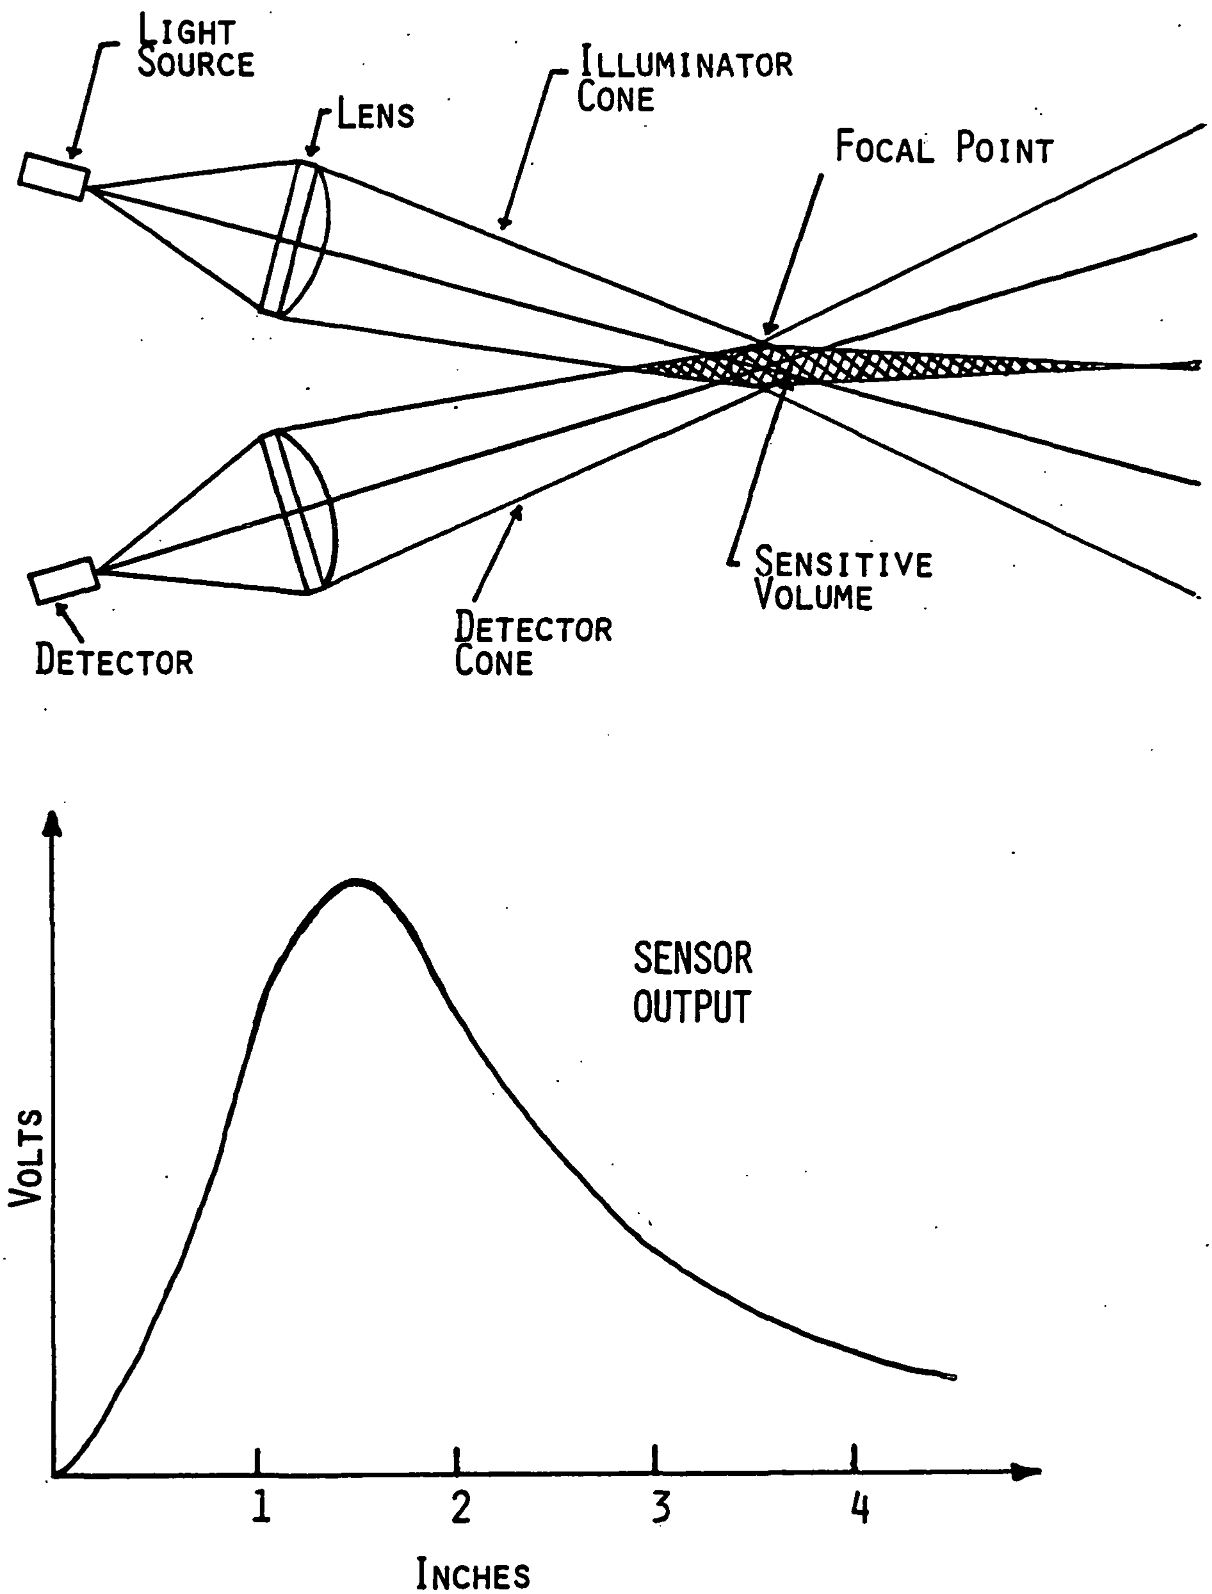

FIGURE 12: PROXIMITY SENSOR CONCEPT

FROM BEJCZY, A.K., "SENSORS, CONTROLS, AND MAN-MACHINE INTERFACE FOR ADVANCED TELEOPERATION," SCIENCE, V. 208, PP. 1327-1335, 20 JUNE 1980.

jaws) and a detector, focused on the same region, which measures the reflected light. Thus the output of the sensor is a function of the distance to an object (within the sensitive volume). When more than one of these sensors are used, information can be deduced about the alignment of the jaws with the target.

A key problem in utilizing these auxiliary sensors is the presentation of the information to the operator. In principle the tactile information could be relayed to a device which would stimulate the operator's sense of touch, but no satisfactory device yet exists. Perhaps (as suggested by Dr. Marvin Minsky of MIT) one could be developed along the lines of a project at Stanford, which has built a unit which translates printed shapes (letters) into patterns of vibration on its surface, allowing the blind to interpret standard printed material.

In the absence of a tactile display, the most likely means of presentation is a graphic (visual) display, for tactile and proximity information. Difficulties arise when several proximity sensors, or a large array of tactile sensors are used: there is too much information for the human to effectively utilize. Fortunately, although all of the sensors may be contributing useful information, the human is usually controlling only one or two parameters at a time. This makes it possible to use a display format which allows him to quickly extract the information he wants. For instance, a bar graph display is often more useful than a column of numbers. At one time the controller may only need to know that the highest pressure being exerted on any part of the target is below a certain limit, and at another time he may just want to be sure there is no slipping taking place. Eventually, "smart" telepresence systems may exist which will have some understanding of the task that is being accomplished and present to the operator only the information he needs.

An illustration of this is provided by JPL's experimental event-driven display for payload handling with the shuttle RMS. Successful ground tests of this system were conducted at the Johnson Space Center under simulated payload-handling conditions. In this system, the data from four proximity sensors attached to a four-claw mechanical hand were integrated into a visual display showing range, pitch, and yaw error values, and indicating whether a successful grasp of the target could be performed. This display enabled the operator to finely control the grasp to prevent preloading the target.

Another possibility for some types of information is an audible display. Experiments performed with aircraft simulations have shown that pilots can control one function displayed aurally together with a different function displayed visually better than if both control functions are displayed visually on separate displays. Since audible displays do not take up any of the operator's attention when they are not emitting sound, they are also particularly useful for signalling contingency events such as excessive force application, collision of the arm with obstacles, malfunctions, etc. Each type of warning would have its own distinctive sound pattern. Thus the operator would not have to be looking at the relevant display to be immediately aware of the problem.

Related to the topic of audible displays are those of computer-synthesized speech and computerized voice recognition, Communication by voice is a natural way to control functions which now require keyboard entry or another human operator. One example is the video operator, who responds to verbal commands from the manipulator controller. A sophisticated computerized voice recognition system could take over this function. Current systems have limited vocabularies and must be "trained" by the individuals who will be using them, but are capable of reliable performance within these limits.

As previously discussed, significantly better manipulator performance is possible when force control can be used for assembly tasks. A variety of man-machine interface designs for telepresence have been investigated, using different types of sensors and controls, and achieving varying degrees of success.

Each design embodies a compromise between complexity and performance. The early telepresence systems developed for nuclear applications were designed for position control only; open loop in the sense of force control. The operator controlled the joint actuators by switches, in the simplest version. Each switch controlled a different degree of freedom and allowed a single velocity to be given to the joint. Operations were quite slow, as only one degree of freedom was used at a time, due to the difficulty of combining motions into the desired resultant.

Some improvement was obtained with a proportional velocity control in a joystick. This allowed simultaneous motions in more than one degree of freedom, and reduced task times. The next step in sophistication was to introduce a CID (Control Input Device), which was often an exoskeleton fitting over the operator's arm, containing the same number of joints as the manipulator. The joint settings in the CID were used to command the joints in the manipulator. This is called a master-slave manipulator because the arm is kinematically similar to the CID and tries to duplicate its position. With this system all of the degrees of freedom can be controlled simultaneously.

