


AD A116236



Man-Machine Systems Laboratory

Department of Mechanical Engineering
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
	AD-A116 236	
4. TITLE (and Subtitle) Supervisory Control: Problems, Theory and Experiment for Application to Human-Computer Interaction in Undersea Remote Systems		5. TYPE OF REPORT & PERIOD COVERED Technical Report Feb 1, 81-Jan 31, 82
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Thomas B. Sheridan		8. CONTRACT OR GRANT NUMBER(s) N00014-77-C-0256
9. PERFORMING ORGANIZATION NAME AND ADDRESS Massachusetts Institute of Technology Cambridge, MA 02139		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS NR 196-158
11. CONTROLLING OFFICE NAME AND ADDRESS Engineering Psychology Group Office of Naval Research (Code 442EP) Arlington, VA 22217		12. REPORT DATE March 1, 1982
		13. NUMBER OF PAGES 138
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		16. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) DTIC SELECTED JUN 29 1982 H		
18. SUPPLEMENTARY NOTES None		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) supervisory control teleoperators robotics man-machine systems remote control manual control human performance remote undersea vehicles human factors man-computer interaction unmanned vehicles computer control		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This technical report provides an overview of a four-year theoretical and experimental research project on human-computer cooperation in remote control of undersea manipulators and vehicles. The report puts forth a theoretical framework for viewing supervisory control in terms of a hierarchical interaction of three fundamentally (continued on next page)		

DD FORM 1473

EDITION OF 1 NOV 68 IS OBSOLETE
S/N 0102-014-6601

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Cont'd

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

different types of control loops: (1) interaction of a human supervisor with a (high level) computer; (2) interaction between the high level and one or more low level computers; and (3) interaction of the low level computers with the environment as mediated through artificial sensors and effectors. The perspective is one of coordinated allocation of human and artificial resources to various functions at various levels and times within given tasks. Numerous problems are discussed in conjunction with human and computer cognition, display command and control, matching capabilities to tasks and evaluation aspects.

Experiments performed in our laboratory on human supervisory control of remote undersea manipulators, vehicles and dynamic processes are reviewed briefly in relation to the theoretical constructs. These draw upon already published detailed reports of each experiment.



Accession For	
NTIS	✓
DTIC 1	
Unannounced	
Justification	
By	
Distribution	
Availability Codes	
Dist	and/or
A	Special

B

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

PREFACE

In July 1978 a first year's report on this contract was published, "Human and Computer Control of Undersea Teleoperators", by Sheridan and Verplank. That report reviewed such human performance theory as seemed relevant, and suggested various problems of applying supervisory control to undersea manipulators and vehicles. At that time the field had no overall theoretical framework for viewing supervising control, no satisfactory command languages, and little confidence that supervisory control was really good for much that was practical. Though no claim can be made that all problems of supervisory control are "solved", we now claim both clearer vision and greater confidence in these techniques than in 1978.

"Doing science" in supervisory control has not been easy. These man-machine systems are complex in structure and function. They are discontinuous in time and space, and non-linear in input-output functional relations. It is not convenient to have smooth variation in experimental variables. One must build hardware and software configurations which embody fixed levels of those man-machine variables. Recognizing that technology must sometimes lead science in order to provide the apparatus for the scientific questions to be asked, we have forced certain "working prototype" man-machine systems into existence. Our purpose in doing this was always to demonstrate and test principles rather than to develop and test hardware.

We have recognized the need to develop a unified and integrated theory of human supervisory control of computerized semi-automatic machines. In working toward this goal we seem to keep uncovering more and more different but seemingly relevant segments of theory from both engineering and psychology. Supervisory control does not fit neatly under any of the conventional categories. It includes optimal control in engineering but is not optimal control. It includes attention, decision making, and communication in psychology but is not any one of these. It is neither computer science nor cognitive science, but much of each is relevant. For any given large-scale supervisory control system, such as an air traffic control system or a large process plant or a military command/control system, descriptive models from many viewpoints would be appropriate. I doubt however that any single normative theory will ever unify supervisory control. In part this is, as explained in the text, because the criteria of good system performance themselves come from within the system - namely the supervisory operator who plays goal-setter.

I thought hard about how to organize this material. Since we are at an early stage of modeling supervisory control it was evident that even the organization of a report such as this forms a crude taxonomic model of the subject. Simultaneous with this report writing I am collecting material for a book on supervisory control. With the blessing and encouragement of Gerald Malecki and Martin Tolcott, the contract monitors of this research, I offer this report as a "trial balloon" for some of the ideas I plan to pursue in more detail in the book. Where the examples of supervisory control in this report concentrated on undersea manipulator and vehicle control applications, the planned book will also include much application material from the process control industry, primarily nuclear power, and the commercial aviation industry. There is also the intention to deal in greater depth about communication with computers regarding human goals and values in system operation, since that is the contribution of the human operator farthest from takeover by computers.

ACKNOWLEDGEMENT

I wish to thank the Office of Naval Research, particularly Gerald Malecki of Engineering Psychology, for consistent support, encouragement and faith that something would come out if they waited long enough.

I also wish to thank, of course, the students and colleagues in the MIT Man-Machine Systems Laboratory whose work this report reviews. They are listed on page viii. It is knowing and working with such wonderful people that makes it most worthwhile.

TABLE OF CONTENTS

v

SUPERVISORY CONTROL: PROBLEMS, THEORY AND EXPERIMENT FOR APPLICATION TO HUMAN-COMPUTER INTERACTION IN UNDERSEA REMOTE SYSTEMS

	Page
Abstract	i
Preface.	iii
1. <u>What is Supervisory Control</u>	1-1
1.1 Definitions.	1-1
1.2 Supervisory control in different applications.	1-2
a) process control	1-2
b) vehicle control	1-3
c) manipulators and discrete parts handling.	1-5
d) information processing.	1-8
2. <u>Modeling Supervisory Control</u>	2-1
2.1 General model of supervisory control system.	2-1
2.2 Multiplicity of supervisory functions.	2-5
2.3 Rasmussen's knowledge - rule - skill hierarchy	2-8
2.4 Allocation of resources: off-loading of rule and skill behavior on HIS and TIS computers	2-12
2.5 Normative dynamic attention allocation in supervisory control. .	2-17
*2.6 Experiments in doing it yourself vs assigning a computer to do it	2-20
3. <u>Considerations of Cognition and Internal Modeling</u>	3-1
3.1 The idea of an internal model, mental or computerized.	3-1
3.2 Loci of internal models in systems	3-3
3.3 Uses of the internal model in systems.	3-3
3.4 Identifying the operator's internal model.	3-7
3.5 Failure detection/isolation as disparity between reality and internal model	3-8
*3.6 Experiments with a particular model-based failure detection/ isolation technique.	3-10
3.7 Deciding when to stop learning and start acting.	3-13
3.8 The general planning problem as matrix inversion: use of internal models.	3-14

TABLE OF CONTENTS

Continued

	<u>Page</u>
3.9 Mutual understanding, by man and computer, of each other.	3-14
3.10 Group decision in supervisory control	3-15
4. <u>Considerations of Sensing and Display</u>	4-1
4.1 Application of conventional human factors criteria to displays for the supervisor.	4-1
4.2 New opportunities and problems with computer-generated CRT displays.	4-2
a) paging structures and access	4-2
b) integration and format	4-4
c) user-scaled trends and predictions	4-8
d) multi-variable data search	4-8
e) zoom control of degree-of-abstraction.	4-9
f) too many options: killing the operator with kindness	4-12
4.3 Special teleoperator display problems	4-12
a) bandwidth constraints.	4-12
b) teleproprioception	4-13
c) teletouch.	4-13
d) superposing video and graphics	4-13
*4.4 Experimental evaluation of frame-rate, resolution, grayscale constraints	4-14
*4.5 Operator-adjustable frame-rate, resolution and grayscale under fixed bandwidth constraint.	4-19
*4.6 Demonstration of computer-graphic arm-hand teleproprioception aids. .	4-20
*4.7 Demonstration of computer-graphic aids for tele-touch	4-27
*4.8 Experiments with computer-graphic predictor displays for remote vehicle control with transmission delay and slow frame rate	4-31
4.9 Emergency displays.	4-34
*4.10 Experiments in "smart-display" of failure detection/isolation	4-35
5. <u>Considerations of Responding and Control</u>	5-1
5.1 Application of conventional manual control to supervisory control . .	5-1

TABLE OF CONTENTS

Continued

	<u>Page</u>
5.2 New opportunities and problems with computer-based supervisory command and control	5-3
a) teaching, initializing and commanding.	5-3
b) interoceptive vs. exteroceptive control.	5-3
c) analogic vs. symbolic control devices and codes.	5-4
d) virtual control panels and devices	5-5
e) authority vs. deference, trading vs. sharing	5-6
f) zoom control of degree-of-abstraction.	5-7
g) prompts, edits, threats and defaults	5-7
h) too many response options.	5-8
5.3 Special teleoperator control problems	5-8
a) absolute and relative coordinate frames.	5-8
b) relative motion compensation	5-9
c) manipulation accommodation	5-9
d) bandwidth, time delay and visco-inertial lag	5-11
*5.4 Implementation of a supervisory controlled telemanipulation system. .	5-12
*5.5 Experiments to compare supervisory control of telemanipulation with direct control.	5-13
*5.6 Refinement of command language for supervisory telemanipulation . . .	5-18
*5.7 Experiments with relative motion compensation in telemanipulation . .	5-23
*5.8 Experiments with supervisory control of undersea vehicles	5-24
5.9 Emergency control and ultimate authority.	5-26
5.10 Telling a computer how to decide.	5-27
6. <u>Matching Capabilities to Tasks and Evaluating Supervisory Control Performance.</u>	6-1
6.1 Tasks, tools, and supervisory control	6-1
6.2 Criteria for evaluating supervisory control	6-3
6.3 Some methods for analyzing supervisory control.	6-6
6.4 Supervisory errors and system reliability	6-8
6.5 Mental workload in supervisory control.	6-10
6.6 Simulation, test-retest reliability and validity.	6-12
6.7 How far to go in automation, and when to stop	6-13

TABLE OF CONTENTS
Continued

	<u>Page</u>
7. <u>Conclusions and Recommendations</u>	7-1
Conclusions	7-1
References	7-3

*Asterisks represent experimental studies, reported in detail in reference documents describing research supported by the ONR contract:

T. Sheridan and W. Verplank (1978)

T. Brooks (1979)

V. Ranadive (1979)

M. Takahashi (1979)

B. Deghuee (1980)

N. Moray (1980)

K. Tani (1980)

D. Fyler (1981)

H. Hirabayashi (1981)

C. Winey (1981)

S. Messner (1982)

U. Tsach (1982)

J. Tzelgov (1982)

W. Wood (1982)

D. Yoerger (1982)

1. WHAT IS SUPERVISORY CONTROL

1.1 Definitions

Supervisory control of a process means a human operator communicates with a computer to gain information and issue commands, while the computer, through artificial sensors and actuators, implements these commands to control the process.

Any process can be brought under human supervisory control and thus be subsumed under this definition, including vehicles (aircraft, spacecraft, ships, submarines, ground vehicles of all kinds), continuous product processes (oil, chemical, fossil and nuclear power plants), discrete product processes (manufacturing, construction, farming), robotic/teleoperator devices where not included above, and information processing of all kinds (air traffic, military command and control, office automation, etc.).

A dictionary definition (Webster unabridged) may be appropriate to consider here. To supervise is to: "look over in order to read, peruse, coordinate, direct, inspect continuously and at first hand the accomplishment of, oversee with powers of direction and decision the implementation of one's own or another's intentions, superintend."

This definition, with words like "direct" and "oversee with powers of direction and decision"--and "superintend" imply goal setting, value-assigning and planning. Other words like "coordinate", imply guiding toward a given goal - in a computer systems context we might call this programming or parameter adjusting. The remaining words, "look over---," "peruse", "inspect---" have to do with continual observation or monitoring. These are clearly different functions, and this multiplicity of functions explicitly shows up in computer-aided supervisory control, as we shall see.

Intrinsic to supervisory control is the idea of teleoperation - man performs a sensing and/or manipulation task remotely by use of artificial sensors and actuators. This can be spatial remoteness, as with a remotely controlled vehicle or manipulator undersea or in space. It can be temporal remoteness, due to a time delay between when an operator issues commands and when he receives feedback. Or it can be functional remoteness, meaning that what the operator sees and does and what the system senses

and does bear little superficial resemblance. Teleoperation can be either the motivation for or the result of supervisory control, as will be made evident.

1.2 Supervisory Control in Different Applications

This section provides brief comparisons and contrasts among different applications of supervisory control: process control, vehicles, manipulators, information systems and human organizations.

(1) Process Control

The term "process" usually means a dynamic system such as a fossil or nuclear power generating plant or a chemical or oil production facility which is fixed in space and which operates more or less continuously in time. Further, there is an implication that the product is spatially continuous, and flows through pipes (e.g. chemicals) or along guides (wire extrusion, papermaking) or along wires (electricity). Some operations may be done on a "batch processing" basis. Typically time constants are slow - many minutes or hours after a control action is taken until most of the system response is complete.

Most such processes involve large structures with fluids flowing from one place to another and involve use of heat energy to affect the fluid or vice-versa. Typically they are multiple-man, multiple-machine systems, where at least some of the people are moving around from one location of the process to another. Usually there is a central control room where many measured signals are displayed and where valves, pumps and other devices may be commanded to operate.

"Supervisory control" has been emerging in process control for several decades. Starting with electro-mechanical "controllers" or control stations which could be adjusted by the operator to maintain certain variables to within limits (a home thermostat is a common example), special electronic circuits which replaced the electro-mechanical function gradually took over. At any time in such systems the operator can switch to manual control and put himself in the control loop. Usually each "control station", a three-by-six inch unit on the control panel, displays both the variable being controlled (e.g. room temperature for the thermostat) and the control signal (e.g. flow of heat from the furnace). Usually many such manual control devices are lined up in the control room, together with manual switches and

valves, many status lights, many dials and recording displays, and as many as 1500 alarms or annunciators - little square windows which light up and indicate what plant variable has just gone above or below limits. From a pattern of these (e.g. 500 the first minute of a loss-of-coolant accident, 800 the second minute, by recent count in a large new nuclear plant) the operator is supposed to divine what is happening.

Starting a few years ago the full computer has found its way into process control. Instead of multiple, independent, conventional PID controllers for each variable, the computer can treat the set of variables as a vector and compute that control trajectory which would be optimal (in the sense of quickest, or most efficient, or by whatever criterion is important). Because the interactions are many more than the number of variables, the variety of displayed signals and the number of possible adjustments or programs the human operator may input to the computer-controller are potentially much greater than before. This does not mean the operator necessarily sees greater complexity, however, since many of the computer's functions may be inaccessible or unknown to him since he is more or less confined to the control room (Typically there are other persons - technicians or "auxiliary operators" who work out in the plant while it is operating and who check, maintain or repair equipment or who perform control tasks on verbal orders from the control room operators. These are tasks which cannot be done from the control room because all control is not remote). Thus there is now, and this has accelerated since the events at Three-Mile Island, a great need to develop computer-interactive displays which integrate complex patterns of information for the operator and allow him to issue his commands in a natural, efficient and reliable manner. The term "system state vector" is fashionable, and the problem is how to display fewer "chunks" of information (using G. Miller's well-known terminology) to convey more meaning about the state vector of variables, where it has been in time and where it is likely to go in the near future.

(2) Vehicle Control

Unlike the processes described above, vehicles move through space, carrying their operators with them, or being controlled remotely. There are various types of vehicles which have come under a significant degree of supervisory control in the last 30 years.

We might start with spacecraft, not because spacecraft are the most costly and sophisticated (which they are not necessarily) but because in a sense their task is the simplest. They are launched to perform well-defined missions, and their interaction with their environment, other than gravity, is nil. In other words there are no obstacles and no unpredictable traffic to worry about. It was in the spacecraft, especially Apollo, where human operators (astronauts) who were highly skilled at continuous manual control (test pilots or "joystick jockeys") had to adapt to a completely new way of getting information from the vehicle and giving it commands - this new way was to program the computer. The astronauts had to learn to use a simple keyboard with programs (different functions appropriate to different mission phases), nouns (operands, or data to be addressed or processed) and verbs (operations or actions to be performed on the nouns). This was a rude shock to the astronauts and they complained, but they learned it.

There were, of course, a certain number of continuous control functions performed by the astronauts. They maneuvered the attitude of the vehicle (both command module and lunar lander) as well as its sextant to bring stars and landmarks into proper view for navigation fixes. They maneuvered the lunar lander's attitude and velocity for rendezvous with the command module in lunar orbit. But, as not generally appreciated by the public, control in each of these modes was heavily aided. Not only were the manual control loops themselves stabilized by electronics, but there were non-manual, automatic control functions being simultaneously executed and coordinated with what the astronauts did. This parallel or "shared" control mode is different from pure hierarchical control; this distinction will be clarified further in subsequent sections.

In both commercial and military aircraft the trend in the last decade or two has been to more and more supervisory control. Commercial pilots now are called "flight managers", indicative of the fact that they must allocate their attention among a large number of separate but complex computer-based systems. Military aircraft are called "flying computers", and indeed the cost of the electronics now far exceeds the basic airframe cost. Inertial measurement has become common in the new jumbojets as well as military aircraft, and this means the pilot can tell the computer to take him to any latitude, longitude, and altitude on the earth, and the vehicle will do so to within a fraction of a kilometer. But in addition there are many other supervisory command modes intermediate between such highest level commands and the lowest level of pure continuous control of ailerons, elevators and thrust. The pilot can set his autopilot to give him display of a smooth command course at fixed turn or climb rates to follow-manually, or he may

have the vehicle slaved to this course. He can set the autopilot to achieve a new altitude on a new heading. He can lock on to radio beams or radar signals for automatic landing. In the Lockheed L1011 there are at least ten identifiable separate levels of control. It is important that the pilot have reliable means to break out of these automatic control modes and revert to manual control or some intermediate mode. For example, when in automatic landing the pilot can either push a yellow button on the control yoke or he can jerk the yoke back and return the aircraft to his direct manual control.

Air traffic control poses interesting supervisory control problems, for the headways (spacing) between aircraft in commercial airspace are getting tighter and tighter, and efforts both to save fuel and to avoid noise over densely populated urban areas require more radical takeoff and landing trajectories. New computer-based communication aids will supplement purely verbal communication between pilots and ground controllers, and new display technology will help the already overloaded ground controllers monitor what is happening in three-dimensional space over larger areas, providing predictions of collision and related vital information. The CDTI (cockpit display of traffic information) is a new computer-based picture of weather, terrain hazards such as mountains and tall structures, course information such as way-points, radio beacons and markers, runways and command flight patterns, as well as the position, altitude, heading (and even predicted position) of other aircraft. It makes the pilot less dependent on ground control, especially when out-the-window visibility is poor.

Ships and submarines are more recently converting to supervisory control. Direct manual control by experienced helmsmen sufficed for many years, but that long tradition has been broken by the installation of inertial navigation, which calls for computer-control and provides capability never before available, and by the trends toward higher speed and long time-lags produced by larger size (e.g. the new supertankers). New autopilots and computer-based display aids, similar to those in aircraft, are now being used in ships.

(3) Manipulators and Discrete Parts Handling

In a sense, manipulators combine the functions of process control and vehicles. The manipulator base may be carried on a spacecraft, a ground vehicle or a submarine, or its base may be fixed. In either case the hand (gripper, end effector)

is moved relative to the base in up to three degrees of translation and three degrees of rotation. Typically it may have one additional degree of freedom for gripping, but some hands have differentially moveable fingers or otherwise have more degrees of freedom to perform special cutting, drilling, finishing, cleaning, welding, paint spraying, sensing or other functions.

Manipulators are coming to be used in many different applications, and the type of supervisory control and justification for it may differ depending upon application.

Historically the first purely manual remote manipulators were built for the nuclear "hot lab" in the late 1940's. The basic design configuration developed in that period is still popular. There followed development of simple manipulators for lunar roving vehicles, for simple undersea operations and for hazardous operations in industry, such as heat treatment of parts.

The serious difficulties of direct manual control from the earth of a manipulator on the moon (causing a minimum of a three-second radio transmission time delay in the control loop) were demonstrated by Ferrell and Sheridan (1967). To avoid instability one had to wait the full three seconds after each of a series of incremental movements, requiring a very long time to complete even the simplest manual task. The longer the time delay the larger the task completion time. In 1964 I had first proposed having a computer on the moon which could be programmed from earth to implement segments of the task locally in semi-automatic fashion, responding to its own sense of touch to branch into different program segments without the impediments of time delay in this local computer-control loop, and stopping and waiting for further instruction when it finished the programmed segment. We proposed calling this "supervisory control". Such local autonomy could also make the manipulator stop when colliding with an obstacle, or take other self protection actions with quick response. Normally the three second penalty in sending the program up to the computer (as compared to the human operator and time delay being inside the direct control loop) would pose no significant instability.

The time delay problem also shows up in a more recent application: remote manipulation undersea doing oil rig or pipeline inspection, repair, surveying the bottom, etc. If a power/signal cable connects the human operator in a surface

ship with the remote vehicle on the ocean bottom or at some considerable depth there may be several kilometers of heavy cable payed out. Even if this cable is made neutrally buoyant, the ocean currents cause considerable drag forces on the remote undersea vehicle/manipulator. Even worse, when working around undersea structures the cable may become entangled, often with loss of expensive equipment as well as time. One solution is to use acoustic transmission of signals for at least part of the distance, and let the vehicle/manipulator carry battery power. But with sound traveling at 5000 ft/sec, a 2500 ft. sound path poses a significant one-second time delay, so we are back again to a need for supervisory control.

There may be other reasons for supervisory control of the undersea vehicle/manipulator which are more compelling. There are things the operator cannot sense or can sense with only great difficulty and time delay (e.g. the mud may easily be stirred up, producing turbid opaque water which prevents the video camera from seeing), so that local sensing and computer response may be more reliable. For monotonous tasks (e.g. inspecting pipelines or structures or ship hulls, surveying the ocean bottom to find some object) the operator will not remain alert for long, and provided adequate artificial sensors can be supplied for the key variables, supervisory control should be much more reliable. Finally, the human operator may have other things to do, so that supervisory control would enable him to check periodically to update the computer program or help the remote device get out of trouble. A final reason for supervisory control, and often the most acceptable, is that if communications, power or other systems fail there can be fail-safe control modes that the remote system branches into to get it back to the surface or otherwise render it recoverable.

Many of these same reasons for supervisory control occur in other applications of manipulators. Probably the greatest current interest in manipulators is for manufacturing (so called industrial robots), including machining, welding, paint spraying, heat treatment, surface cleaning, bin picking, parts feeding for punch presses, handling between transfer lines, assembly, inspection, loading and unloading finished units, and warehousing. Today repetitive tasks such as welding and paint spraying can be taught (programmed) by the supervisor, then implemented with the control loops closed only in terms of arm positions or velocities. If the transfer line is sufficiently reliable, welding or painting non-existent objects

seldom occurs, so that touch or vision are usually not required. Manufacturing assembly, however has proven to be a far more difficult task.

The computer can perform in ways that are very difficult for human workers. For example, if a known mix of products is coming down the assembly line in a known order, the computer can then treat each product according to its appropriate (different) program without any forgetting or confusion. A human worker would become very confused.

In contrast to assembly-line tasks where, even if there is a mix of products, every task is preperformed, there are many emerging applications of manipulators with supervisory control in which each new task is unpredictable to considerable extent. Some examples are in mining, earth-moving, building construction, building and street cleaning and maintenance, trash collection, logging, and crop harvesting where large forces and power must be applied to external objects. The human operator is necessary to program or otherwise guide the end effector in some degrees of freedom (to accommodate each new situation) while in other respects certain characteristic motions are preprogrammed and only need to be initiated at the correct time. In some medical applications such as microsurgery the need is somewhat opposite from the above - to minify rather enlarge motions and forces, to extend the surgeon's hand tools through tiny body cavities to cut, to obtain tissue samples, to remove unhealthy tissue, or to stitch. Again, the surgeon controls some degrees of freedom (e.g. of an optical probe or a cauterizing snare) while automation controls other variables (e.g. air or water pressure).

(4) Information Processing

Supervisory control is so inherent in information processing that there has been hardly any comparison of it to other applications of control such as those described above. This is partly because the sensors and actuators are less explicit; the "task" already lies within the computer to some extent.

We can identify the trend toward supervisory control in data storage and searching, where manual storage and search through file cabinets and through books in libraries has given way to computerized storage and search. However even in computer aided search the operator may not know ahead of time exactly what he/she wants and will still be involved in continual observing and browsing.

Manual means of preparing or modifying data are being replaced by text editing, where programs are called upon in accordance with Figure 1 to modify the state of the task - in this case a document.

Electronic mail may also be viewed as supervisory control. This is a combination of text-editing and data searching, where the computer is called upon to deliver prepared messages to one or many individuals or to display messages which the operator may wish to read and is entitled to read.

Finally, new computer-aided video teleconferencing can also be considered to be supervisory control where the conference leader is the supervisory controller. The "task" in this case is to enhance natural communication between individual conferees (who are seated in front of video cameras and displays arbitrarily distant from one another) by allocation of audio and video and recording channels among individual voices, faces, documents and other shared materials. Eventually computer data-bases, models and other computational processes will surely be brought into such teleconferences to serve the same functions that guest experts, administrative assistants, vote-counters and others now serve in face-to-face meetings.