With the introduction of computers to do fast real-time computation of geometrical transformations, strict kinematic similarity is not necessary between the master and slave arms. For instance, when control of the end-effector position and orientation is required, and details of the joints can be arbitrary, any sort of mechanical linkage can be used to support the operator's hand control (for direct position control). The desired end-effector position is read by the computer, which calculates the necessary joint positions. This is known as resolved-motion control, and permits greater freedom in design of the CID, while retaining many advantages of master-slave designs. One of the control modes of the shuttle's RMS is of this type; in this case the CID consists of two hand controllers -- one for rotation of the end-effector, the other for translation. The end-effector velocity is proportional to the deflection of the hand controllers, which are similar to joysticks. The resulting control system is much more compact than a master-slave would be, and better suited to the purpose of the RMS. A backup control system for the RMS consists of individual joint drive switches, the simplest system described above.

The ability of a control system to do real-time geometrical transformations permits another refinement, known as display-referenced control. In this scheme, the control system uses the current orientation of the primary television camera to interpret the manual input from the operator. The result is that, from the operator's point-of-view, the controls always bear the same relationship to the display. Thus, for instance, movement of the control joystick away from the operator would always produce vertical motion of the manipulator on the display, no matter what the current camera angle is.

All of these systems are still open-loop with respect to force control: the operator has no means of sensing the forces on the arm. For parts transfer this is not a big loss, as a position-control strategy is all that is needed. The shuttle RMS, for instance, is not intended for manipulation but for payload handling. In the early days of telepresence design, such manipulators were also used for tasks requiring compliant motion. In these cases, experienced operators "closed the loop" by observing deflections of the arm visually, to get a rough idea of the forces. Some passive compliance was built into these arms for that purpose.

The first teleoperators designed for true force control were master-slave manipulators which were modified to become force-reflecting. For a typical electric-actuated manipulator with revolute joints, this means sensing the currents through the motors (which are proportional to the torques, for DC motors) in the slave arm and back-driving motors in the joints of the master arm. Only a fraction of the force on the slave arm is applied to the master, to make the operator's work easier. The force-reflection idea can also be used with resolved-motion manipulators in which the forces on the end-effector are detected and applied to the operator's hand controller.

Force-reflecting ("bilateral") manipulators have cut the performance times for typical assembly tasks significantly. Figure 13 shows the general results of several studies illustrating this. The performance of various systems was measured for the same task, and compared to the reference time of an

TIME COMPARISON TO PERFORM TYPICAL TASKS

	LASL	MIT	NASA	MBA	CEA
TWO-ARMED MAN (UNSUITED)	1	1	1	1	1
TWO-ARMED MAN (SUITED)				8	
TWO-ARM SERVO/OR MECH. MANIP.	8	8-10	8	8	2-8
ONE-ARM SERVO/OR MECH. MANIP.	16		16		
ONE-ARM EMM (POSITION CONTROL)	80	40-50	64	55	10-30
ONE-ARM EMM (SWITCH CONTROL)	480	80-100	640	55	50-100
CRANE (IMPACT WRENCH)	>500	>100	>600	>500	>100

FIGURE 13

(FROM OAK RIDGE NATIONAL LABORATORY PUBLICATION ORNL-WS 8314R) unsuited man (using two arms). It can be seen that the addition of force control to a single manipulator arm reduces the average task time by about a factor of 4 below that required using position control alone.

Master-slave or resolved-motion manipulators with force-reflection are a proven technology for tasks of complexity equal to that of the anticipated space telepresence tasks. One aspect of space operations presents quite a challenge to these systems, however: time delays. As previously discussed, in a space-to-ground telepresence loop, round-trip signal time delays may be as long as two seconds. Several studies have investigated the effect of time delays on various man-in-the-loop manipulator control schemes.

For purely position-controlling manipulators, investigators at M.I.T. have found that, with delays of 0.3 seconds or more, the operators spontaneously adopt a "move-and-wait" strategy. This involves moving the master arm to a best guess for the desired position, then waiting out the time-delay interval to see the results of the move. This process is repeated until the task is completed. The number of "waits" involved depends on the complexity of the task. It was found that this strategy was effective in accomplishing the tasks, although errors were more frequent than in the case of no delay (particularly for complicated tasks). The extra time needed to accomplish the task with delay was repeatable and could be predicted from no-delay performance. For short delays (0.3 seconds in this test) some attempts were made to sustain continuous movement (thinking ahead, in effect), but these results were slower and more error-prone than the move-and-wait tries.

The results with force-reflecting manipulators are less encouraging. It is possible to use a move-and-wait strategy, as with position control, but the major advantages of force control are lost. When driven into an immobile object, even at slow speed, the manipulator arm can generate large forces before the operator is informed of the contact and can take corrective action. Also, when the force information is presented directly to the controlling arm, time delay can cause a serious instability problem. Figure 14 shows a graph of typical unstable control movements following a small disturbance.

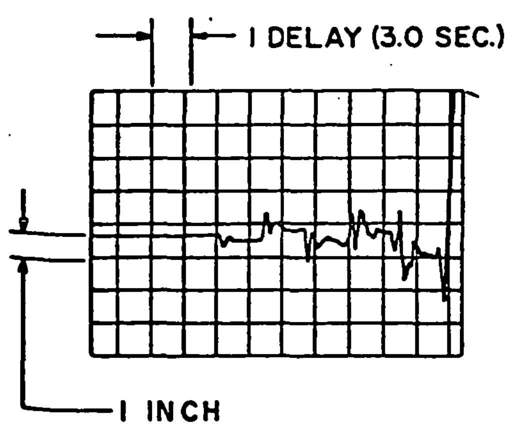

Figure 14: Unstable Control Movements

From Ferrell, W.R., "Delayed Force Feedback," Human Factors, v. 8, pp 449-455, October 1966.

Unexpected disturbances are more likely to cause unstable response. The instability can be reduced by decreasing the level of force fed back to the operator, but this reduces sensitivity and does not eliminate the problem. One way to get around this difficulty is to present the (delayed) force information to the operator's idle arm, or display it in another form, such as visual or auditory. These alternate forms of display are not as natural for the operator, however, and though they provide stability they cannot compensate for the operator's basic inability to close the force-control loop when there is a time delay. The performance of such a telepresence system is limited to that of a position-controlled manipulator using a move-and-wait strategy.

The prospects for using a classical direct-driven manipulator for assembly-type tasks in the presence of a significant (tenths of a second) transmission time delay are poor. Several ways have been suggested to cope with this problem. The simplest is to introduce passive compliance into the manipulator arm. This limits the forces generated by collisions between the arm and the worksite environment. It also gives the system a tolerance for position errors during compliant motion. The tolerance is fixed by the design and must be a compromise between the rigidity desirable for some tasks and the compliance needed for others.

With the inclusion of a processor at the manipulator, more sophisticated methods can be used. Automatic adjustment of forces can be performed by the onsite processor, which receives sensory information and controls the arm with no time delay. The simplest application is to limit the forces to a preset level. When the arm sensors indicate that the force limit in some direction has been exceeded, further commands to move in that direction are inhibited, and the arm is controlled to regulate the force to that level until a command to move in the opposite direction is received. The force limits can be set to the desired values for each task with a command from the remote control station. If slip sensors are incorporated in the manipulator hand, the onsite processor can also be used to automatically adjust grasping force to maintain a firm grip on the target object.

Passive compliance sets an overall limit on the forces exerted on the environment due to small position errors of the manipulator. An onsite processor allows this limit to be changed at will. The logical extension of this idea is to close the force-control loop in the onsite processor. In this scheme the force information is not fed back to the operator. Instead, he uses a force-sensing hand controller rather than the usual position-sensing type to directly specify the desired forces at the end-effector. The onsite processor then adjusts the position of the manipulator until the desired forces are obtained. For motion in a direction which is not constrained by the environment, the onsite processor would limit the arm velocity to a value proportional to the commanded force.

These ideas ameliorate the adverse effects of time delays by reducing the penalties associated with the errors that are made. A more attractive approach is to reduce the number of errors that occur, and their magnitude. Only then will the system be as effective as one with no delay.

With the addition of more computing power at the control station, a predictive display becomes feasible. This idea is basically to fool the human operator into producing the inputs he would give if there were no time delays. Such a system begins with an accurate dynamic simulation of the manipulator arm, using the equations of motion. If the delay is two seconds, for example, the simulator would have accurate two-second-old information on the state of the arm, as well as a record of the inputs since then. From this, a running estimate of the current state of the arm is computed. This can then be referred to the point-of-view of the primary TV camera and a line drawing of the arm generated for display on a screen. If the simulation is accurate, this eliminates the need for a move-and-wait strategy with parts transfer (unconstrained) tasks.

The simulation can be improved by including the effects of the environment, such as keeping track of the mass of any object being carried by the manipulator. The next level of sophistication would involve a world-model of the worksite environment, including the relevant parameters -- dimensions, masses, locations, etc. -- of all the objects therein. The information required for this world-model may come from design data on the satellite being serviced, in which case it would be preprogrammed into the simulator. Alternatively, the world-model can be generated and maintained in real-time by a computer vision system, which would analyze the TV images and combine this information with other available data (from a laser range-finder, for instance) and the original design information.

A good world-model, combined with a faithful dynamical simulation of the arm, could produce an accurate prediction of the manipulator state during compliant motion. Predicted forces would be fed back to the operator, just as in the usual master-slave or resolved-motion system with force-reflection. If the fidelity of the prediction is sufficient, a move-and-wait strategy would be unnecessary for any type of task. The magnitude of the positioning errors which occur would be reduced significantly, so that the remaining error could be handled by a small amount of passive compliance.

Such a system, while simple in concept, requires some sophisticated techniques from the fields of computer science and artificial intelligence. The key issues in this area are computer vision and knowledge representation, which will be described in detail later in this report. At this point, we will simply note that no complete system has yet been demonstrated.

Much of the research in artificial intelligence is aimed at producing autonomous systems, capable of planning and executing complex tasks on their own. The control of a manipulator arm is one of the traditional problems addressed. Some of this work is finding applications in telepresence now, and its role is bound to increase.

The autonomy of a telepresence system is the degree to which it can function independently. Increasing the autonomy has two goals: to reduce the operator's workload, and to improve performance. For space telepresence in the near future, operator workload and productivity are not the critical items, since the actual manipulation time has very little effect on the cost of a mission. Performance capability, however, is crucial; particularly in the presence of degrading factors such as time delay. In a typical space telepresence application, a large investment hinges on the successful outcome of the manipulation, and anything which increases confidence in the system is worth quite a bit.

Autonomy is increased by removing the human operator from the lower levels of the control hierarchy. The executive level of control is performed automatically in all telepresence systems considered here, simply because the distances involved are too great for direct mechanical linkage. The tactical level of control is routinely automated in the resolved-motion controllers already discussed. These controllers are given the desired motion of the end-effector, which they then distribute to the individual joint controllers of the arm. The next level of autonomy is to replace the human in strategic control, which takes an overall plan of action and derives the sequence and timing of the individual motions required. Systems with some autonomy at this level are known as supervisory control systems, since the human takes a back seat during whole sequences of movements.