2. MODELING SUPERVISORY CONTROL

2.1 General Model of Supervisory Control System

Figure 1 characterizes supervisory control in relation to manual control and automatic control. Common to the five man-machine system diagrams are displays and controls interfaced with the human operator, and sensors and actuators interacting with a process or "task". The first two systems on the left represent manual control. (1) is without computer aiding while in (2) significant computer transforming or aiding is done in either or both sensing and acting (controlling) loops. Note that in both (1) and (2) all control decisions depend upon the human operator. When either the minor (3) or major (4) fraction of control is accomplished by control loops closed directly through the computer we call this supervisory control. If, once the control system is set up, essentially all the control is automatic (5), that is, if the human operator can observe but cannot influence the process (other than pulling the plug), it is no longer supervisory control.

The five diagrams are ordered with respect to degree of automation. The progression is not meant to imply either degree of sophistication or degree of desirability.

Figure 2 shows a more general model of a supervisory control system than Figure 1. The human component is still left as a single entity. There are two subsystems, the human-interactive subsystem (HIS) and the task-interactive subsystem (TIS). The HIS generates requests for information from the TIS and issues high level commands to the TIS (subgoal statements, instructions on how to reach each subgoal or what to do otherwise, and changes in parameters). The TIS, insofar as it has subgoals to reach, instructions on how to try or what to do if it is impeded, functions as an automaton. It uses its own artificial sensors and actuators to close the loop through the environment and do what is commanded.

Note that the HIS and TIS form mirror images of one another. In each case the computer closes a loop through mechanical displacement (hand control, actuator) and electro-optical or sonic (display, sensor) transducers to interact with an external dynamic process (human operator, task). The external process is quite variable in time and space and somewhat unpredictable.





Figure 1. Five man-machine systems ordered as to degree of automation



- 1) Task is observed directly by human operator's own senses.
- 2) Task is observed indirectly through artificial sensors, computers and displays. This TIS feedback interacts with that from within HIS and is filtered or modified.
- 3) Task is controlled within TIS automatic mode.
- 4) Task is affected by the process of being sensed.
- 5) Task affects actuators and in turn is affected.
- 6) Human operator directly affects task by manipulation.
- 7) Human operator affects task indirectly through a controls interface, HIS/TIS computers and actuators. This control interacts with that from within TIS and is filtered or modified.
- 8) Human operator gets feedback from within HIS, in editing a program, running a planning model, or etc.
- 9) Human operator orients himself relative to control or adjusts control parameters.
- 10) Human operator orients himself relative to display or adjusts display parameters.

Figure 2. General model of supervisory control system

The numbered arrows identify individual cause-effect functions, with explanations of the loops at the right. It is seen that there are three types of inputs into the human operator: (1) those which come by loop 1 directly from the task (direct seeing, hearing or touching), (2) those which come by loops 2 and 8 through the artificial display and are generated by the computer and (3) those which come by loops 10 and 9 from the display or manual controls without going through the computer (i.e. information about the display itself such as brightness or format, present position of manual controls, which is not information which the computer has to tell). Similarly, there are three types of human outputs: (1) those which go by loop 6 directly to the task (the human operator by-passes the manual controls and computer and directly manipulates the task, makes repairs etc; (2) those which communicate instructions via loops 7 and 8 to the computer, and (3) those which modify the display or manual control parameters via loops 10 and 9 without affecting the computer (i.e. change the location, forces, labels or other properties of the display or manual control devices).

Correspondingly there are three types of force and displacement inputs into the task: (1) direct manipulations by the operator via loop 6; (2) manipulations controlled by the computer via loops 3 and 7; and (3) those forces which occur by interaction, over loops 4 and 5, with the sensors and actuators and are not mediated or usually intended by the computer or operator. Finally there are three types of outputs from the task: (1) information fed back directly to the operator over loop 1; (2) information fed to the TIS computer via loops 2 and 3; and (3) information (in the form of forces and displacements) which modifies the sensors or actuators via loops 4 and 5 without being explicitly sensed by the computer.

When the task is near to the operator, the HIS and TIS computers can be one and the same. When the TIS is remote usually HIS and TIS computers are separated to avoid problems caused by bandwidth or reliability constraints in telecommunication, loops 2 and 7. This problem will be discussed in detail in section 4.3.

Multiplexing switches are shown in loops 1, 2, 7, and 6 to suggest that one HIS may be time-shared among many TIS, i.e. many tasks, each with its own local automatic control or robotic implementer. In fact, more and more this is coming

to be the case in supervisory control. In some process plants there are over 1000 TIS, some being primitive feedback controllers, some being simply programmed but highly reliable microcomputers. The sheer number of TIS causes a multiplexing or switching overhead cost.

This is a descriptive model of supervisory control; that is, it is intended to fit what is observed to be the structural and functional nature of a wide variety of situations we discussed earlier. The variables on the lines of Figure 1 are all measurable; there are no intervening variables, no suppositions about what is going on that we cannot observe readily. This is why we have not (yet) elaborated the human operator beyond a single entity. We will do such elaboration on the model and discuss its implication next.

Being a descriptive model it is by definition not a normative model. We have not (yet) imposed any notions of how the system should work, or of what optimal behavior consists, or how close actual behavior compares to optimal.

It is important to note, also, that we do not intend to develop a model of the human operator independent of the rest of the system. McRuer and Krendall in 1965 abandoned trying to model the human operator in a simple control loop as an invariant entity per se and turned instead to finding invariance in the series combination of human controller plus controlled process. Our approach, similarly, will be to find invariance in supervisory control plus tool box of computers, sensors and effectors plus task. We will note various functions that either human or computer can do, but some are best done by human and some are best done by computer. Which does what function will evolve over many years in the future and will always depend on circumstance. For now the intent is to provide a qualitative description of the combination.

2.2 The Multiplicity of Supervisory Functions

The essence of supervision, as noted in conjunction with the dictionary definition earlier, is that it is not a single activity, as we are accustomed to characterize various sensory-motor or cognitive or decision-making skills, or communication or controlling behavior. Supervision implies that the primary or direct activity, whatever it is, is normally being done by some entity (man or machine) other than the

supervisor. There may be a single primary task, or many such tasks. The supervisor, from outside, performs those many functions necessary to insure that the single entity does (or multiple entities do) what he, the supervisor, intends. Thus there may be multiplicity of function for two reasons:

1. For each primary task there are many different things to do to ensure that the primary entity (what was called the TIS in the model description) does what the supervisor intends that it do.
2. When there are multiple primary tasks, while the basic functions may be similar from one primary task to another, the data are different, and the initial conditions are different in performing the same function on each.

Our supervisory control model shows the supervisory computer to multiplex among, or alternately connect to, different TIS's or primary tasks. It also shows multiple connections to and from the human operator. It does not make clear that in switching from one TIS to another the initial conditions ("getting one's bearings") are different with each switch. Nor does it make clear that the human supervisor is continually switching functions even while dealing with a single TIS.

But this seems to be the essence of the supervisor: switching functions within one task and switching tasks. The remainder of this section elaborates this point.

Earlier, in conjunction with the dictionary definition, the ideas of "planning", "programming" and "observing" emerged as different components of "supervising".

Missing explicitly from the earlier dictionary definition but implied nevertheless are two additional functions. The first is taking over from the "other" entity, the TIS in our case, seizing direct control when indirect control by supervision fails. The second is to learn from experience.

Summarizing and elaborating to suit our present context, the supervisor, with respect to each task (and each TIS), must perform five distinct functions listed in Table 1.

Table 1. Functions of the supervisor

1. <u>Plan</u> a) be aware of what tasks are to be done, what resources are available, what resources (TIS) are committed to what tasks, and what resources are uncommitted b) decide on overall goal or goals, including objective function or tradeoffs among goals, and including criteria for handling uncertainties c) decide on strategy or general procedure, including logic of authority (human, HIS computer, TIS computer) in various situations d) consider known initial conditions and various combinations of probable inputs and possible actions and their consequences in view of system constraints and capabilities e) determine best action sequence to do what is intended under various situations f) decide what is to be considered abnormal behavior including automatic recovery from trouble, and what defaults or contingency actions are appropriate.
2. <u>Teach</u> (a, b, c and d could also be considered part of planning) a) estimate what the computers (HIS and TIS) know of the situation b) decide how to instruct the HIS to instruct the TIS to execute intended and abnormal actions c) decide how many of intended and abnormal actions TIS should undertake in one frame, i.e. before further instruction d) try out part or all of that instruction in (operator's) own mental and/or HIS computer model without commitment to transmit to TIS e) impart instruction (program) to HIS computer with commitment to transmit to TIS f) give command to HIS to start action
3. <u>Monitor</u> a) decide on what TIS behavior to observe b) specify to HIS computer the desired display format c) observe display, looking for signals of abnormal behavior and performing on-line computer-aided analysis of trends or prediction or cross-correlation as required d) observe task directly when and if necessary e) make minor adjustments of system parameters when necessary, as the automatic control continues. f) diagnose apparant abnormalities or failure, if they occur, using computer aids
4. <u>Intervene</u> a) decide when continuation of automatic control would cease to be satisfactory and minor parameter adjustments would not suffice either b) go physically to TIS or bypass all or portions of HIS and TIS computers to effect alternative control actions or stoppage or recovery c) implement maintenance or repair or modifications of TIS or task d) recycle to (1), (2) or (3) as appropriate
5. <u>Learn</u> a) decide means for collecting salient data and drawing inferences from it over repeated system runs b) implement these means c) allow for serendipitous learning d) periodically take stock of learning, modify system hardware and software, and anticipate future planning of operations e) develop understanding of and trust in the system

While the explanation of these five functions in Table 1 is not a consensus and some of these steps are manifest to a greater or lesser degree in any actual supervisory control situation, the point is that the necessary sequencing through these differing functions makes the human supervisory controller essentially different from the human in-the-loop, one-continuous function controller and decision-maker. These essential differences are summarized in Table 2 for eleven categories.

As implied above the allocation of attention by the human supervisor is both between functions for a given task and between tasks. In skilled or overlearned activities a person can engage in many at once (provided the required sensors and effectors are not overtaxed with respect to simple mechanical or signal processing considerations). Thus one can drive a car, talk, scratch his nose and look for a landmark at the same time. But one cannot do multiple simultaneous tasks each of which requires "new thinking" unless the time requirements are such that one can shift attention back and forth. In view of these facts we initially may characterize the attention allocation of the human supervisor as, first, selecting among alternative tasks to be done, and second, selecting his proper function with respect to that task.

2.3 Rasmussen's Knowledge-Rule-Skill Hierarchy (Rasmussen, 1976)

Before proceeding further with elaboration of the supervisory control model it is appropriate to discuss and interpret some ideas of Rasmussen, to whom the author is indebted for several key ideas.

Rasmussen has proposed a hierarchical model of human behavior in the operation of systems, Figure 3, a simple version of which is shown in Figure 4. The principal feature is that there are three levels of behavior which appear to varying degrees in different tasks:

1. Knowledge-based behavior. This the most abstract level. Features of the environment are used for identification of what the problem is. Decision is made, in consideration of goals, concerning what task to undertake next. Planning is carried out regarding the procedure for doing that task.

Table 2. Differences between conventional and supervisory controller

Conventional in-the-loop human controller, decision-maker	Supervisory Controller
1. operates with given goals or objective function, and given problem(s)	to a large extent determines goals and objective function, and finds the problem(s)
2. performs most duties inside the control loop	performs most duties outside the control loop
3. mode of interacting with the system sensors, actuators and computers is fixed	mode of interacting with the system sensors, actuators and computers is flexible
4. often, though not always, is directly in touch with task through own bodily sensors and effectors	mostly is remote from task, must use artificial sensors and effectors to control task
5. performs the same function continually	performs very different functions at different times
6. stays at same "hierarchical level of operation" with respect display and control information	can operate at different hierarchical levels at different times as required
7. activity is usually paced by the dynamics of process being controlled and control mechanism	activity is usually self-paced
8. cannot arbitrarily stop activity without system becoming unstable or control otherwise degenerating	can arbitrarily stop activity and system normally remains stable
9. gives continuous or continual commands spaced near to one another in time, each of which has relatively little information	gives intermittent commands, possibly spaced a long time apart, each of which can carry much information
10. display and control coding tends toward analogic	display and control coding tends toward symbolic
11. has little cognitive overhead devoted to managing his own time and attention	much cognitive overhead devoted to managing his own time and resources




Figure 3. Rasmussen's general decision model

From: R. Pew, D. Miller, & C. Feeher, Evaluation of Proposed Control Room Improvements Through Analysis of Critical Operator Decisions. Bolt, Beranek, & Newman; Cambridge, Mass. EPRI NP-1982, Aug. 1981.




Figure 4. Simplified version of Rasmussen's decision model

2. Rule-based behavior. Here a specific pattern is recognized. Associations are made with respect to that pattern (or state of the system or environment). A sequence of already stored discrete rules or steps is brought to bear for implementation.
3. Skill-based behavior. This is the least abstract level. Observed features of the environment are used continuously to drive the automatic sensory-motor behavior

Note in Figure 4 that the input at each level derives from the first step of the next lower level, and the output of each level depends on the last step in the next higher level in order to achieve implementation. Another way of stating the same idea is that knowledge-based behavior can be viewed as the outer control loop, closed through the environment, driven by higher goals (1) as independent reference input. Similarly rule-based behavior can be viewed as the middle control loop, closed through the environment, driven by procedure or subgoals (2) as independent reference input. Skill-based behavior can be viewed as the inner (lower) control loop, closed through the environment, responding to specific task rules as independent reference input.