Supervisory control schemes can be divided into two classes: traded control and shared control. Traded control systems are the most common and will be discussed first.

Traded control implies that, at any given time, either the human is directly controlling the strategic level of manipulation, or a computer is. Typically the human would define a subtask for the computer and it would take over for a while, with the human maintaining control only in the sense that he could interrupt the routine, at will, and resume direct control. The complexity of the subtasks allowed and the detail in which they must be defined indicate the sophistication of the computer system.

Even a conceptually simple system can be a great asset to the operator. One such system, called MMIT, has been assembled at MIT for dealing with problems encountered by the Navy in their underwater manipulations. To use this system, the operator defines a set of points in space that he wants the manipulator's tip to pass through. This can be done textually, with a keyboard (using some predefined coordinate system), or by directly controlling the arm and indicating the desired points to the machine (by demonstration). When the set of points is complete, the computer generates a trajectory which passes through them in sequence, and stores this trajectory for execution when desired.

An example of a task benefitting from such a system is cleaning sediment from a weld with a water jet. The moment the jet is activated, the surrounding water will become murky with silt, making it impossible to see the weld and follow it with the jet. With supervisory control, however, the path can be defined in advance while the water is clear and then executed automatically when the jet is turned on.

Industrial robots provide another example of this type of supervisory control. The trajectories are defined textually or by manually moving the manipulator arm. They can then execute the same motion repeatedly with only occasional human supervision.

More flexibility is attained by a system which can alter its behavior depending on sensory information. It can be given an instruction to "rotate wrist clockwise until torque equals ten N.-m.," for instance. This allows the supervisory control system to perform tasks which require active compliance. With the capability to react to force or tactile sensory data, a very sophisticated supervisory control system can develop, beginning with a vocabulary of simple task elements. The simplest elements may include tightening a bolt that the end-effector is grasping, or exchanging one end-effector for another in a rack. Quite complicated tasks can be specified by combining these simple elements into procedures (essentially computer programs written in a manipulator-oriented language).

For space telepresence the computer could be located at the worksite, avoiding time-delay problems and reducing the amount of communication to the control station. To deal with the operator's time-delay using supervisory control, the procedures defined should at least be comparable in length to the time-delay. This enables the system to operate without intervention for the period between when the command is given and when the operator can see the results. Supervisory control decreases the frequency at which the operator must command the system, reducing the time wasted in waiting for return signals, thus speeding up the whole operation. The onsite processor can also react more quickly to a developing problem (if it has been programmed to do so) than the remote operator could, and minimize the consequences.

The programming of complicated procedures and task vocabularies would take place long before the required manipulation, to allow time for checkout on ground-based simulators. This information could be programmed into the manipulator before launch, or for a space-based system, uplinked over a period of time prior to the specific mission. Simple procedures could be defined during the manipulation, as with the MMIT system, when needed by the operator.

Supervisory control systems such as these have limitations. The definition of a broad task vocabulary is a formidable programming job. To create a program for a very complicated task, taking into account all of the possibilities that may occur, would take the operator far longer than simply performing the manipulation himself, even with time-delays. Unless the procedure is to be executed many times, it is not worth the trouble.

The addition of computer-interpreted vision is the logical next step in autonomy. This is a big step, and it requires reorganization of the system around a world-model. A world-model, as discussed previously for predictive display systems, contains descriptions of all of the objects in the workspace (including the manipulator) with their interrelationships and all of the parameters relevant to manipulation. For example, the description of an access panel would include its location and orientation on the satellite, and the locations of all of the bolts which secure it, the size of the bolts and the direction they turn, the location and degrees of freedom of any hinges, etc. In short, all of the information a human would use to perform a manipulation. The world-model is a representation of the visual data (as well as data from other sources) in a form usable for manipulation. A world-model is hard to construct and maintain without vision data, and vision data is hard to interpret without a world-model, so the two components are complementary. A more detailed exposition of the relationship between them (also known as the high-level and low-level aspects of vision processing), and the current state-of-the-art of such systems follows in the section on Computer Vision.

A telepresence system with an internal world-model can begin to take over some activities at the highest level of control: the planning level. This may be thought of as an extension of supervisory control, but the distinction is important and these systems will be referred to as planning systems.

Implicit in the world-model are the tasks that can be performed. The goal of each task can be expressed as a state of the world-model, just as the initial state can. In order to plan the manipulation to get from the initial state to the desired state, the computer needs a set of rules, or reasoning tools, which allow it to predict the effects of its actions. For a telepresence system, the rules would embody the equations of motion of the manipulator arm and its interaction with the environment.

Computer systems capable of inference using a set of rules and a data base (world-model) are called expert systems. The development of such systems is a well-established field in Artificial Intelligence research. In this case we require a system which is "expert" in the dynamics of a particular manipulator (the one which it controls). With this capability, the telepresence system can "intelligently" interact with the worksite environment to attain its goal.

The flexibility of an autonomous system can be increased by improving its ability to learn (or adapt). The simplest planning systems can learn, in the sense of modifying their world-model based on sensory information. A more powerful type of learning would enable the system to modify its own reasoning tools. Its rule base could be changed, based on experience and deliberate experiment. This would simplify initial rule base programming and give the system an ability to deal with unforeseen circumstances and malfunctions. Ultimately, the software could be "trained" for each mission by simulation, much as humans are now.

To use an autonomous planning system, the human operator would supply to the machine (textually) essentially the information he would need himself. A satellite repair mission, for example, would require design data on the satellite and a description of the repairs. With computer-aided design (CAD) systems becoming common, the design information may already be available in machine-usable form.

Very little communication would be needed between the computer at the worksite and the human operator at the control station during the manipulation. The human is relegated to a supervisory role throughout complicated tasks. Such a telepresence system could approach or exceed the capabilities of a human present at the worksite. Autonomous systems are obvious choices for exploration missions, where communications are a problem and large time delays prohibit direct human control.

Fully autonomous planning systems are the logical limit of traded control telepresence, and are still in the early stages of development. An alternative to traded control is shared control, which makes use of some Artificial Intelligence ideas but doesn't require the sophistication of full autonomy. With shared control, the human operator gives the computer a description of the goals, in world-model terms, as well as direct manipulation input at the strategic level, using a hand controller. The computer modifies the direct commands as necessary to conform with the higher-level plan. The need for modification may come from the existence of time delays or just misjudgments by the operator. High-level information can come from textual input or from a world-model maintained at the control station, designed to simulate key features of the work-site, and manipulated by the human operator.

This world-model does not need to be complete or particularly accurate, since the operator's commands are not directly controlling the manipulator. The world-model at the control site is derived from sensory information, but is simplified. Objects may be represented by simple geometric shapes, and their locations need not be precise. The human is presented a graphic display of this world-model, and he directly controls a simulated manipulator in it. The manipulator simulation is again simple, using linearized, approximate equations of motion. The human operator performs the desired task in this world-model, which is an idealized version of the real workspace.

The onsite computer also maintains a world-model, but this one is as accurate a reflection of reality as possible. It is from this model that the simplified version is constructed and relayed to the control station. The human operates his controller, and his inputs together with their effects on the simplified world-model are transmitted back to the onsite computer. Note that the communication (both ways) involves only the simplified world-model, requiring much less information than the fully accurate one. The human inputs are transformed to account for geometric differences between the simplified model and the accurate model, and these become the nominal control signals for the real manipulator arm. The simplified model from the control station contains the important features of the desired state of the workspace, such as "manipulator is aligned with bolt." This is the high-level information. The onsite computer compares this information to its accurate world-model. If there are no discrepancies, the nominal signals are used, unmodified, to control the arm. If reality diverges from the plan, however, the system uses its manipulation rule base to correct the control signals. The rule base does not need to be as comprehensive as one in an autonomous planning-level system, since deviations from the desired state will be detected as soon as they occur, while they are small.

Since a shared control system makes use of the human operator's strategic skills, it is less complicated than a fully autonomous system. However, it offers several advantages over direct or even supervisory control. One example is in the application to a telepresence problem involving time delays, There is no operational dependence on the source of time delays in a control loop, so suppose we have a system in which the link from worksite to control station is immediate, but a two second delay occurs on the return path. The simplified control station world-model should then include a two second prediction, just as previously described for a predictive display system. The simple world-model is propagated forward from the one received, using the last two seconds of control inputs, to produce the version seen on the operator's display. Thus, the operator is working two seconds in advance of reality. His input and the world-model are transmitted back to the worksite (with time delay) where it represents a high-level description of the desired state, just as in a shared control system with no time delay. The prediction errors can be handled by the onsite computer just like any other errors, and the signal is corrected to produce the desired result.

A shared control system incorporating prediction is superior to a predictive display alone, since the force control loop is closed in the onsite computer. Also, with shared control, errors caused by prediction inaccuracies are corrected as they occur. For example, consider a case in which the manipulator is intended to pass between two closely-spaced obstacles. The prediction errs in such a way that the operator believes the movements will succeed, but in reality his commands would cause a collision with one of the obstacles. With a simple predictive display system, the collision would occur and the operator would be informed of it, after the time delay, by a sudden discontinuity in his display. On the other hand, in a shared control system, the onsite computer would detect the misalignment and correct it before the collision could occur.

Shared control is a particularly attractive idea for space telepresence, since it can cope with large time delays, yet it does not require as much development as an autonomous planning-level system. A near-term shared control system will probably be more capable and reliable then a near-term planning system, and its need for more operator involvement is not a big drawback for space applications, where the additional cost would be negligible.

Both shared control and planning systems depend on the construction and maintenance of a data base containing information about the worksite and task, called a world-model. In the next section the required technologies are discussed.

7. Computer Vision

Telepresence was first defined by Dr. Marvin Minsky (of MIT's Artificial Intelligence Lab) as the transference of human cognitive and operative skills to a remote worksite via a machine system interface. Telepresence will eventually evolve into a fully autonomous teleoperator system with the human as task-specifier and supervisor of the machine. Such a teleoperator will possess its own planning, decision-making and problem-solving skills. A computer will act as the representative of the human worker at the maintenance, construction, or exploration site.

Computer vision is a major step on the path to autonomy. It enables the computer to use for itself the greatest source of information available about the state of the worksite. Our own experience shows that vision is a powerful tool for manipulation tasks.