Rasmussen points out that the wishes of the person drive the system from the top toward disaggregation of behavior, i.e. more and smaller responses, while physical constraints drive it from the bottom toward aggregation. Events at the top are symbols (theoretical, teleological). Events at the mid-level are signs (semantic)

2.4 Allocation of Resources: Off-Loading of Rule and Skill Behavior on HIS and TIS Computers, Respectively

We may now make use of Rasmussen's hierarchical model to add a third dimension to supervisory attention or effort allocation. The first two were allocation among tasks and among supervisory functions within tasks. The new allocation is among knowledge, rule or skill-based behaviors.

Figure 1, remember, treated the human operator as a single component in the system. However our purpose in this report is to understand and predict human behavior, how operators behave in present supervisory control systems, how they might behave if we made improvements, and consequently how to make such improvements. Thus we need to consider what might be happening (functionally, not physically) inside operators' heads.

We may view this allocation as illustrated by Figure 5. This shows the human operator's "problem space" as being subdivided four-dimensionally into:

1. a discrete set of tasks which, presumably, he decides to take on and about which he knows at least some of the requisite information about initial state, desired state (goals) and physical constraints;
- 2) for each task, the five functions of plan, teach, monitor, intervene, and learn which the operator must perform, as described in Table 1;
- 3) whether the required level of behavior is, according to the scheme proposed by Rasmussen, knowledge-based, rule-based or skill-based;
- 4) for each task and function, and level what requirements there are for sensing, cognition and responding (abbreviated S-C-R).

Examples of the S-C-R behaviors of the operator include:

sensing

- a) finding and orienting with respect to a display (loop 10 of Figure 2)
- b) reading and interpreting display (loops 2 and 8)
- c) observing the task (environment) directly (loop 1)
- d) reading printed documents
- e) receiving information from other persons

cognition

- a) recalling information
- b) running thought experiment in (hypothetical) internal model
- c) making a decision about where to seek information or what action to take

responding

- a) finding and orienting with respect to a control (loop 9)
- b) teaching or programming the HIS (loops 7 and 8)
- c) intervening to do direct control (loop 6)
- d) generating paper documentation
- e) giving information to other persons

It is tempting to ask the question "How do the different activities within this four-dimensional problem space get coordinated?" and for an answer appeal to the




Figure 5. Four-dimensional attention or effort allocation of the human supervisor

notion of a higher level box in the diagram which is a "coordinator". This would not be fair play. The buck must stop somewhere. Without trying to specify precisely what the coordination is or how it happens, its highest level must occur within the knowledge-based, cognitive, planning activities (cells) associated with various tasks.

I imagine the various cells to be a network of interconnected computational elements each having multiple inputs and outputs. Some, the S elements, control sensing transducers. Some, the R elements, control transducers capable of mechanical work. Some, the C elements, are only computers. Their interaction is somehow determined by the following properties which characterize each cell or element:

1. a set of conditions on the inputs which determine which of several algorithms will be run, if they are allowed to run at all;
2. the current or anticipated state of busyness which may inhibit responding to some or all current inputs;
3. outputs from the algorithms which represent either completed work or "default" signals that result from incomplete data or execution.

Minsky's "frame theory" is appealing here (Minsky, 1975). However I don't intend a mechanistic or explanatory model of the human supervisor - only to suggest the attributes of the "problem space" in a qualitative descriptive model.

While all tasks and all supervisory functions (i.e. all combinations of these) must receive some attention, not all behavior levels and not all S-C-R categories need exist in combination with the former. There will be empty cells in the array.

Subordinate to the operator, the HIS computer must similarly be allocated. In this case its behavior is not differentiated with respect to the five supervisory functions. However it must keep tasks and S-C-R subtasks separated.

Examples of S-C-R behaviors of the HIS computer are:

sensing

- a) receiving information from the TIS (loop 2)
- b) receiving program changes from the operator (loop 8)

cognition (all loop 8)

- a) fetching information from its own memory
- b) running a fast-time simulation on itself
- c) making a decision

responding

- a) generating a display for operator (loop 8)
- b) commanding the TIS computer (loop 7)

Note that only the "rule-based" and "skill-based" levels are filled in for the HIS computer. This is because, at the present time, what we are calling knowledge-based behavior in computers is only experimental and not ready for systems application.

Examples of S-C-R behaviors of the TIS computer are

sensing

- a) measuring information from the environment (loop 3)
- b) receiving program changes from the HIS computer (loop 7)

cognition (all loop 3)

- a) fetching information from its own memory
- b) making a decision (including conventional signal processing, optimal control calculations, etc.)

responding

- a) sending information to the HIS computer (loop 2)
- b) commanding the actuators (loop 3)

Note that only the "skill-based" level is filled in for the TIS computer, though one might argue that local small computers do upon occasion perform in what we defined as "rule-based" behavior.

Thus we observe that the human supervisor, capable of all three levels of behavior, can "off-load" rule and skill-based activities onto the HIS computer, and the HIS computer can in turn off-load skill activities onto the TIS computer.

This implies that at any one time the human, HIS computer, and TIS computer allocations probably will not be the same. Once any TIS is programmed it can run with neither the HIS computer nor the human supervisor attending to it. In fact, as multiple TISs come to be used the likelihood will be greater that both HIS computer and human supervisor are occupied elsewhere. (Once the pilot has the autopilot set and he observes that it is working properly he turns his attention to other "flight management" tasks). The HIS computer can be processing some data or even generating a display while the human operator's attention is diverted to something else. Clearly, some of the time it is expedient that human, HIS and TIS all be doing their own separate activity. At other times they might be waiting for each other, and this clearly is not efficient (e.g. when the robot is waiting for further instruction before going on with the assembly task). Thus, one problem is to decide when they should anticipate each other, when they should flag each other away from other pursuits, and when they simply should go their own separate ways.

As Rasmussen points out, there is greater breadth of focus of the problem at the higher behavior level, while focus becomes narrower and more specific at lower levels. A person simply cannot keep track of much detail and view all the interrelationships; to think more broadly one must code or "chunk" information in more compact ways.

The design principles which emerge immediately from such consideration are:

1. when rule-based and skill-based behaviors are appropriate and the HIS and TIS computers can perform these as well as the operator can, especially if he is busy with other things, the operator should "off load";
2. since he then necessarily loses contact with the detail he should be provided the capability to "zoom down" to the detail or "zoom up" to a higher, broader more general level at will.

2.5 Normative Dynamic Attention Allocation in Supervisory Control

The supervisory operator, communicating through a flexible interactive HIS computer, which in turn communicates with a number of specialized TIS computers, sensors and effectors, is provided a marvelous box of tools. Decisions of what tools to select

and when to use them are only partly at the discretion of the operator. To some extent both what to use and when to use them are forced by dynamic events. More than just a tool assigner, the supervisory operator is also a "problem finder" - he is actively looking for tasks to which to apply his tools.

In the present context it is relevant to raise again the question of whether people can do more than one activity at a time. As noted earlier, if those activities are automatic functions or well learned behaviors such as pedaling a bicycle, singing etc. there is hardly any interference with conscious thought. But at any one instant a person can only think consciously about one activity. In a paradigm described below we use the latter assumption - that the operator's cognitive attention can be directed to but one place at a time.

The dynamic nature of the supervisor's attention-allocation and decision-making may be characterized as shown in Figure 6. In an experiment we did this is what our subject saw on his display. Each block represents a task, the solid blocks being known tasks, the dotted ones expected tasks. Blocks appear at random times and distances from the vertical deadline at right, move at uniform velocities to the right and disappear when they hit the deadline. The distance of each block from the deadline therefore is the available time for doing that task. "Doing the task" requires the operator to hold his cursor in a corresponding space at right, and as it is "done" that block diminishes in width. Thus the initial width of the block is the time required for the operator to do the task. The height of each block is the earning per unit time of work. Thus the area of all blocks (tasks) "done" is the total reward. The supervisor's job is to allocate his cursor among the various task demands (blocks) some of which may overlap so that he cannot "do" them all. His objective is to "do" as much net area as possible and thus maximize his return.

Parameters of the paradigm are the statistics of block (task) arrivals in time and location relative to the deadline, heights and widths, rate of width reduction when "doing", and whether reward is only for completion or whether there can be partial rewards. There can also be costs or time delays imposed for moving the cursor from one block (task) to another - as there is in the real world.




Figure 6. Computer display in experiment on dynamic attention allocation




Figure 7. Wood's display for experiment on doing task yourself vs. assigning a computer to do it

The question arises, in conjunction with this paradigm or similar ones which are models of our supervisory attention allocation problem, is there a normative procedure, is there a best way to perform the allocation against which to calibrate human behavior or some mix of human and computer interaction?

Under simple and well defined circumstances the answer is obviously yes. For the same discrete blocks paradigm as described above the optimization algorithm can be derived by dynamic programming on a moment-by-moment basis (Tulga and Sheridan, 1980). Then one can get experimental answers from a simulator regarding better and worse strategies, and make a comparison to human subjects operating the simulator. Experimental subjects, it turned out in our experiments, did not differ so much from the optimal in total score, though their strategies were often quite different.

But even the above paradigm poses the situation in completely determined form. The times, penalties, capabilities and costs are given. In the real supervisory control problem, some or all of these factors may be obscure. Indeed, to the degree to which the supervisor sets his own objective function, i.e. he decides upon relative importances, no optimal can be determined. Normative models are successful in human performance modelling for simple tasks where goals and costs or rewards are clear to all concerned, including the experimental subject.

We are left with the problem of analyzing what the operator might select to do and use if he decides on particular goals or objective functions, and how we might improve his computational aids, displays and controls to help him do it. This is not unlike the improvement of any tool. We want to improve the tools in the tool box, having only a rough idea what the operator may wish to build and what are his criteria for successful construction. So we consider the range of constructions we think might he like to make, we consider his own constraints in interaction with various tools we might provide him, and we build him some prototype tools and see what he can do and how he likes them.

2.6 Experiments in Doing It Yourself vs Assigning a Computer to Do It

The paradigm described above, though it characterizes the multi-task attention allocation feature of supervisory control, does not include one important aspect of supervised automation. That is, it does not allow the operator to initiate, set-up, program or teach an automatic device to go to work on a task component such that, as he then directs his attention elsewhere, that task is done automatically.

A very rudimentary form of this latter problem was studied by the writer. (Sheridan, 1970). This is the case where the operator intermittently observes (samples) a single continuous variable and then tries to set a control to a "best" value until he samples again. Each sample incurs a penalty, so that he must trade-off sampling costs against the costs of control error. Dependent upon the signal statistics and the costs of error and of sampling, an optimum (expected value) supervisory sampling rate can be obtained. Experiments showed that experimental subjects sampled slightly more often than the optimum.

As part of the current ONR research effort Wood (1982, Figure 7) is currently experimenting with a dynamic blocks-diminishing multi-task similar to that described in 2.5. A key difference, however, is that the operator can either do the task himself or can assign a machine (shown on the subject's block display as a little bulldozer pushing the block) to do it. A second difference is that if the block (task) is not serviced by either human operator or machine it just sits there, and poses a cost per unit time if it is not completed. Finally, in Wood's experiment the operator may search freely among four "work areas" for tasks to do; he can see on his graphic display into but one of these at a time. In other respects the set-up is similar to that of Tulga and Sheridan, i.e. with key parameters for task arrivals in each work area, for service time and for lump reward upon task completion.

Wood has two versions of his experiment. In a first version the operator must both decide which work area to look into and decide whether to do the task himself or to assign a single machine (which is then tied up until completion of that task). In a second version of the experiment the strategy for search through work areas is fixed and there are multiple machines to assign.

Several interesting findings have emerged from experimental results thus far. When task arrivals were slowed down by a factor of four from a rather brisk pace to give the subject some time to consider his decisions, performance did not improve. It was found that even when the machine was set to be much less efficient than doing it himself (its "doing" rate was as slow as 0.1 of the manual rate) operators still tended to assign the machine. But when a significant cost was attached to machine use they gave it up. Apparently increasing the machine's wage had a larger effect on operator decision to use it than correspondingly decreasing its productivity.

A normative model was developed, including expected value comparisons of moving to other work areas or staying in the present work area and assigning a machine or doing the task one's self. For N work areas, looking ahead one step results in $2N + 2$ probabilistic outcomes to be evaluated before deciding what to do next. The model can look ahead M steps on this basis, resulting in a $(2N + 2)^M$ outcomes to be evaluated at each step. Clearly look-ahead more than several steps is impractical.

It is essential for the operator of Wood's simulation to have a "mental picture" of: (1) the meaning of parameters for arrival, service delay, rewards and costs; (2) which work areas have tasks; (3) where machines are in use and how long they will be busy; and (4) how long since he last examined each work area. This "mental picture" aspect of supervisory control is discussed more fully in the next section.

3. CONSIDERATIONS OF COGNITION AND INTERNAL MODELING

3.1 The Idea of an Internal Model, Mental or Computerized

Cognition is the act or process of knowing. An active process is implied, not simply passive memory and recall, though memory and recall are involved in cognition as well as in sensing and responding. The principal cognitive feature of human supervisory control, I assert, is the "internal model".

The idea is not new; it has its origins in antiquity. But probably in the 1950's the development of "model reference adaptive control" and of the "observer" in control theory first formalized the idea and simultaneously stimulated both control engineers and cognitive and man-machine scientists to make use of the idea. Figure 8 illustrates a model reference control system. The general idea is that the reference input to the real control system and the disturbances, insofar as they are measurable, are fed into a computer model whose transfer function is a norm, i.e. what behavior is desired. Any discrepancy between outputs of reference model and actual process becomes a basis for additional negative feedback control. As indicated, this tries to make the actual control system conform to the model.

Figure 9 illustrates a formal observer, as currently used by control theorists. The dotted outline at the right represents the standard-form system equations, $\dot{\underline{X}} = \underline{A} \underline{X} + \underline{B} \underline{u}$ and $\underline{y} = \underline{C} \underline{x}$, where $\underline{x}$ is the state vector, $\underline{u}$ the control signal, $\underline{y}$ the measured system output, and $\underline{A}$, $\underline{B}$ and $\underline{C}$ are linear matrices. The actual $\underline{x}$ cannot be measured. The control logic contains an exact-as-possible model, plus a matrix $\underline{G}$ which operates on any discrepancy of the model output $\underline{y}'$ from the actual process output $\underline{y}$ plus a linear gain $\underline{L}$ which can be optimized when $\underline{x}$ is accurately estimated (i.e. "observed").

The basic idea is that the control signal, $\underline{u}$, is fed both to the real-world process and to the model. Any discrepancy between outputs of the two provides a basis for improving the model. The model then allows "observation" of intervening variables, a form of measurement based on a model instead of the actual thing to be measured.

A person's internal or mental model, sometimes also called his "problem space" or "task image", is a hypothetical representation in the brain of real events and their relationships, such that he can predict what events will cause other events.




Figure 8. Model reference control system




Figure 9. Control system containing a formal observer

It therefore is more than just memory, though it uses memory. It is the cognitive analog of a dynamic computer simulation: given certain equations and certain inputs it shows what output will result. In its simplest form it is a static input-output function or a table: "if this input, then that output". It could be a more complex arithmetic computation. Usually, however, static models are insufficient, and some differential or difference equations are necessary, having variables which are functions of time.

3.2 Loci of Internal Models in Systems

The on-line model is internalized in the system in various forms, sometimes resident in a computer, sometimes resident in the head of an operator, and sometimes resident in the configuration of a display or a control which is geometrically isomorphic with the task or controlled process. Figure 10 illustrates this in the case of controlling a teleoperator. In this case the multi-axis control handle, which looks like a spaceship, allows the operator to move the control (this is sometimes called a "local model") in the way he wants the real spaceship to move. That it is desirable to have displays be geometrically isomorphic with the process may be obvious. Surely, then, as an operator uses a local model control and/or display it would be important that his model-in-the-head correspond.

There can be multiple Loci of internal models in large systems, including both on-line computers, which get reinitialized and updated by discrepancy feedback or other means, and off-line computers, which don't. Certainly every operator involved in a task has some internal model of it. Thus, ideally, there would be some continuous comparison made between the outputs of both types of computer models plus the results of the operator's thought experiments, all run with the same input stimulus as is driving the process itself. All of these would be compared to the actual output. Figure 11 illustrates the idea. Discrepancies would be the basis for investigation: either some model is incorrect or something has changed in the structure of the process. Perhaps a failure has occurred.

3.3 Other Uses of the Internal Model in Systems

Figure 12 shows a variety of uses of the internal model concept. Technique number 1 uses the model as a norm against which to assess process performance. This is an important means for failure detection (discussed further in 3.5 and 3.6). Number 2 is the "observer" as explained earlier. In Number 3 the parameters K_i of the

"INTERNAL MODELS"




Figure 10. Loci of internal models in a teleoperator system. Each numbered location in one way or another represents the state of the system, is updated either by the operator or by the hardware/software, and is referenced by the operator in planning and controlling.






Figure 11. Continual comparison of computer and human internal model outputs to actual process output




Measurement of process performance or failure detection

Use of model as norm or reference to assess process performance. Discrepancy e can also be taken to mean "failure" somewhere in the process.




Observer

Use of model to estimate state variables not directly measurable. g is "innovations signal" which helps to drive model to conform to process.




Parameter tracking

Use of model for identification of parameters. After sufficient adjustment of model parameters (to drive g to zero) those are asserted to be best estimates of process parameters.




Model-reference controller

Use of model for making noisy or nonstationary y/u' transfer function appear constant.



Predictor

Use of model for predicting y at future time. Fast model is repetitively reset with initial conditions from process (dotted lines), then run "into the future" on fast time-scale.



Training aid

Verbal protocols and special tests to determine accuracy of operator's mental model. Note that $+$ and $-$ are reversed from technique number (1) above to indicate that the process is now the reference.

Figure 12. Six techniques for using the internal model

model are adjusted in correspondence to $\partial K_i / \partial t : - dE/dt$ to force the model to conform to the process. This is called "parameter tracking" (see Sheridan and Ferrell, 1974). In number 4, assuming the process is noisy or slightly time-varying, the discrepancy signal can be directed to the actual process instead of the model to force the former to conform to the latter (opposite of 2). This was described earlier as "model reference control." Number 5 involves running the model on a time-scale M times faster than real-time for some fixed interval τ , resetting the model's state variables to correspond to those of the process, then iterating the fast-time run, doing so repetitively. Each new output trace then becomes a prediction of up to $M\tau$ into the future of what the process will do if the present input is sustained. Number 6 suggests the idea of using various inputs to the process (this could be a simulation) to see if the operator can predict what will happen. An experimenter would compare the actual process output (which is hidden from the operator) to what the operator says it should be, and would use various verbal protocols to find out why the operator believes the system will behave as he says it will.

3.4 Identifying the Operator's Internal Model

Discovering just what is the operators internal mental model, problem space or task image in any given situation is a challenge that cognitive scientists are keenly pursuing. "Knowledge-based" computer systems are being applied to this problem, where the interaction of the operator with the computer, the kind of information he seeks and the path by which he searches are revealing of his mental model. Verbal protocol techniques (Bainbridge, 1974) make use of key words and drawings. More formal psychometric techniques offer promise. No technique is satisfactory as yet.

It has been claimed that some nuclear power and other plant operators see their tasks in terms of the console itself, i.e. given certain signals on the displays they operate certain controls and follow certain procedures - the mental model is in terms of displays and controls. Other operators presumably "see through" the console, and as they look at displays and operate controls they envision tanks, fluid, pumps etc. Most would claim the latter "transparency" is preferable, though too "scientific" a mental model may distract the operator into contemplations which are not appropriate, especially in times of crisis.

3.5 Failure Detection/Isolation as Disparity Between Reality and Internal Model

Failure detection is perhaps the most important function of the human supervisor. The problem of inferring that a failure has occurred and isolating the locus of failure is a complex one, and there have been many behavioral insights derived from experiment. For example, when asked to explore cause-effect relationships in a network and discover the source of a failure, people do not make proper use of non-failure data (Rouse, 1980). However it is not our intent to review this extensive literature, since it takes us into "problem solving" and away from supervisory control.

Figure 11 suggests that discrepancy between internal model and reality is a good basis for suspecting failure. The question is - how can failures be detected and located?

It is not sufficient to simply detect a discrepancy between a measured variable in a process and a corresponding variable in a model. Given the same inputs, such a discrepancy may be due to a difference between corresponding process and model parameters located anywhere, and not necessarily near the point of comparison. A small difference at one point can integrate to a large difference at another point.

Figure 13 shows one method (Sheridan, 1980) to avoid this problem. This is to use a "disaggregated" model", i.e. component elements of the overall model, all driven on-line by variables measured at corresponding points in the operating real system. Then component element outputs are compared to the real system at corresponding points. If there is sufficient discrepancy over and above expected measurement noise then it is assumed not only that there is a failure but also that this is where the failure is located. This method, however, may not detect some failures.

This method also presupposes that the model elements are accurate representations of reality, i.e. are robust or apply to a wide variety of situations, and indeed that the system variables can be measured at sufficiently many places.

It is often difficult to derive models of systems which are accurate fits to input-output characteristics beyond the normal range. Once any component fails its output becomes abnormal immediately and because of interconnection after a short time many other variables tend to go into the abnormal range, even though their immediate




Figure 13. Comparison of outputs of disaggregated model to corresponding measured variables in actual process to detect and locate failures.

components have not failed. These facts militate against use of fully interconnected models. When the model elements are disconnected as in Figure 13 their inputs can be made to be the same as their corresponding elements in the actual process. Then comparability of outputs is primarily a question of whether the model elements are valid for the values of the variables.

3.6 Experiments With a Particular Model-Based Failure Detection/Location Technique

Implementing the fully disaggregated model scheme outlined in the previous section is seen as tedious for a variety of reasons: Not all actual system variables are easy to measure. Further, there is no limit to how far the model can be disaggregated, making for an unlimited number of required measurements of and comparisons with the actual system. The latter includes both effort and flow variables (e.g. voltage and current, pressure and flow, etc.) at each point. Is there some most efficient way to choose variables to measure and to disaggregate the model? The answer is yes, we think, and is represented in current work by Tsach (1982).

Tsach's technique puts emphasis on selecting one measurement point for each state variable of the actual process, and measuring both effort and flow at that point in the process (Figure 14). The model is set up in the computer so that it can be cut or disaggregated into left and right submodels at that point. Causality direction is established for the actual process effort e_s and flow f_s at that point, one arrow necessarily pointing to the left and one to the right. Assume effort e_s is input to the left submodel. Its flow output f_m is then compared to the flow f_s at the corresponding point in the actual process. Since the power at that state variable is $e_s f_s$ and since e_s is common to both actual process and left-hand model, this f_m to f_s comparison is equivalent to comparing the power transfer at that point as measured by the actual process and the left submodel. If there is no difference we claim that, except for very special circumstances, there is no failure in the actual process to the left of that point. If an observed difference is greater than some allowance for noisy measurement, we conclude that there is a failure somewhere on the left.

A corresponding comparison is made for the right-side model, in this case using the system co-variable f_s as input to the model and e_m as its output to be compared to e_s . Figure 15 shows data from a typical simulation experiment.






Figure 14. Tsach's technique for failure detection/location: cutting the model once for each state variable and using power co-variables to detect discrepancies from actual process. This technique is more sure to detect failures than the straightforward use of disaggregated models.



Pressure (a) and flow (b) measured at point 1 in the system compared with the values predicted by the model. In (b) QS and QM are different, hence the failure is to the right of cut 1.



Torque (a) and angular velocity (b) measured at point 3 in the system compared with the values predicted by the model. In (a) TS and TM are different, hence the failure is to the left of cut 3.

Figure 15. Typical time histories of model-process comparisons using Tsach's technique

By Tsach's technique a first such cut is made somewhere "in the middle" of the process. If a failure is observed on one side or the other, the model on that side is again cut roughly in half, and the failure location is further isolated. This is repeated until the inter-state-variable failure is located to a satisfactory degree. This "split-half" isolation procedure is the most efficient one from an information-theoretical viewpoint.

Note that Tsach's technique works whether power is fixed or continuously changing in time, and accommodates easily to AC power systems where effort-flow products are relatively steady even though the variables themselves are fluctuating rapidly.

Tsach has found ways to cope with problems of noisy measurements, inexact models, highly interconnected systems which are not so easily "cut in the middle", non-linearity, and thermofluid systems where three variables (i.e. pressure, flow and temperature) determine transferred power. His aim is to show that detection/isolations can be done quickly - before the variables go too far into abnormal ranges for which the model is no longer valid.

This whole technique is viewed as a "smart front end" to a man-machine interaction, i.e. where a human operator is alerted if any discrepancy is large enough and then proceeds to participate in or at least confirm the discrepancies. This raises the questions of what further processing of these measurement - model discrepancies should be made, how this should be displayed, and just what the operator's role should be in such a failure detection/location process.

3.7 Deciding When to Stop Learning and Start Acting

Sophisticated new supervisory control systems provide great challenge in understanding their complexities combined with a wealth of tools for sensing, processing and displaying information. One danger is that operators will be inhibited from acting until they have "all the information".

Obviously the longer one waits, the more information becomes known and the more certain a decision will be about what to do. At the same time, the longer one waits the less effective a given control action may become in resolving a problem. Thus, there is a tradeoff decision, one which may have to be made under the stress

of time. Simple discriminant reaction time may be said to result from such a tradeoff. In operations research this is called the "optional stopping problem" (see Sheridan and Ferrell, 1974). There is promise for aiding the operator's optional stopping decision by application of internal models, especially those which can run in fast-time and make predictions.