The input to a vision system is usually light from the worksite. For space telepresence, a television system will always be available to convey images to the human supervisor (even in a fully autonomous system), so this is the most likely source of input for the computer vision system as well. The vision system's output is a description of the worksite in terms which are useful for manipulation: the objects visible, their geometry and spatial relationships.

Certain aspects of the vision problem make it similar to another long-standing AI problem -- the interpretation of natural language. One such characteristic is the large amount of data which must be processed in real time. For a modest vision system with a 256 x 256 array of data points in the image, each containing an 8-bit number updated at the rate of 30 Hz., the machine must handle nearly 16 million bits per second. Even this does not begin to approach the resolution of human vision. The human brain devotes billions of neurons to this task, and according to a JPL study (Gennery, D., Cunningham, R., et al., "Computer Vision," JPL Publication 81-92, November 1981), "it is possible that no existing sequential computer comes within six orders of magnitude of being powerful enough to see as well as a human being." Vision processing demands efficient algorithms.

Another similarity to the language interpretation problem is the necessity of having prior information. To assign meaning to speech, for instance, knowledge of the language used, the meaning of accents, inflections, and idioms is essential. In addition, it has been found that all sorts of extraneous information is required to resolve the ambiguities which commonly occur in human speech. To correctly assign adjectives to their nouns in a complicated descriptive sentence, for example, the computer often needs to know something about the object's properties, in order to select the most likely alternative. In vision processing, it needs to know the types of objects that may be seen, and how their projected images depend on distance, orientation, lighting, etc.

These features make vision processing (as well as language interpretation) a difficult task. A basic AI question is how to best fit computational structures to a given problem domain. This question is unsolved in general, but intensive study of the vision problem has produced a variety of approaches which have been effective in reducing the computational load to manageable levels for some applications. Vision systems are already being used in industry, enabling robots to recognize parts and locate them for manipulation. More sophisticated experimental systems abound in research laboratories.

The functions of a vision system will be described sequentially, starting with the hardware receiving the light from the scene. The image data produced consists of a two-dimensional array of pixels (picture points). Each pixel corresponds to a small area of the focal plane, and represents the value of one or more quantities in that area. Typically the quantity represented is intensity of light, although other types of information can be used. Depth information from a laser rangefinder scanning across the scene is one example. Each pixel can be associated with a single number (usually 8 bits), as in monochromatic (black and white) television, or several. For color vision the quantities may be intensities in each of two or three primary colors, or an alternative group of characteristics known as hue, saturation, and brightness.

The pixels are usually arranged in a rectangular or hexagonal array. The numbers are updated by sensors (30-60 times a second, for TV). The distillation of scene data into pixels represents the lowest level of vision processing. Since this is implemented in hardware, the resolution that will be needed must be known in advance and designed into the machine.

The next set of procedures is known as the low-level processing of image data. A variety of algorithms are grouped under this heading, all of which extract relevant feature information from the image. The output of this level of the vision system is usually a list or a directed graph showing the spatial relationships of the image features detected. The ideal is to have simple algorithms (which can be executed quickly) reduce the huge amount of image data to a relatively small set of data about features, while preserving all of the important information. This process is called segmentation.

Many types of feature-extraction operators have been devised. Edge operators detect areas of the image with a large gradient in brightness. An edge usually indicates a depth discontinuity or a shadow in the scene. Line operators detect bright or dark lines (essentially two edges back-to-back). Another common feature for detection is a corner, where two edges or lines intersect. Texture, defined as a local variation in pixel values, is another feature which can be characterized.

Features like edges, lines, and texture are detected using a "window" (usually a 3 x 3 or larger array of pixels) scanned across the image. The change in the pixel values within the window determines the existence, direction, and magnitude of a feature at that location. Windows of different sizes and shapes can also be convolved with the image data for smoothing and feature enhancement.

Region growing is the next step in image segmentation. Edge- and line-follower algorithms piece together continuous boundaries and discard isolated edges. Region growing groups together adjacent pixels which share common properties of brightness, color, texture, or other features. At the end of this phase, the entire image has been reduced to a set of regions, each with their own set of characteristics. A directed graph can then be constructed to express the geometric relationships between regions. The graph may indicate, for instance, that region A contains region B, and regions C and E are non-overlapping subsets of region B.

High-level vision processing relates the image data to objects in a world-model. Most of the variation between different types of vision systems is in their usage of high-level information. A general-purpose real-time vision system requires sophisticated techniques to reduce the computational burden to a manageable level, while a system for performing a specific task in a controlled environment can be much simpler. A typical problem in industrial robotics may require locating a part of known shape on a flat table, with the camera position fixed directly overhead. In this case the simplest approach might work: an exhaustive test of all the possible object orientations, checking to see if they match the image. In a situation with many possible objects, having arbitrary orientations and distances from the camera, the time required for this approach becomes prohibitive; some technique must be used to reduce the number of comparisons to be made.

The high-level part of a vision system is also called the recognizer or classifier. It is guided by a database which contains descriptions of the objects it is designed to recognize. In an application-oriented system for telepresence, this might include the components of the satellite to be examined, the manipulator arm itself, and, for other objects not characterized in advance, a set of generalized shapes. These would typically be blocks, cones, or ellipsoids with parameters which can be adjusted. Any object not specifically recognized could be represented by some combination of these generalized shapes.

Recognition consists of matching object descriptions in the database with features in the image. As many methods for doing this exist as there are researchers in the field. The most common approach is to give each feature in turn a "likely" interpretation, then examine the whole set of these assignments for compatibility. The assignments are then revised and checked again until certain confidence criteria are met. The final set of feature identifications is then completely compatible with a three-dimensional interpretation of the scene, and should represent the most likely possible interpretation.

The "likelihood" function, which is used to give features their initial assignments and to compare alternatives, is the most important part of this method. If initial assignments are made which are nearly all correct, the vision system will run much faster than if many iterations are necessary. For this reason, an efficient vision system makes use of prior information wherever possible (as humans do). In a continuously-operating system, for instance, the scene may change little between one image and the next. In this case an initial set of assignments can be quickly derived from those of the previous image.

Tracking and verification are two types of vision problems which exploit the existence of prior information. Tracking refers to the continuing detection of one or more moving objects against a fixed background, while verification implies that a model for the scene already exists and must be checked.

The simplest vision systems have a fixed likelihood function which is predetermined by the programmers based on the anticipated scenes. For instance, in an industrial robot system where parts are all laid out flat on a table: assume that each part has one 135° corner. The high-level part of the vision system would then automatically assign that corner of a part to each 135° corner feature detected by the low-level system. The compatibility check would subsequently eliminate any erroneous assignments caused by, for instance, the edges of two adjacent parts making a 135° angle.

More sophisticated and general-purpose vision systems must be able to change their own likelihood function based on the state of the scene. Ultimately, if the computer's world-model has an understanding of the dynamics of the worksite, the state of the worksite can be predicted from knowledge of the previous state, and the vision problem reduces to verification. If the features are correctly identified on the first iteration, no more are necessary.

This procedure is an example of the highest level (the system's world-model and dynamical knowledge-base) helping out a lower level (the high-level part of the vision system). Communication and cooperation between levels can be an effective technique for reducing the amount of computation required. For instance, the high-level part of the vision system could reduce the amount of low-level processing required by directing the feature detection algorithms to look only for certain key features. The high-level system selects these features based on its ongoing recognition attempts. One example would be a vision system looking at an access panel. When one corner has been tentatively identified as the corner of a panel, the high-level system would direct the feature-detectors to look for the other corners and the latch. If the access panel design is familiar to the system, the low-level processors can be given the most likely places to look for these features. By saving the time needed to detect and classify all of the edge and texture information for the area of the access panel, the vision system can operate much faster. In a telepresence system with some autonomy, the highest level of the computer deals with the manipulation goals. Knowing these goals, it can direct the vision system to concentrate on the portions of the scene which are most important at that time, to give high-resolution information where it is needed most, and spend less time working on the rest of the scene.

In this way, a certain amount of high-level processing can be traded for quite a bit of low-level processing. There is a limit to this, which is reached when the computational time needed to predict the existence and location of a feature exceeds that needed for a low-level feature operator to scan the whole region of interest. In general an optimal distribution of computation between all of the levels exists, involving communication both up and down the hierarchy.

Figure 15 summarizes the processing sequence for the low-level and high-level aspects of a machine vision system.

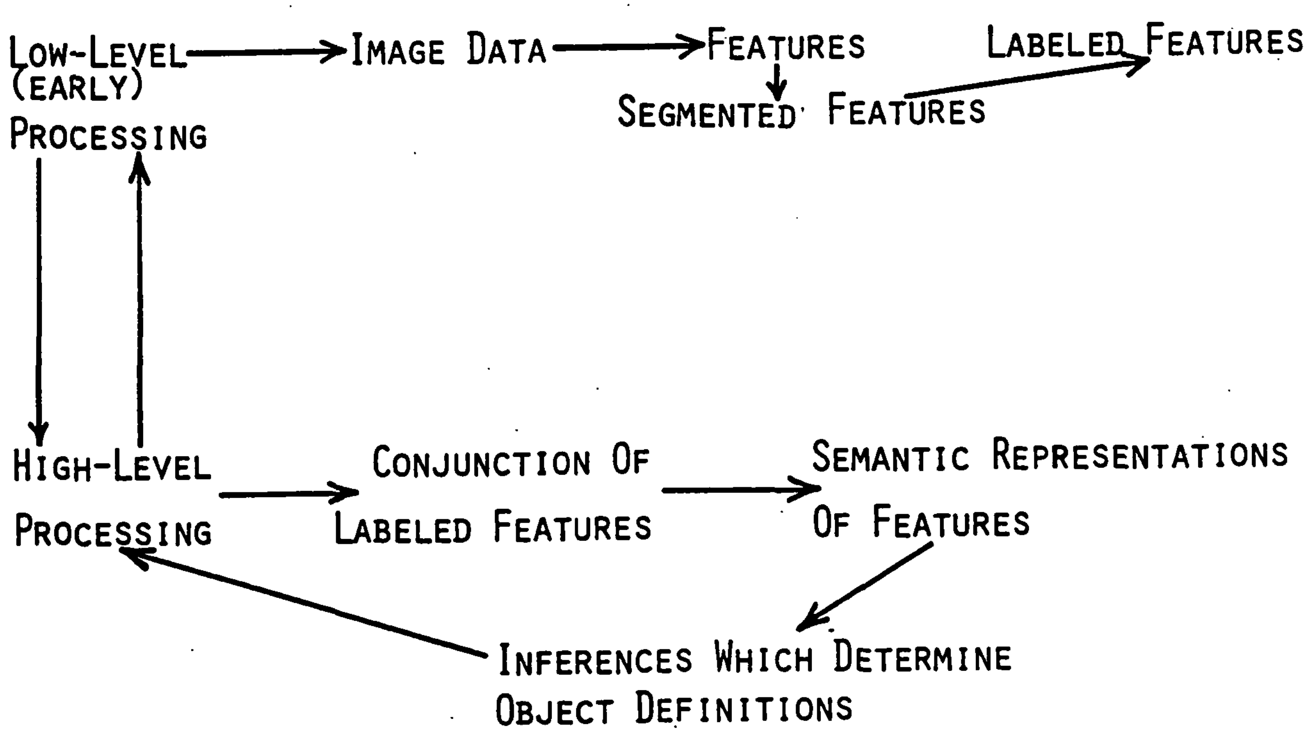

FIGURE 15

Figure 16 shows the arrangement used in a particular vision system application. The hierarchical algorithm depicted is used in the detection of tumors from radiograph images of a lung. Technical details of the method can be found in the original source.

	PREPROCESS	SEGMENT	CONTROL
Pass 0 (Digitize radiograph)	The digitizer has a hardware attachment which produces the optical density.		
Pass I (Find lung boundaries)	In 64 X 56 consolidated array, apply gradient at proper resolution	In 64 X 56 array, find rough lung outline; in 256 X 224 array, refine lung outline	TOP–DOWN
Pass II (Find candidate nodule sites and large tumors)	In 256 X 224 array, apply high-pass filter to enhance edges, then inside lung boundaries; apply gradient at proper resolution	In 256 X 224 array use gradient-directed, circular Hough method to find candidate sites; also detect large tumors	BOTTOM–UP
Pass III (Find nodule boundaries)	From 1024 X 896 array, extract 64 X 64 window about each candidate nodule site, then in window apply high-pass filter for edge enhancement; then apply gradient at proper resolution	In 64 X 64 full-resolution, pre-processed window, apply dynamic programming technique to find accurate nodule boundaries	TOP–DOWN

FIGURE 16: A HIERARCHICAL TUMOR-DETECTION ALGORITHM

FROM D.H. BALLARD, C.M. BROW, "COMPUTER VISION", PRENTICE HALL, INC., ENGLEWOOD CLIFFS, N.J.

Quite a bit of work has been done on all aspects of the vision problem (see bibliography). Since there are several phases of vision processing which can become computationally very expensive, a workable vision system for a given application tends to be as simple as possible, and "cheats" by using prior information wherever it can. For this reason no truly general-purpose vision system yet exists, but it is possible to deal with a restricted problem domain (such as maintenance work on a known satellite design) with accuracy and reliability. Careful design of the workspace can considerably enhance the performance of a machine vision system by, for instance, judicious use of color-coding and surface patterns.

A few words about Artificial Intelligence in general are not out of place here. A.I. has been characterized by Dr. H. Simon as, "the science of weak analytical methods." Moreover, A.I. is an empirical assembly of analytical methods for the symbolic representation of problems and corresponding computational procedures for establishing optimal problem solutions. A.I. then is an assembly of analytical methods, out of which some synthesis gives rise to what we call reasoning in a cognitive sense. (If this sounds a bit like Alchemy you are beginning to get the true flavor of A.I. as it stands). A machine system demonstrates intelligent - reasoned behavior by systematically constructing for a given problem domain, a problem representation and some corresponding methods for generating a problem solution(s).

To say that A.I. is a science composed of weak analytical methods points to the fact that A.I. is an empirical science in search of a formal theory for its unification. The statement does not mean to imply that the analytical methods themselves are weak due to some logical inconsistency inherent in the methodology of analysis. Simply, it is the case that A.I. is not a unified science at this time. A partial unification may evolve from a close examination of A.I. problem solving skills as they interact with human problem solving skills in the context of the Man-Machine interface for space telepresence/teleoperation.

So what can A.I. do to minimize the complexity of problems which the human will encounter when employing the Man-Machine interface to perform space operations? Smart machine systems will reduce the complexity of operational problems by the systematic application of its methods in such a way as to give rise to reasoned solutions to complex problems. A.I. may provide a complement to human problem solving capabilities in the unforgiving space environment.

8. Summary

In a telepresence system, the requirements for the man-machine interface depend on the capabilities of the other components (including the human) and on the nature of the work to be done. The mechanical and control design constraints have been described in sections 3 and 4. Human characteristics were discussed in section 5, and the anticipated tasks described in section 2.

Overall system architecture is directly related to the interface design, since several different levels of control may be required from the operator. Four basic types of telepresence system architecture are depicted in Figures 17-20, covering the spectrum of arrangements discussed in section 6.

The first (Figure 17) is the simplest control structure, with a direct link between the control input device (CID) and the manipulator servos. The CID often takes the form of a master arm, which the operator controls and the manipulator servos are slaved to. The sensors shown in the figure include proprioceptors (joint sensors) and exteroceptors (such as proximity sensors). This type of system comprises virtually all of those in actual use (as opposed to experimental efforts) with a man in the loop.

Figure 20 shows a supervisory control system. Two processors are shown; for a simple version only one is necessary (the one at the control station). If time-delays are present