3.8 The General Planning Problem as Matrix Inversion; Use of Internal Models

Any planning or design activity may be said to be a prediction of what manifest result $\underline{y}$ would be produced by an input $\underline{u}$, given a system, say $\underline{y} = \underline{S} \underline{u}$ where $\underline{S}$ is a matrix of polynomial terms which characterize the input-output equations of the system. Formally this means that $\underline{u} = \underline{S}^{-1} \underline{y}$, that is, given the result to be achieved the proper input is identified by inverting the matrix $\underline{S}$.

Vermeulen (1981) showed how computer-graphic aiding and computer simulation could be combined to aid an operator in making economic and technical decisions when a data-base of complex "influences" $\underline{S}$ was available (in this case a static internal model). He also showed how difficult is the simultaneous trial and error adjustment of the many components of $\underline{u}$ to find the best combination of results $\underline{y}$. His experiments were both with abstract problems and with the realistic task of designing a gas turbine to meet various conditions of cost, power etc. and given constraints due to the laws of physics. But when constraints are in a form such that the $\underline{S}$ matrix can be inverted the task is simple and direct.

While the matrix inversion works only in simple problems, Vermeulen's formalization of the design problem in vector-adjustment terms is a useful concept for supervisor control. The supervisor, in planning his actions, needs to know "if I do $\underline{u}$, what result $\underline{y}$ will occur?" If he has a good simulation of $\underline{S}$ he can try various $\underline{u}$ combinations to see which one is best. But since in general $\underline{u}$ is a vector of very many dimensions, trying all $\underline{u}$ is impractical. Possibly some compromise between the direct trial-and-error simulation experiments and the formal matrix inversion is possible. This problem requires more study.

3.9 Mutual Understanding by Man and Computer of Each Other

In Section 3.5 it was suggested that if operator's and computer's internal models don't agree on what is the current system state, that is a cause for some concern. As we envision it, such a comparison can be made, with the appropriate computer and display aids.

At a deeper level lies the problem of man and computer agreeing on goals and criteria, capabilities for doing the tasks at hand and what they can expect from each other. Moray (1981) has suggested that if the operator has insufficient knowledge of how the computer works he is likely to countermand it, and if human capabilities and limitations are not somehow built into the programs of the HIS, it is likely to countermand the supervisor. Figure 16 shows three "partially intersecting sets" of understanding by the human operator, HIS and TIS - understanding both of one's own function and of the other's computational capabilities.

How to measure and provide this mutual understanding is not altogether clear at present, but it is an issue of great current interest within computer science.

3.10 Group Decision in Supervisory Control

As supervisory control systems become larger aggregations of equipment and become interconnected in more complex ways it is inevitable that there be several or many supervisory controllers interacting. This is not unlike managers of different organizations meeting together to avoid conflict and to coordinate efforts to the end of greater satisfaction for everyone.

The difference is that the "meeting" of supervisors is itself part of the system operation. Rather than sit around a table in a conference room they communicate with each from their normal operating positions. In the future, rather than communicate primarily over telephones where participation must be simultaneous, their participation will be mediated by computer, such that they can:

- 1) leave an electronic message, which may be called up on the computer by the recipient(s) at their convenience;
- 2) jointly call-up and examine the same displays;
- 3) "vote" by indicating judgement of importance of given options or indicate numerical estimate of cost, scalar distance, probability, degree of confidence, etc.;
- 4) possibly have the computer interpret the vote, e.g. take automatic action if there is consensus, take an action which is an average or some other function of the votes, bar individual action where there is insufficient agreement.

Computer-mediated conferencing and messaging systems are now in common use, but they have not been integrated into operating systems of the types we are discussing here.




Figure 16. Partially intersecting sets of understanding by human supervisor, HIS and TIS

4. CONSIDERATIONS OF DISPLAY

4.1 Application of Conventional Human Factors Criteria in Supervisory Control

Changing to supervisory control and using CRT, LED, LCD, synthetic speech and other displays generated by computer does not change the applicability of many human factors criteria developed for old fashioned displays.

First are considerations of detectability or visibility, whether the operator will notice the display in the first place. This means sufficient size, illumination, contrast, and steadiness for a visual display, and proper loudness, pitch and duration for an auditory display. Next are considerations of readability, which have to do with recognition of the code: letters, numbers, scale marks, other symbols for a visual display, speech or special tones for an auditory display. Finally are considerations of interpretability, which concern the operator's finding out what to do.

A given static CRT display is surely worse in these respects as compared to a conventional display of the same size. The CRT's advantage is that it need not be static - it can change and display various images, and for this reason far fewer CRT's are needed than conventional display instruments.

In designing formats most of the same tradeoffs still apply with respect to what should be made available to the operator and where it should be located. For example:

- 1) availability determined by tradeoff between frequency of use vs importance when used (this includes normal use vs emergency use);
- 2) location determined by tradeoff between causality (flow mimic) vs physical co-location vs temporal order of use vs other association in use vs likeness between variables;
- 3) amount of information determined by tradeoff between how much detail vs abstractness, i.e. how much is too much for what purpose.

In other words, the same compromises must be made with CRT's as with conventional displays. Banks et al (1982) and Snyder et al (1980) provide excellent reviews of CRT display considerations.

4.2 New Opportunities and Problems with Computer - Generated CRT Displays

Computer generated cathode-ray tube (CRT) displays are most important primarily because the other computer display technologies are relatively new and little used as yet. CRT displays are common in supervisory control, though of course they may also be present in systems in which there is no significant supervisory control. New technical opportunities become problems for human factors since the question is then how best to use them. Some of these opportunities/problems are:

a) Paging Structures and Access

The main advantage of the CRT is that it can be programmed to display anything, not only changes in text or numbers or symbols or pictures or other graphics relative to one format, but completely different formats for different situations.

Each separate format is called a page, and each page has its own display generation program for translating a set of measured system variables into a continuously or intermittently changing image. There may be many pages which have the same type of graphics (e.g., diagrams of different parts of a system, lists of alarms, sequential text or tables of procedures) while the details of the text, line diagrams, bar charts, and pictures with filled-in color will differ radically from one another.

It is important that the operator have in his head as well as on a placard on the control panel (not in the computer) some diagram and associated instructions for how to access (i.e. what commands to give to call up) various pages.

The complement of accessible pages should be planned by the system designer into some sort of a tree or hierarchical schema, such as is shown in Figure 17. Pages at the top of the hierarchy are those referred to often or most important, especially in emergencies. They are used by the operator to orient himself or to acquire general status information. Some may be dedicated displays, i.e., no other pages can be called up on that CRT. For multiple use CRT's the recall to the top level display (or one of a small set of such displays) should be a very simple operation which is not likely to be forgotten in times of stress.

Page tree structure may change as a function of "mission phase" or "system control mode" such that during takeoff or plant start-up one page tree structure is available, during level flight or plant operation another set, during landing or plant shut-down still a third, etc. Page trees for emergency information (e.g. alarm data) should remain fixed.




Figure 17. Paging hierarchy for computer generated display

There may be various ways to "page down", i.e. call up lower level displays on the tree or schema. One is to have a menu at the bottom of each page which indicates what is available at the next level and what code to key in. Alternatively a cursor may be placed on part of a diagram or in a designated box, etc. to instruct the computer to give a more detailed page focusing on that information. In addition to the step-by-step menu arrangement, which is always good to have for operators who forget or get confused, there should be some keyboard commands by which the experienced operator can call up displays directly - without stepping through the tree.

b) Integration and Format

Instead of having to present related variables on separate conventional display instruments which, at best, are placed side by side, the CRT can combine these variables into an integrated display. This means that an aircraft primary flight CRT display can combine attitude, altitude, heading, speed, rate of climb and even provide a picture of the runway for landing, all on one picture. Or it can show temperature and pressure of water in a nuclear reactor as a point, moving relative to a saturation curve (Figure 18), rather than require operators to put that information together from separate displays (which operators had to do at Three-Mile Island).

While such integration is still an art, it is clear that it is a critical issue in supervisory control. In process control there are efforts to develop CRT "system state vector" displays which somehow provide an overview of plant state, and indicate how far any one aspect may be deviating from the normal. One approach is the polar polygon plot, where the radial deviation of each of, say eight, vertices from the center shows the relative magnitude of that variable, normality being a regular polygon (Figure 19). A more radical approach is the Chernoff cartoon face, where eyes, nose, mouth or other features each represent different variables, and size of eyes, slant, etc. represent the values of the variables (Figure 20).

Diagrams of systems can be drawn with pictorial symbols, lines representing fluid or force or information, colors which change with status, texture which moves to show activity, flashing to get attention, text or numbers superimposed, etc. The flexibility of such animated diagramming is so great that the designer's tendency is to pack too much into one display. A useful rule-of-thumb is that at least 80%




Figure 18. A simple "integrated display" showing the variation of the coolant (dot) in a closed power system relative to the critical saturation line.




Figure 19. The polygon state vector plot. Radial deviation of any component variable outside the "safe zone" annulus is quickly apparent. Variables 1-6 are the same in both 6 and 8 variable plots




Figure 20. The Chernoff face. Different facial features of computer-generated image represent different variables. In upper figures note change in nose from left to right, then (lower left) change in both eyes and eyebrows, and finally (lower right) further change in eyebrows

of the area of such a diagram should be "background". Some computer aided display design programs have been developed which allow symbols to be designed, sized, located on the page, interconnected, labeled, etc. with great ease and flexibility in modification.

As noted earlier, the integrated diagram, picture or alphanumeric text display should, insofar as possible, try to capture and correlate with the operator's internal model, "problem space" or "task image" (as best it can be understood).

A popular characterization of the good display is that it be "transparent", in the sense of a window on the system variables. Metaphorically the operator wants to "look through" the display to the systems, not be obstructed in time and understanding by its "opacity".

c) User-Scaled Trends and User-Adjustment

The CRT allows flexibility in plotting any variable on any time scale and any magnitude scale, comparing it to any other variable or set of variables, making cross-plots, smoothing, coding by color, etc. Various prediction routines can be used, from simple Taylor series to predictions based on models as suggested in 3.3 and 4.8.

While there should be standard or "default" plot formats which the operator can fall back on, there will be times when he wants to deviate from these. However, given too much freedom to adjust formats, diagrams, text and other displays to suit his whim and fancy, the operator may become distracted and the display may be left in an inappropriate state for other operators. Thus there is some reasonable limit to display flexibility. Also, perhaps, after some time period in a special mode, displays should automatically return to some standard.

d) Multi-Variable Data Search

The supervisory operator may upon occasion need to search a multi-attribute data base, where a number of pictures or descriptions of events or objects are represented in memory, coded by their attributes (e.g. time, cost, frequency of use, performance, etc.) in a multi-dimensional array. If the operator does not know the address (i.e. the attribute coordinates) precisely he must search visually. Search through a two, or even a three-dimensional array is straightforward enough on a CRT, but four or more attributes poses interesting problems.

One approach is to display the "center object plus nearest neighbors" arranged axially on spokes of a wheel (one spoke for each attribute dimension) or "center object plus extremes" for each attribute (Barrett, 1981, Figure 21). Another approach is to move the search with respect a given attribute by naming or pointing to the attribute (ibid.). Still another approach is to select two (or three) dimensions in a reference square (or cube) and move a pointer around in a reference space, as a means of calling up the center and nearest neighbors in this low-order array (Knepp, 1981, Figure 22).

e) "Zoom" Control of Degree-of-Abstraction



Paging up and down in a tree to access different levels of detail as well as parts of a flow chart was discussed earlier. Goodstein and Rasmussen (1981) and others have proposed allowing the operator to vary the display in degree-of-abstraction in accordance with his cognitive needs. This correlates with whether the required behavior is knowledge-based, rule-based or skill-based.

A knowledge-based display would show system status relative to goals and criteria. It would provide the "big picture". It would indicate structure. It might allow the operator to query a data base on why certain equipment was designed as it was, what failures there have been, what other approaches there are, what administrative information bears on a particular problem, etc. Knowledge-based displays call for knowledge-based behavior, which usually takes time and requires the right talent and disposition.


A rule-based display would show status or system state relative to procedural steps. It is concerned with task progress in a time-line or flow chart, with identification of trends or patterns, with success and failure in completing subtasks.

A skill-based display would show instantaneous values of individual variables, position on a trajectory, and what to do immediately. The focus is narrow and detailed.

We have discussed displays in each of Rasmussen's categories. "Zooming" may mean turning one's attention from one such type of display to another. The "zoom" control may not be continuous, but in steps. One hope is to allow the operator to be able to simply order a zoom "up" and have the computer decide what more abstract, bigger picture display to provide. Another approach would be to have a standard order or tree, as was discussed for level of detail of equipment or flow diagram.



Single object with change direction indicated verbally



Single object with change direction indicated pictorially.