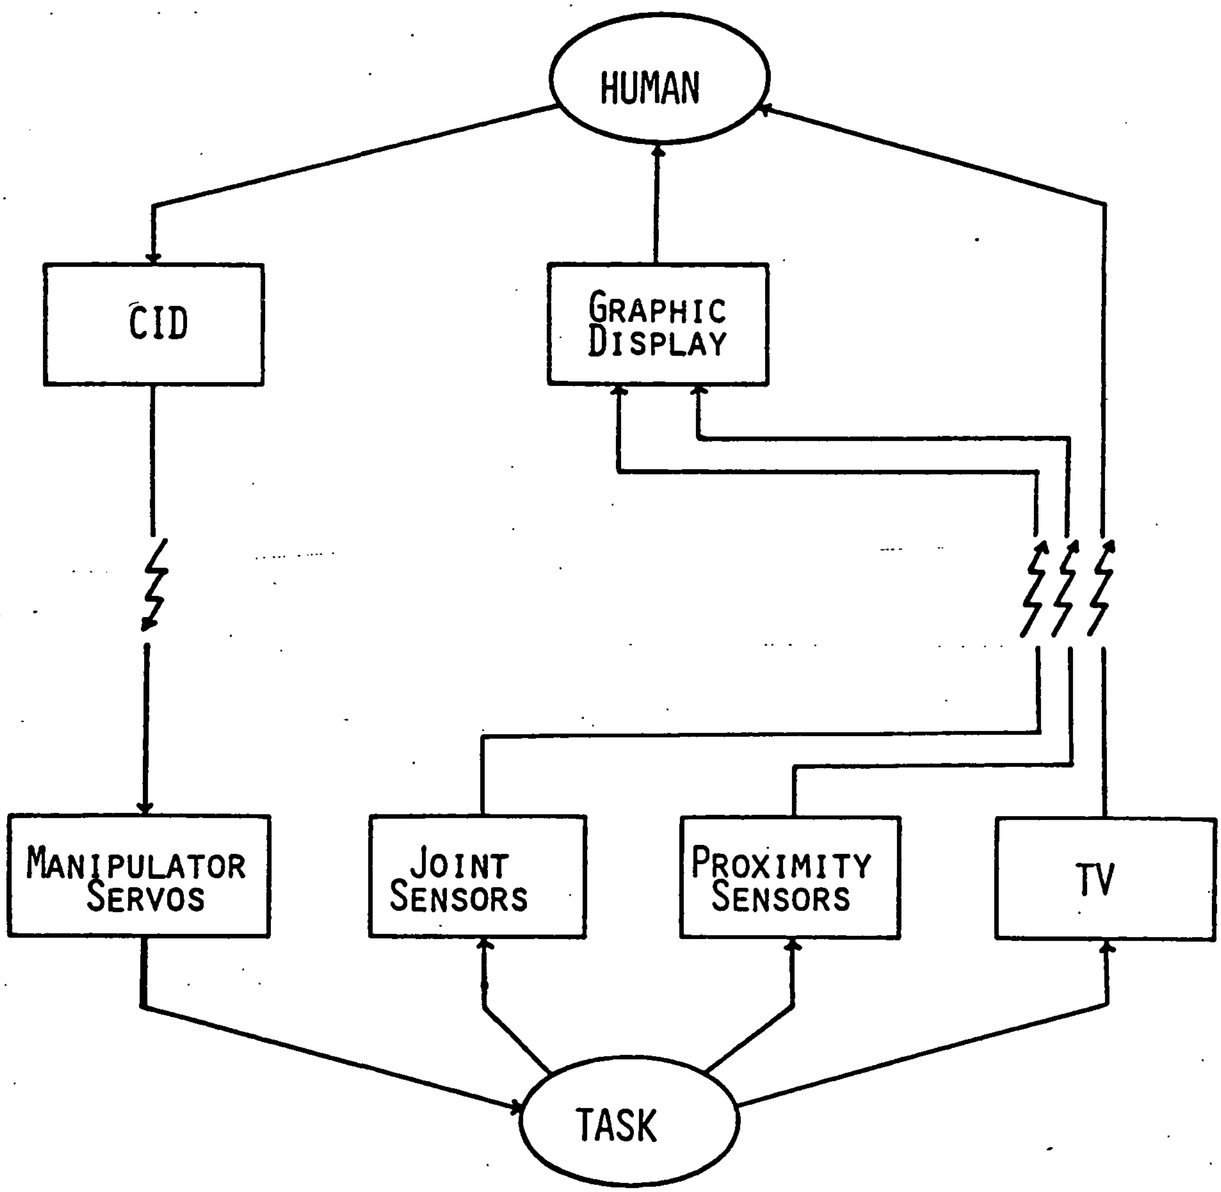

FIGURE 17: DIRECT CONTROL OF MANIPULATOR

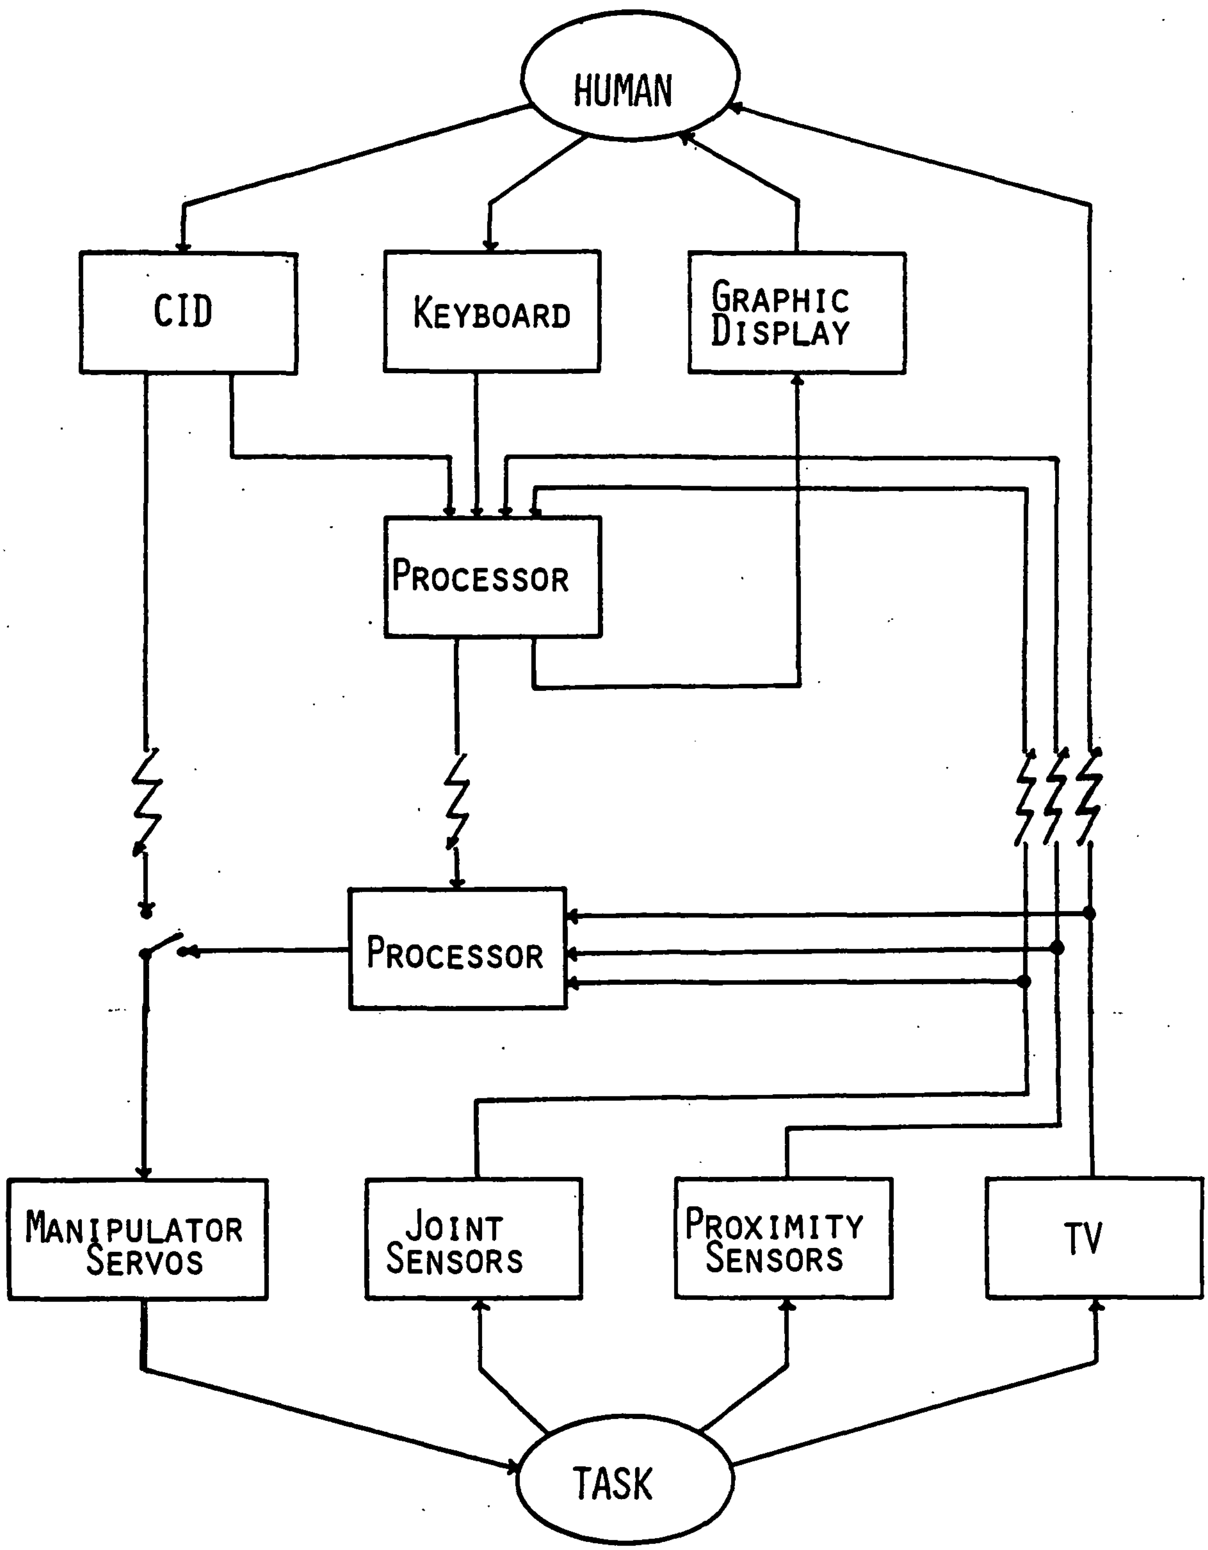

FIGURE 18: SUPERVISORY CONTROL

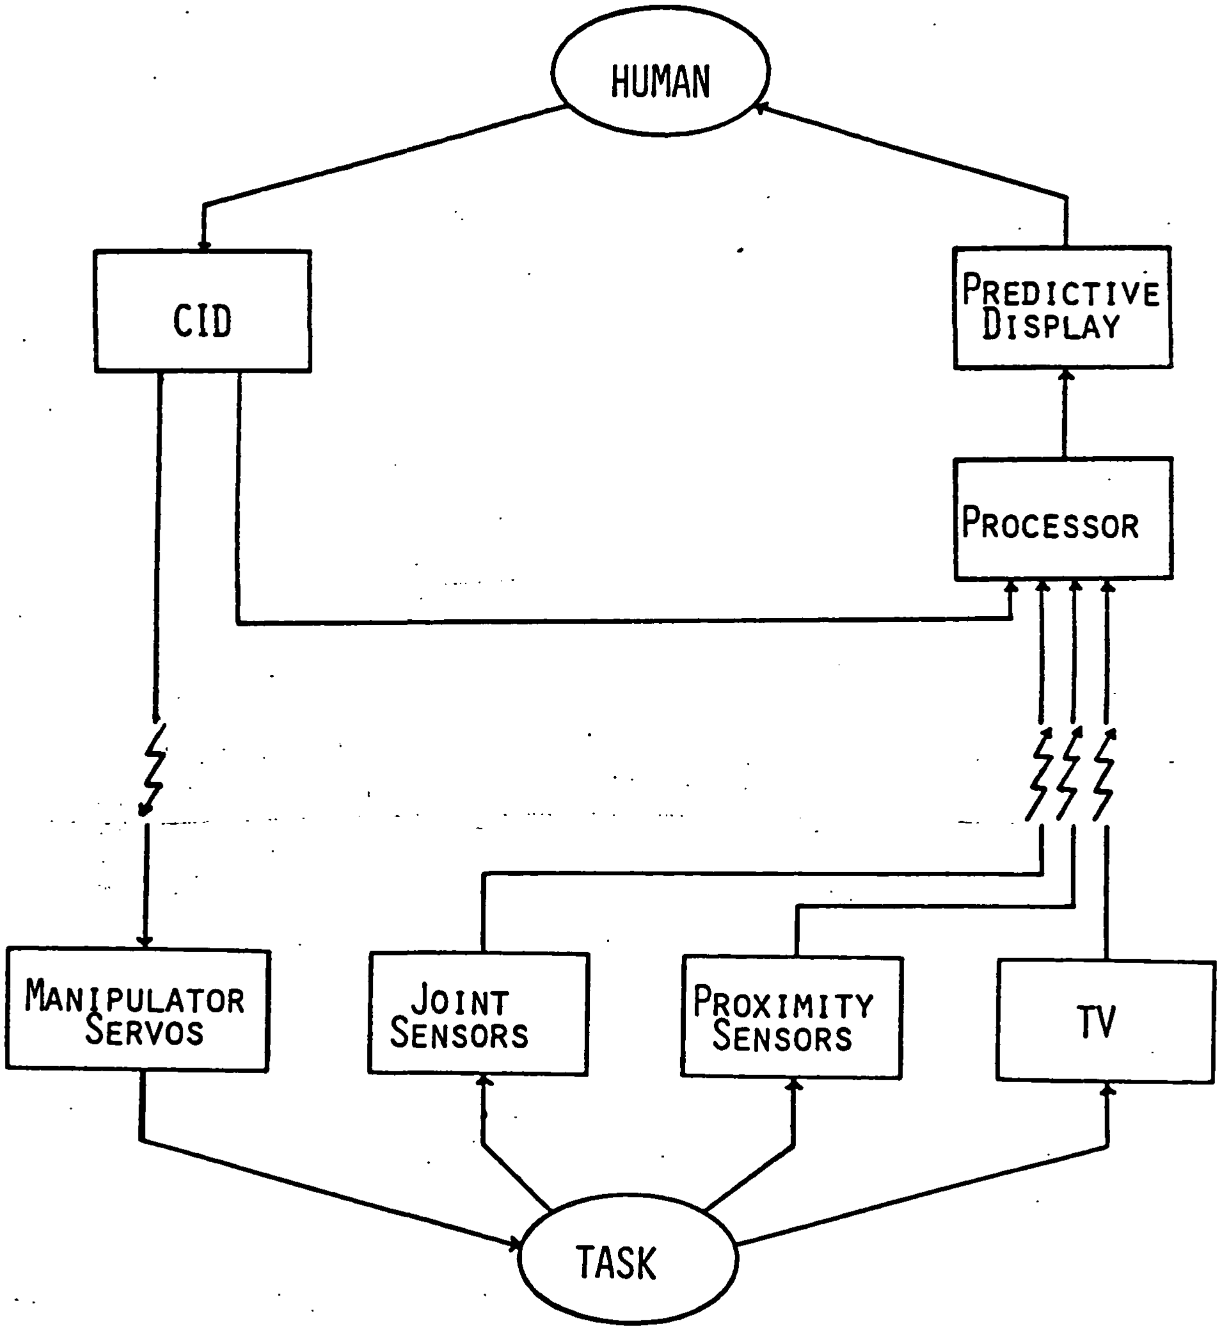

FIGURE 19: DIRECT CONTROL WITH PREDICTIVE DISPLAY

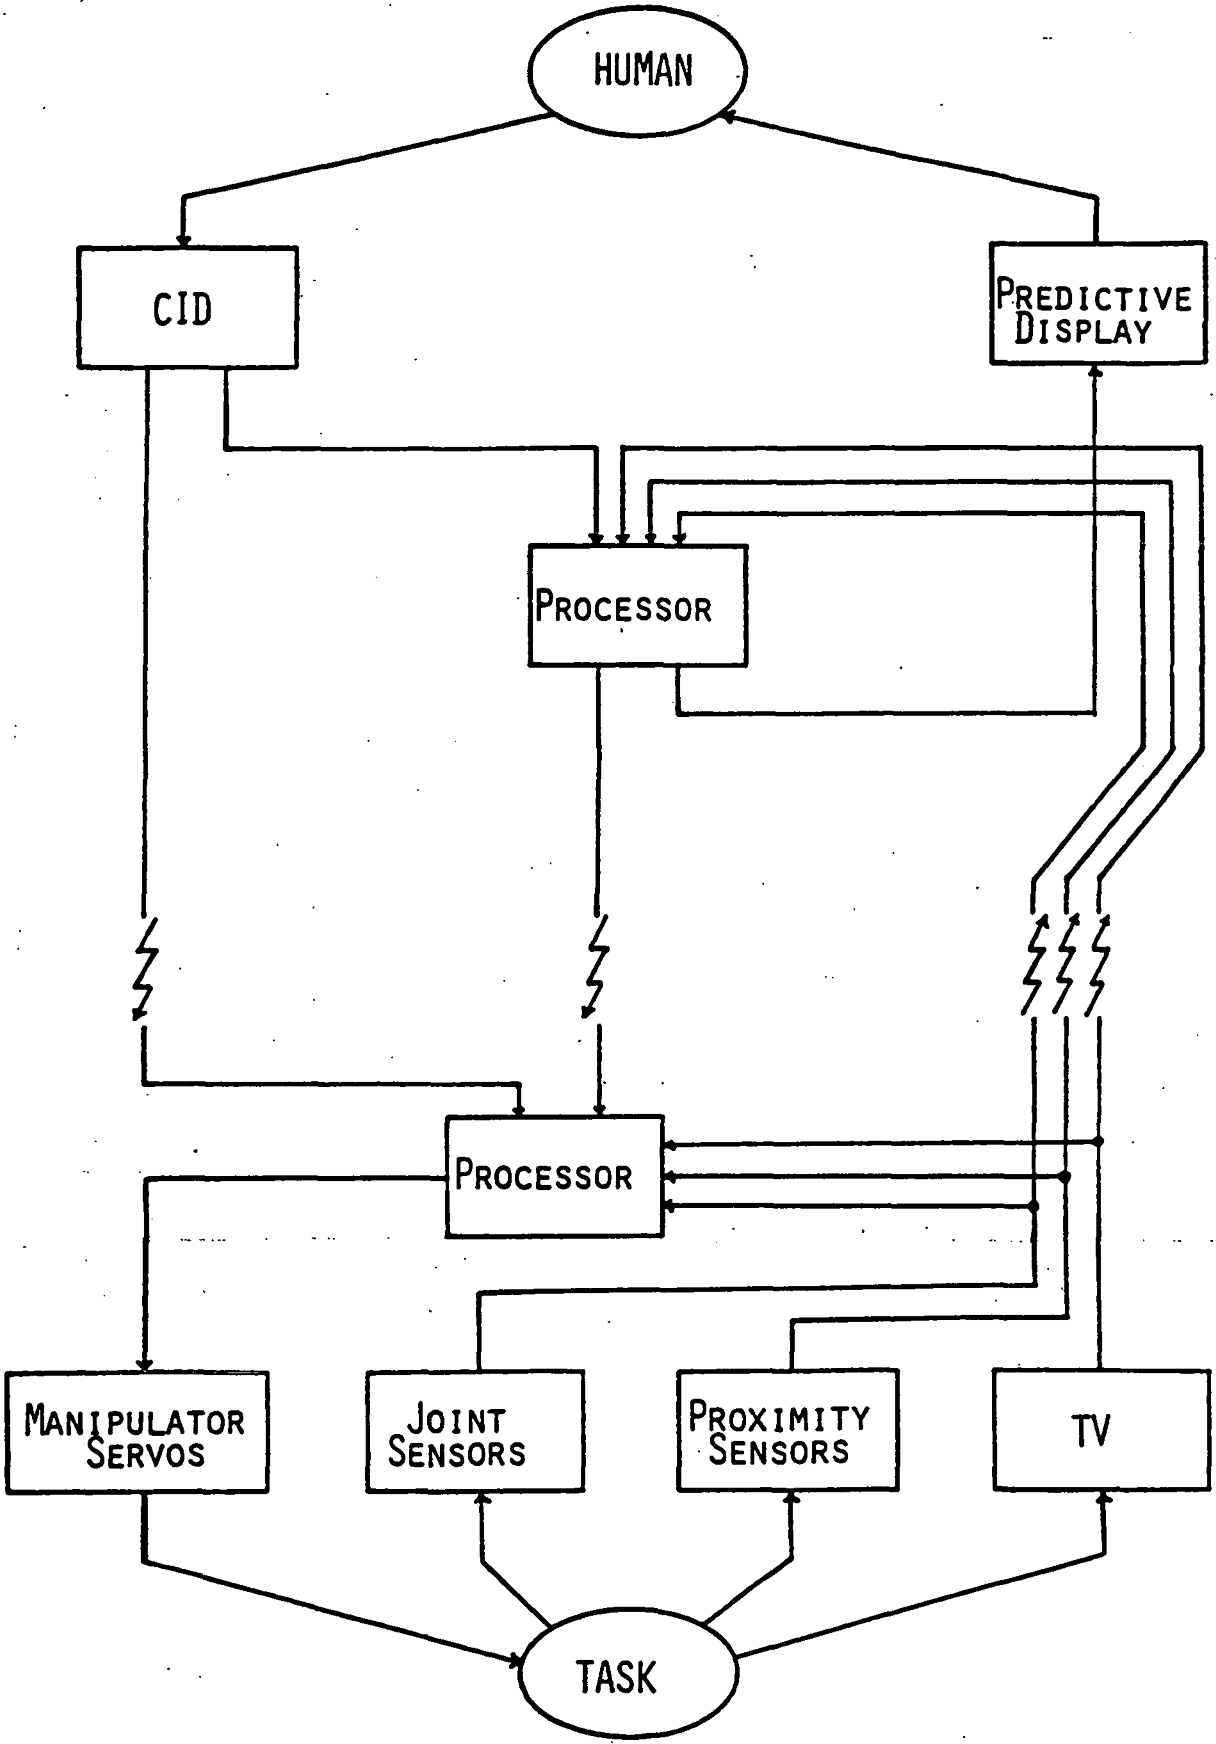

FIGURE 20: SHARED CONTROL WITH PREDICTIVE DISPLAY

or communications costs are high the second processor (onsite) is a useful adjunct. When the supervisory control system is in continuous control, the manipulator is said to be autonomous. Industrial robots are examples of such an autonomous system, capable of dealing with a very restricted problem domain. Ultimately, fully autonomous systems may exceed the capabilities of man-in-the-loop telepresence, though the evolution will be gradual. For most projected tasks an intermediate combination of man and machine control will be most effective.

Predictive displays are specifically intended to cope with a time-delay in the control loop, an expected feature of space telepresence. Figure 19 shows the structure of a simple version which, as discussed in section 6, allows the execution of part-transfer tasks as if there were no time-delay. A more advanced configuration (Figure 20) uses an onsite processor to implement shared control, wherein the nominal control inputs (from the human) are augmented to compensate for prediction errors. Such a system could accomplish assembly tasks despite the presence of time-delay.

Most of the systems mentioned use some artificial intelligence (AI) technology to complement human capabilities, particularly for precise and repetitive tasks. Manipulator control is a problem of tremendous magnitude from the AI point of view. Heuristic methods must be used rather than exhaustive algorithms, for controlling complex behavior. The difficulty is reduced if the immediate subgoal is close to the current situation, as with shared or supervisory control.

As small computers become more powerful and space telepresence tasks more demanding, AI techniques will increase in importance, taking their place alongside control theory and kinematics in the standard repertoire of the design engineer. This will have particular impact on the human factors aspect of telepresence, enabling machines to perform functions traditionally reserved for the human operator. A more efficient merger of man and machine will be the result.

9. Bibliography

Task Definition

Garin, J., "REMOTEX and Servomanipulator Needs in Nuclear Fuel Reprocessing Plants," Proc. Workshop on Remote Maintenance in Energy Systems, March 1981, University of Florida.

Description of advanced manipulator development at Oak Ridge National Laboratories.

Grimson, J., "Applications of Robotics in Accelerator Research Facilities," Proc. Workshop on Remote Maintenance in Energy Systems, March 1981, University of Florida.

A brief summary of the state of the art, current needs, and directions for future development, particularly at FermiLab.

NASA, "AXAF Utilization of STS and Satellite Servicing," NASA presentations at review panel meeting, April 1982.

Includes review of past satellite servicing experiences, economic justification for servicing, and descriptions of possible servicing schemes.

Plantier, M., Bodmer, R., "Teleoperation and Automation: A Survey of European Expertise Applicable to Docking and Assembly in Space," Final Report ESTEC Contract #4402/80/NL/AK(SC), EUROSAT S.A., May 1981.

Shows the need for teleoperation in space, from the European point of view, and gives an overview of current technology.

Trucks, H.F., "Space Telescope Maintenance and Refurbishment," Space Telescope Project Office, NASA MSFC.

Describes the level of maintenance allowed for in the design of the space telescope project.

Wernli, R.L., "Deep Ocean Applications of Manipulators and Work Systems," Proc. 27th Conf. on Remote Systems Technology, 1979.

Wernli, R.L., "Robotics Undersea," Mechanical Engineering, August 1982.

Description of typical tasks for an underwater manipulator and systems developed by the Navy for performing them.

<u>Manipulator Design</u>

Flatau, C.R., "Design Outline for Mini-Arms Based on Manipulator Technology," M.I.T. A.I. Memo No. 300, May 1973.

Horn, B.K.P., "Kinematics, Statics, and Dynamics of Two-D Manipulators," M.I.T. A.I. Lab. Working Paper 99, June 1975.