Figure 21. Barrett's computer-generated displays for visual search in multi-attribute space. Upper display is "Center object plus nearest neighbors" in four dimensions by which a bolt varies (length, diameter, head size, thread size). With lower displays user touches label (left) or extremum of scale (right) to make center object (bolt) change in desired attribute and direction




Figure 22 Knepp "reference cube" display for search in multi-attribute data base. Rotation of both viewpoint and position in reference multi-attribute space (left) determines what is displayed in X-Y object display (right).

f) Too Many Options: Killing the Operator with Kindness

"Future shock" in supervisory control is a real hazard. An operator can be confronted with too many display options, to the point where he doesn't remember how to access or interpret them, or where similar system events appear different because on different occurrences he has observed them through different displays. Both such situations could lead to unwarranted operator errors.

Yet another hazard is that operators become fascinated and distracted by the novelty and "gadgetry" aspects of the displays themselves and pay too little attention to what they are telling him about the system.

4.3 Special Teleoperator Display Problems

Since this report is concerned primarily with undersea teleoperation applications of supervisory control, a number of display problems that are more particular to this application will be discussed.

a) Bandwidth Constraints

The communication channel between TIS and HIS may be bandlimited, e.g. ship-to-undersea-teleoperator acoustic transmission (1000-50,000 bits/sec.) depending on distance and other factors), earth to outer space radio transmission, or ordinary telephone transmission. While 1-50 K bits/sec. is sufficient for audio and most control signals it is not sufficient for video transmission. Regular broadcast video requires up to 50,000,000 bits per second, where each of 250,000 pixels (picture elements) must be refreshed 30 times per second with 6 to 8 bits of grayscale. This is 1000 to 10,000 times more bits per second than good audio channels. Since the product of pixels resolution, frame-rate and bits of grayscale equals bit rate, one can reduce any or a combination of these factors to send video over a constrained channel. For example a 30 x 30 pixel picture with 2 bits of gray (4 levels) at 2 frames per second would be approximately 4000 bits per second.

In view of these problems, why not always use electrical cable with its high bandwidth? The reason is that long cables, e.g. several miles, get very heavy, and even if they are made neutrally buoyant, their inertia and drag can be huge - especially in view of ocean currents. Further, they tend to become tangled in structures, rocks etc. on the ocean bottom.

b) Teleproprioception

This coined word means self (proprio) awareness (ception) at a distance (tele), in other words, the operator knowing where in position and orientation his manipulator and/or vehicle are relative to the environment.

This is a problem because the operator is using discrete switches or a continuous joystick to control video feedback on where his arm or vehicle are in space. The video usually has a narrow field of view, the light can be poor to negligible in turbid water, and shadows on visually unfamiliar terrain can be confusing. Further, even though the operator can see what is on his video display, unless part of the vehicle or the base of the manipulator is in the view he may have no sense of which way the camera is being pointed.

For manipulator control a master-slave position servo is helpful because the operator can see and feel the position of the master. For vehicle control a "local model" in the form of a map, updated from time to time with estimated vehicle position, is helpful in the same way.

c) Tele-touch

When water is very turbid video simply will not work. Then one is left with sonar (pingers, side-scan or two-dimensional acoustic imaging) or just "feeling around" with a manipulator. Sonar is proven, but its resolution is limited. Manipulator positioning can be accurate, but touch sensors and displays are still under development. Some special teletouch techniques are described in 4.7.

d) Superposing Video and Graphics

The head-up landing display in the aircraft cockpit permits the pilot to look through the windscreen to the runway and in the same location at the same visual accommodation see a projection (using a prism or half-silvered mirror) of computer-generated graphics. The latter might include symbols for runway, attitude, altitude, airspeed, etc. Thus the pilot can visually search through the clouds for the actual runway at the same time he flies the aircraft relative to the computer-generated CRT display, without having to shift his eyes in visual angle and distance accommodation.

Similarly, with human operators of teleoperators, the video is important but may at times have to be augmented by computer-generated displays. In this case it is not a matter of looking through a windscreen to a real environment but more having to move the eyes from one display to another. Thus the superposition of the two images would seem to be helpful, but it remains to determine experimentally how best to do this.

4.4 Experimental Evaluation of Frame-rate, Resolution and Grayscale Constraints.

As noted in section 4.3a above, to avoid problems of the tether becoming a large drag on the submersible vehicle and/or getting tangled up in structures that one wishes to inspect remotely, one may employ acoustic communication. Even if a tether is dropped from a surface vessel down to within a few hundred feet of the submersible, acoustic signal transmission for the remainder of the distance can circumvent the problems cited above. However this can only be done at the cost of having to reduce the bandwidth considerably relative to that for a wire (tether).

Thus one is left asking, for a given fixed communication bandwidth, how best to trade between the three variables of frame-rate (frames per second), resolution (pixels per frame) and grayscale (bits per pixel) the product of which is bandwidth (bits per second).

These tradeoffs were studied by Ranadive (1979) in the context of master-slave manipulation. The experimental subject was asked to perform two remote manipulation tasks using a video display as his only feedback and our Argonne E2 seven-degree-of-freedom servo manipulator (in this case with force reflection turned off). Figure 23 illustrates the experimental situation.

The first task was to locate a nut on a fixed bolt or knob and take it off by unscrewing it. (We abbreviate this task "TON" for take-off-nut). The second task was to pick up a cylinder and place it sequentially within the bounds of three fixed squares on the table which were numbered 1, 2, and 3, where the order of the placement, e.g. 3-1-2, was randomly drawn for each new trial. (We abbreviate this task "1-2-3"). Performance on each task was simply defined as the inverse time required to do that task correctly, and combined performance was the average of these inverse times.




Figure 23. Experimental configuration for Ranadive frame-rate, resolution, grayscale tradeoff experiments



Figure 24. The same picture at various degrees of resolution (pixels)

The video display was systematically degraded with a special electronic device which allowed frame-rate to be adjusted to 28, 16, 8 or 4 frames per second, resolution to be adjusted to 128, 64, 32 or 16 pixels linear resolution, and grayscale to be adjusted to 4, 3, 2 or 1 bits per pixel (i.e. 16, 8, 4 or 2 levels of CRT intensity). Figure 24 shows the effect of resolution reduction.




Figure 25. Learning curves in Ranadive experiment

Two subjects were used, both engineering students. They were trained for 10 hours in all combinations of display tasks and visual variables. When subjects first saw the video pictures with which they had to perform remote manipulation tasks, they refused to believe that they could succeed. Much to their surprise, however, they discovered that they were able to perform with a considerably degraded picture. Figure 25 illustrates learning curves for the two subjects. During the data collection phase of the experiment subjects were allowed to practice on each display combination until "ready".

The data collection runs were ordered so that two of the three video variables were kept constant while the third was varied randomly among the levels for that variable. Ten times were collected (ten trials were run) for each combination (each data point).

Figure 26 shows the results. On the top row are shown the performance effects of frame-rate, resolution and grayscale while holding the other variables constant. Note that for frame-rate beyond 16 frames per second improvement depends on resolution and grayscale; performance improves smoothly for increases in resolution; for grayscale there is no improvement beyond 2 bits if the frame-rate is high enough.

On the bottom row constant level-of-performance tradeoffs (in this case using the TON task only) are shown for each of the three pairs of video variables. These iso-performance curves (solid lines) are compared to iso-transmission lines, i.e. combinations of the two parameters which produce constant bits per second. It is seen that there is a remarkable correspondence. This means that for this experiment, and within the range of video variables employed, man-machine performance corresponds roughly to bits per second of the display, regardless of the particular combination of frame-rate, resolution or grayscale.

Another result, though not tested systematically, was that subjectively much more noise appeared on each video picture at the slowest frame-rates than at faster frame-rates. It is believed that this was due to visual-psychological smoothing rather than anything electronic occurring at higher frame-rates.




Figure 26. Results from Ranadive experiment. The top row shows performance as one of the three variables is adjusted, the other two kept fixed. The bottom row shows isoperformance and isoinformation transmission tradeoffs for two variables at a time, the third held constant

4.5 Operator-Adjustable Frame-rate, Resolution and Grayscale Under Fixed Bandwidth Constraints

Assume limited-bandwidth acoustic transmission is to be used as the means for communication between a human operator and an undersea teleoperator. It is clear that video will pose a far greater bandwidth requirement than other signals, probably more than all other signals combined. Therefore it seems reasonable to allocate fixed channels to other variables as required, then to reserve the bandwidth remainder to the combination of video signal requirements. That is, frame-rate, resolution and grayscale would not each have fixed bandwidth allocations; rather, provision would be made to trade-off between these as required, retaining their product as close as possible to the maximum.

To make this idea more understandable, assume that a given human operator of a teleoperator needs to get an accurate picture of a static object. He would like high resolution and sufficient grayscale, but frame-rate could be anything. In contrast, suppose the operator needed to monitor the speed at which a well known object moved against a background. Only enough resolution and grayscale would be necessary to get a good definition of what is object and what is background, but frame rate would have to be high. Either condition could be obtained by adjustment.

Deghuae (1980) used an experimental computer-based aiding device which allowed the operator to make this three-way adjustment in-situ, i.e. he could modify the FRG tradeoff himself while performing a master-slave manipulation task of the type performed in Randive's experiments. In particular, the same master-slave manipulator was used with the force feedback turned off, and the same take-off-nut (TON) task, but a many-peg-removal task was chosen instead of the 1-2-3 task. Scoring was inverse time, the same as in Randive's experiments.

Four subjects were trained about ten hours each on each of the tasks. The same levels of the video variables were used as before. Two maximum bit rates (products of frames per sec, pixels per frame and bits per pixel) were used, one being 11500 bits per second, the other twice that or 23000 bits per second. The subject used three keys to call for any FRG combination, up to the maximum. When one factor was increased the other two were automatically decreased to keep the product at the constant maximum. Each subject, for every combination of task and maximum bit rate, performed both with and without the in-situ tradeoff adjustment capability. There were four trials for every cell of the experimental design. Data were analyzed by analysis of variance.

As might be expected, the use of the tradeoff control was significant at the 95% level. However both the task main-effect and the task-subject interaction were significant at the 99% level, a result not particularly surprising. What was more surprising was that the two maximum bit rates did not produce significantly different performance.

There was much variability in performance due simply to the fact that the visual interpretation time was extensive, and the real-time continual decision task of how to set the FRG combination added to this. It is believed that the means of making this adjustment can be better "human engineered", and that this would reduce variability and improve performance. Similarly the lighting was seen to be a critical factor, where amount of light affected grayscale adjustment and shadows provided important cues.

A principal result of this study was confirmation that with some training and some patience an operator can remove a nut with a remote manipulator using video of only 10^4 bits per second and with no force or tactile feedback. From the results an important special use of the adjustment became apparent in this case, namely to periodically but briefly increase resolution and grayscale at minimum frame rate in order to get confirmation that the peg was in the hole, or that another critical task phase had been achieved.

Use of this device is an important aspect of supervisory control, where the computer aid mediates the operator's instructions to provide, in this case, the best display (rather than control per se). This is loop 8 (of Figure 2) working in conjunction with loop 2.

4.6 Demonstration of Computer-Graphic Arm-Hand Teleproprioception Aids

As noted in 4.3b, whether control is supervisory or purely manual, teleproprioception is a particularly difficult problem in remote manipulation. This is due both to the lack of depth cues and to the lack of reference frame for both orientation and translation movements of the manipulator.

Winey (1981) developed a clever computer-graphic display which has two functions:

- 1) It allows a neat controlled experimental measure of how different forms of teleproprioception affect performance.
- 2) It offers promise for in-situ use as an auxiliary to the normal video display of the actual remote hand.

Winey's computer-generated arm-hand is shown in Figure 27. It was programmed on the DEC 11/34 computer and Megatek 7000 vector-graphic display with hardware rotation capability. The position signals from the master arm which normally are fed only to the slave were also fed to a geometric model of the slave, which model drove a graphic likeness on the Megatek CRT display. Of particular interest here are the added features Winey used in conjunction with this model and display, and how they aid teleproprioception.

A first novel feature was the generation of multiple orthogonal projections, as shown in Figure 27, instead of just one view. Two such views are sufficient in theory to provide all the information (but any projection is trivial to generate once the model is worked out). Note that the precise geometry of the remote arm is known and the control inputs are known. Thus there is little uncertainty about the configuration of the actual remote arm-hand relative to its own base - unless the control system is functioning improperly, which would be indicated by large position error feedback signals from the remote arm. In other words, normally there is little need for the human operator to get from the video display any teleproprioceptive information about the position-orientation of the remote arm-hand relative to its own reference frame.

Note also that, through use of the fast hardware rotation capability of the vector-graphic display, the operator may in effect position himself (his viewpoint) anywhere in the sphere surrounding the arm-hand.

This model-plus-graphic-display is exactly the equivalent of the "observer" in automatic control. It is a computer-based internal model driven by control inputs, any states of which can be observed by a person!




Figure 27. Winey's computer-generated arm-hand in two views

A second novel feature of Winey's display was the addition of "shadows" and a surrounding "box", as shown in Figure 28, to aid teleproprioception.