Paul, R.P., "Robot Manipulators," M.I.T. Press, 1981.

Vukobratovic, M., Potkonjak, V., "Dynamics of Manipulation Robots," Springer-Verlag, 1982.

Waters, R.C., "Mechanical Arm Control," M.I.T. Vision Flash 42, March 1973.

These sources discuss in detail the geometry, kinematics, and dynamics of serial manipulators.

Flatau, C.R., "Arrangement of Motions in a Robot Assembly Machine," M.I.T. A.I. Lab. Working Paper 91, November 1974.

This paper examines the advantages of parallel degree-of-freedom manipulator designs.

Whitney, D.E., Nevins, J.L., "What is Remote Center Compliance (RCC) and What Can It Do?," Draper Laboratory, November 1978.

Describes a technique for end-effector design to reduce the difficulty of performing peg-in-hole type (compliant) tasks.

Manipulator Control

Bejczy, A.K., "Distribution of Control Decisions in Remote Manipulation," Proc. of IEEE Conf. on Decision and Control, December 1975.

Bejczy, A.K., "Advanced Teleoperators," Astronautics and Aeronautics, Vol. 17, No. 5, May 1979.

Bejczy, A.K., "Sensors, Controls, and Man-Machine Interface for Advanced Teleoperation," SCIENCE, V. 208, 20 June 1980.

Heer, E., Bejczy, A.K., "Control of Robot Manipulators for Handling and Assembly in Space," Proc. 2nd IFAC-IFIP Symp. on Information Control Problems in Manufacturing Technology, October 1979.

Heer, E., "Report of the JPL Advocacy Group for Autonomous Systems Technology Research and Development," NASA-JPL, June 1981.

Heer, E., "Robots and Manipulators," Mechanical Engineering, November 1981.

These papers describe the work done at JPL on manipulators during the last ten years; proximity sensors, arm linkages, end-effector designs, tactile sensor arrays, displays, and control schemes.

Horn, B.K.P., "What is Delaying the Manipulator Revolution?", M.I.T. A.I. Lab. Working Paper 161, February 1978.

Discussion of the reasons manipulators have not yet lived up to their potential: basically, that research efforts have been too small, fragmented, and discontinuous. Though brief, this paper gives a feeling for what the technical limitations are.

Horn, B.K.P., Raibert, M.H., "Configuration Space Control," M.I.T. A.I. Memo No. 458, December 1977.

This paper proposes a manipulator control system which can cope with nonlinear time-varying behavior better than classical linear feedback methods. Dynamical coefficients are stored in a lookup table indexed by the configuration of the manipulator.

Lowrie, J.W., Fermelia, A.J., Haley, D.C., Gremban, K.D., Van Baalen, J., Walsh, R.W., "Evaluation of Automated Decisionmaking Methodologies and Development of an Integrated Robotic System Simulation," NASA CR-165975, September 1982.

This report touches on many aspects of the design of a current-technology space teleoperator, including joint design, dynamics, and control.

Mason, M.T., "Compliance and Force Control for Computer Controlled Manipulators," M.I.T. A.I. TR 515, April 1979.

Presents a theory of force control for compliant motion based on formal models of the manipulator and the task geometry.

Raibert, M., "A State Space Model for Sensorimotor Control and Learning," M.I.T. A.I. Memo No. 351, January 1976.

Raibert, M., "Control and Learning by the State Space Model: Experimental Findings," M.I.T. A.I. Memo No. 412, April 1977.

These papers present a model for manipulator motor control and learning, and the results of experiments which demonstrate the controller's ability to learn new movements, adapt to mechanical changes caused by inertial and elastic loading, and generalize its behavior among similar movements.

Waters, R.C., "A Mechanical Arm Control System," M.I.T. A.I. Memo No. 301, January 1974.

This paper proposes a system of arm control for a variety of simple actions. Also discussed are force feedback and the computed torque drive method.

Vukobratovic, M., Stokic, D., "Control of Manipulation Robots," Springer-Verlag, 1982.

This book extensively discusses control synthesis for robotic manipulators.

Man-Machine Interface

Black, J.H., "Factorial Study of Remote Manipulation with Transmission Time Delay," M.S. Thesis, M.I.T., December 1970.

Ferrell, W.R., "Remote Manipulation with Transmission Delay," IEEE Trans. Human Factors in Electronics, Vol. HFE-6, no. 1, September 1965.

Ferrell, W.R., "Delayed Force Feedback," Human Factors, vol. 8, pp. 449-455, October 1966.

Hill, J.W., "Comparison of Seven Performance Measures in a Time-Delayed Manipulation Task," IEEE Trans. on Systems, Man, and Cybernetics, vol. SMC-6, no. 4, April 1976.

Sheridan, T.B., "Remote Manipulative Control with Transmission Delay," IEEE Trans. Human Factors in Electronics, vol. HFE-4, no. 1, September 1963.

These papers discuss the problem of time delays in the control loop of a teleoperator. Experiments using manipulators with and without force feedback to perform a variety of simple tasks are described.

Brooks, T.L., Sheridan, T.B., "Superman: A System for Supervisory Manipulation and the Study of Human/Computer Interactions," M.I.T. Sea Grant Report No. MITSG 79-20, July 1979.

Ferrell, W.R., Sheridan, T.B., "Supervisory Control of Remote Manipulation," IEEE Spectrum, vol. 4, October 1967.

Hardin, P.A., Whitney, D.E., Sheridan, T.B., ""AND TREE" Computer Data Structures for Supervisory Control of Manipulation," Proc. 1972 Int. Conf. on Cybernetics and Society, IEEE, 1972.

Whitney, D., "State Space Models of Remote Manipulation Tasks," IEEE Trans. on Automatic Control, Vol. AC-14, no. 6, December 1969.

Winey, C.M., "Computer Simulated Visual and Tactile Feedback as an Aid to Manipulator and Vehicle Control," M.I.T. Man-Machine Systems Lab Report, 1981.

These papers describe various aspects of supervisory control: man-machine interfaces, data structures, planning algorithms, goal description, etc.

Flanagan, J.L., "Computers that Talk and Listen: Man-Machine Communication by Voice," Proc. IEEE, vol. 64, no. 4, April 1976.

Description of computer techniques emerging in the laboratory for voice communication between man and machine.

Freedy, A., Hull, F.C., Lucaccini, L.F., Lyman, J., "A Computer-Based Learning System for Remote Manipulator Control," IEEE Trans. on Systems, Man, and Cybernetics, vol. SMC-1, no. 4, October 1971.

Describes a concept for an adaptive teleoperator, capable of learning by observing the operator's control function. Initially, the output device is totally controlled by the operator, while the computer system acts as a passive observer. As the operation continues, the computer gradually assumes the role of active controller.

Hill, J.W., Salisbury, J.K., "Study to Design and Develop Remote Manipulator Systems," SRI International Project 4055 report. November 1977.

Describes apparatus and results of peg-in-hole manipulation experiments, to determine the advantages of force feedback used with a master-slave manipulator.

Olsen, R.E., "Dexterous Manipulator Laboratory Program," Proc. 18th Conf. on Manual Control, June 1982.

Describes Grumman Aerospace Corp. test apparatus and task boards for teleoperation development.

Shields, N., Piccione, F., Kirkpatrick, M., Malone, T.B., "Human Operator Performance of Remotely Controlled Tasks," Essex Corp. Report No. H-82-01.1, March 1982.

This report summarizes the teleoperation tests conducted at the Marshall Space Flight Center between 1971 and 1981 by the Essex Corporation.

Artificial Intelligence

Arbib, M., "The Metaphorical Brain," Wiley Science: New York, 1968.

Arbib, M., "Brains, Machines and Mathematics," McGraw Hill: New York, 1964.

Ashby, R., "Design for a Brain," Science Paperback, Chapman S. Hall, Ltd.: London, 1960.

Wiener, N., "Cybernetics," The Technology Press, M.I.T.: Cambridge, 1948.

These references give a background on A.I. from a historical perspective.

Bobrow, D.G., Collins, A., "Representation and Understanding," Academic Press: New York, 1975.

Nilsson, N., "Principles of A.I.," Tioga Press: Palo Alto, 1980.

Winston, P., "Artificial Intelligence," Addison Wesley Press, 1977.

These books set forth the basics of A.I., namely, the methods for representing problems, the techniques for solving problems, and the ways in which these methods might be organized into computational systems.

Machine Vision

Ballard, S., Brown, C., "Computer Vision," Prentice Hall, 1982.

This book is a general overview of A.I. techniques for modelling the visual world and reasoning about visual problems.

Gennery, D., Cunningham, R., Saund, E., High, J., Ruoff, C., "Computer Vision," JPL Publication 81-92, November 1981.

The results of a very thorough literature search, which gives technical descriptions of a wide variety of algorithms and concepts used in current research.

Pfister, G.F., "On Solving the Findspace Problem," M.I.T. A.I. Lab. Working Paper No. 113, March 1973.

Discusses a method of solving the problem of finding a volume, in a bounded space containing impenetrable objects, where another object of specified dimensions will fit. This type of algorithm is not specifically a "vision" algorithm, but will be required to make full use of a vision system.

Speckert, G., "Hand-Eye Coordination," M.I.T. A.I. Lab. Working Paper No. 127, July 1976.

Describes a simple method of converting visual coordinates to arm coordinates, which does not require knowledge of the camera position.

Robotics

Ambler, A.P., Barrow, H.G., Brown, C.M., Burstall, R.M., Popplestone, R.J., "A Versatile System for Computer-Controlled Assembly," Artificial Intelligence, vol. 6, pp. 129-156, 1975.

Describes a system which uses TV cameras, a moving table, and a manipulator to make simple assemblies such as a peg and rings and a toy car.

Bolles, R., Paul, R., "The Use of Sensory Feedback in a Programmable Assembly System," Stanford A.I. Lab Memo AIM-220, October 1973.

Describes a system which uses TV and a manipulator arm to assemble a water pump.

Ejiri, M., Uno, T., Yoda, H., Goto, T., Takeyasu, K., "A Prototype Intelligent Robot that Assembles Objects from Plan Drawings," IEEE Trans. on Computers, vol. C-21, no. 2, February 1972.

Describes a set of basic decision-making algorithms for assembly tasks.

Ruoff, C.F., "PACS-An Advanced Multitasking Robot System," The Industrial Robot, June 1980.

Describes a sensor-based parallel processing robot system which can be easily integrated into the manufacturing environment. A minicomputer is used to control two cooperative arms, using vision and force sensors.