A third novel feature was the addition of force reflection from a modeled environment. Here, in addition to the computer model of the arm-hand geometry, Winey added a model of the deformation of environmental surfaces or objects in correspondence to the degree to which some point or points on the arm-hand crossed over their boundaries (Figure 29). He then fed corresponding force-reflection signals to the master. In this way the human operator could feel the elastic forces of surfaces pushed on just as he would if the actual slave pushed on and deformed actual surfaces. In the present case, however, no actual surface exists. The operator is manipulating and actually feeling the forces from a computer model!

A fourth display, used mostly for comparison purposes, was a separate indication (the length of a bar) of the distance between the hand and the sphere. This he called a "proximity indicator".

An experiment was performed to test whether certain of these displays provided sufficient cues to be able to reach out and grasp a sphere as shown in Figure 27, with no other visual cues (i.e. no actual sphere, no actual TV picture. Each of four subjects trained for two hours, then performed 80 repetitions of grasping the sphere under four different computer-generated display conditions: (1) single arm in a box with a shadow; (2) single arm plus hand-to-sphere proximity indicator; (3) front and side (orthogonal) projections of arm; (4) proximity indicator only. As noted in Table 3 of the time-to-grasp results, there were two versions of the task, one where the sphere was static, the second where it was moving. The force feedback was not employed in these experiments.

From Table 3 it is evident that the task was quite easy to do with the arm-plus-shadow, the arm-plus-proximity indicator, and the two-orthogonal-arm view display, with the latter being slightly quicker than the others. Using only the proximity indicator took far longer. Both means and standard deviations were constant across the four subjects. By analysis of variance quite significant differences between subjects and between the two forms of the task were evident, in addition to the (most significant) differences between display types.




Figure 28. Wineys use of "shadows" and surrounding box to aid teleproprioception




Figure 29. Winey's experiment with touch (force feedback from a computer model)

Table 3. Results of Winey's experiment on time to grasp sphere under different display aids

Task: Grasp a Sphere which is:

Computergraphic Display Aid	Subject	Stationary		Moving	
		<u>Avg</u>	<u>Std.Dev.</u>	<u>Avg.</u>	<u>Std.Dev.</u>
Shadow	1	2.87	1.90	3.06	1.35
	2	3.89	0.95	3.09	1.47
	3	3.74	2.05	4.74	2.93
	4	4.93	1.83	5.12	4.00
Front and Proximity	1	3.83	1.76	3.95	1.57
	2	3.77	1.07	3.54	1.26
	3	3.73	1.26	5.56	5.73
	4	3.91	2.90	4.43	2.51
Front and Side	1	3.08	1.14	3.60	1.30
	2	2.34	0.68	2.41	0.76
	3	2.83	1.11	3.77	1.67
	4	3.26	0.99	4.28	2.39
Proximity only (no arm)	1	13.54	10.74	32.45	24.40
	2	12.55	8.16	23.65	18.77
	3	14.07	11.21	35.36	26.51
	4	20.90	11.21	34.44	36.14

Beyond usefulness as an experimental device we believe such display aiding techniques have great potential in real-time teleoperator control, especially where the video channel is of low bandwidth, where there is signal transmission time delay, where water is turbid, or where the manipulator hand is otherwise occluded from view. In these cases by observing such a display aid in addition to the normal video the operator can get a continuous, clear view of the manipulator arm and hand configuration from any viewpoint, and he can get reference distances from any known object. This feedback will generally be more useful than that from the video display - except, of course, for the relation of the arm-hand to unknown or non-fixed environmental objects.

Finally, it may be useful to include the whole vehicle in the simulation, as shown in Figure 30.

As with the display aids described in the previous section this form of aiding for the display may be considered an important part of supervisory control.

4.7 Demonstration of Computer-Graphic Aids for Tele-Touch

Using the vector-graphic display developed by Winey, Fyler (1981) created a novel means for tactile probing and discovery of the shape of an unknown object or environment. This technique offers promise where the water is so turbid that video is useless (and because acoustic imaging is as yet unavailable). It is the analog of a blind person probing in the dark by repeatedly touching at different points on an object or environmental surface in front of him and gradually building up a "mental image" of what is there, continually guiding his touching activity on the basis of what he discovers.

In performing "tele - touch" with a master-slave remote manipulator, if there were no dynamics and if force feedback were perfect it might be asserted that building up the necessary "mental image" would be no different than direct manual groping in a dark room. However every operator knows that is not reality; the master-slave manipulation itself is sufficiently cumbersome that one quickly loses track of where contact has recently been made and what the arm's trajectory has been.




Figure 30. Computer generated display of real-time vehicle plus manipulator simulation


In performing tele-touch where a computer is determining the trajectory rather than a human operator's hand movements guiding a master, building up the "mental image" is still more difficult.

Fyler designed a unique touch-probe, a mechanical device which closes an electrical contact when it encounters a slight force from any direction. Then he programmed the 11/34 computer to determine and store the cartesian coordinates where any contact (touch) is made. He displayed on the Megatek screen a projection from any viewpoint of cumulative points so stored. The operator can make no sense of such a display so long as the points are fixed. But the instant the image of points is rotated the shape and orientation of the one or more surfaces on which the contacts were established becomes immediately evident. What is a "mental image" in the case of direct manual grasping or touching becomes an explicit visual image. Figure 31 provides some (unfortunately static) examples of such displays.


As more points are added, the definition of the surface or object becomes more apparent. It helps somewhat to have the computer connect adjacent points with lines so that the best available "image" in three dimensions is a polyhedron and its planar projection is a polygon (or, if both front and back surfaces of an object are touched, two overlapping polygons). When rotation is effected the polyhedron immediately becomes evident. Rotation may be at a constant rate - usually around an axis near to or transecting the surface or object of interest - or may be controlled manually by a track-ball.

As contacts are made, points are added to the display, and what started out to be a polyhedron with few vertices and faces becomes a smooth surface, or a recognizable object. The first few contacts between the manipulator probe and environment are made more or less at random. However, as the polyhedron takes on form, it is evident to the operator where to place the next few probes to provide the most discrimination and not waste probing effort and time.


Another display trick Fyler demonstrated was to put the polyhedron into the Lexidata raster-graphic display generator's look-up table in such a way that the orientation of any facet of the polyhedron is determined. Then he "illuminated" different facets of the polyhedron on the raster display as a function of



20 Points



50 Points



100 Points

Figure 31. Random touch points on a sphere which generate a polyhedron.
As sphere is rotated the shape is easily perceived

the orientation of each facet - as if the sun or light source were at one angle shining on a polyhedron (Fig. 32). Again the operator was provided a trackball, in this case to let him move the apparant light source to any radial position surrounding the object, the polyhedron in this case being fixed in orientation, not rotating.

4.8 Experiments with Computer-Graphic Predictor Displays for Remote Vehicle Control with Transmission Delay and Slow Frame Rate

Another form of computer-based display aid is the predictor display. This is a technique in which a computer model of the controlled vehicle or process is repetitively set to the present state of the actual system, including the present control input, then allowed to run in fast-time, say 100 times real-time, for some few seconds before it is updated with new initial conditions. During each fast-time "run", its response is traced out in a display as a prediction of what will happen over the next time interval (say several minutes) "if I keep doing what I'm doing now". The technique is about thirty years old, has been much discussed in the human factors literature (see, e.g. Kelley, 1968), and has been applied some to continuous control of ships and submarines. It still holds promise for a variety of future applications.

Sheridan and Verplank (1978) used a predictor display for remote vehicle control when there is significant transmission delay (say more than 0.5 seconds) and slow frame-rate (say less than one frame per four seconds). Both of the latter conditions are likely to be present with long distance acoustic communication.

A random terrain was generated and displayed in perspective, updated every 8 seconds (Figure 33). A predictor symbol appeared on the terrain, continuously changing as the experimental subject controlled the motion of the vehicle, through a one second time delay. Front-back velocity control was accomplished through corresponding position a joystick, and turn rate by the left-right position of the joystick. Also superposed on the static terrain picture was a prediction of the viewpoint for the next static picture and an outline of its field of view. This reduced the otherwise considerable confusion about how the static picture changed from one frame to the next, and served as a guide for keeping the vehicle within the available field of view.



Figure 32. Polyhedron can be "illuminated" from one angle, and this angle can be moved arbitrarily with a track-ball

PREDICTOR DISPLAY




Figure 33 Simulation experiments with predictor displays (from Sheridan and Verblank, 1978). Slow-frame-rate pictures (8 seconds per frame) were simulated by computer-displayed terrain. The path to be followed was a ridge in the terrain. A moving predictor symbol (perspective square) was superimposed on the static picture of terrain. The point from which the next picture was taken was indicated with a "table" (square with four legs) and the field of view was shown with dotted lines.

Verplank showed that by use of the above two display symbols together, relative to the periodically updated static (but always out of date) terrain picture, subjects could maintain speed with essentially continuous control. By contrast, without the predictor they could only move extremely slowly without going unstable.

4.9 Emergency Displays

If the computer has sufficient evidence of system abnormality, the first thing is to get the supervisor's attention. The loudest klaxon horn may not be the best. A sufficiently loud but not terrifying attention-getter followed by computer-generated speech directing the operator's attention to the proper visual display is probably better.

All too often the alarm signal is a false alarm in the sense that either (a) it is triggered by an expected occurrence as part of an otherwise irregular set of events intentionally initiated by the operator, (b) it is an event which is an obvious and necessarily concomitant of a more critical event, or (c) the alarm logic itself is faulty. More sophisticated computer systems can be programmed with sufficient logic to suppress alarms under the (a) and (b) conditions and to check-up on and least raise the question of uncertainty if there is some evidence of logic failure (c).

There is always the danger of multiple simultaneous alarms, where the operator simply loses track of one of them in pursuing the others. This is why it is useful to have computer generated lists with importance priority and/or time of occurrence specified. Flashing or changing colors of CRT displays or parts of displays (e.g. component symbols on diagrams) is another approach.

Various proposals have been made for some operator-adjustable criteria for alarm logic. This would work in the following way. If the operator had little to do he could set the alarm threshold to be low; the slightest abnormality could trigger an alarm. If the operator were very pressed, on the other hand, he could raise the threshold, and only be bothered by alarms of greatest urgency.

Computerized alarm systems of the future should be able to provide not only the basic alarm signal, but also tell how serious are the circumstances, how confident

is the evidence, how urgent is the response, and perhaps even make some suggestions, possibly with computer-generated speech, on where to go for more information or what to do in response.

4.10 Experiments in Smart-Display of Failure Detection/Location

Section 3.6 described Tsach's technique for detecting and locating failures by making comparisons between measured variables in an operating system and corresponding variables in an on-line model. To make this cognition aid to the supervisor effective the information must be properly displayed to him. This can be said of any such cognitive aid. The form of display can be as critical as the "smartness" of the computation.

An experiment was designed (Tsach, 1982) to test how quickly and reliably a person could utilize a display based on Tsach's technique. For a second order system both system (noisy measurements) and model (no noise) outputs resulting from the same input were plotted against time on the horizontal axis and displayed as "moving windows" offset vertically just enough to separate the traces. Subjects pushed a button when they decided the two traces had become sufficiently different that a failure had occurred (actually a parameter of the second order system had suddenly changed).

Eight subjects did twenty runs on each of a $3 \times 4 \times 2$ array of conditions:

- There were three types of smoothing (averaging) on the trace displays
 - 1) the system output was smoothed, the model not
 - 2) the model output was smoothed, the system not
 - 3) both were smoothed
- There were four categories of displays in addition to the traces
 - 1) none
 - 2) an odds ratio derived by conventional Bayesian analysis from short time samples of the two traces (before smoothing). Every few seconds this odds ratio would change, based upon a new pair of samples.
 - 3) indication of how long the odds ratio (probability of failure divided by probability of no failure) has exceeded unity.
 - 4) (2) and (3) displayed together

- Half the time the second order model was cut so that the two state variables were included in the relevant subsystem and submodel. Half the time the cut left the subsystem and submodel as simple coefficients. In the former case the output was naturally much smoother than the input and a little delayed. In the latter case there was only an amplitude ratio change.

Resulting decision times were not significantly different between "no added display" and "time the odds ratio has been greater than one". However displaying the odds ratio or displaying both odds ratio and time made for significantly quicker responses.

With respect to smoothing of the signal traces, if the subsystem and submodel traces contained inherent smoothing (i.e. the state variables were included) additional smoothing did not help. When they were only coefficients and when the system trace alone was smoothed, or the model alone, detection was significantly quicker. It was quicker still for the coefficients when both system and model traces were smoothed.

When the odds ratio was used there were very few errors (misses or false alarms). In the worst condition (no smoothing on the traces and no supplementary display) the hit rate was 0.89 and the false alarm rate 0.14.

Our conclusion is that some smoothing and the use of a supplementary odds ratio derived from a Bayesian calculation makes the display of Tsach's failure detection/location method easy and quick to comprehend and more reliable.

5. CONSIDERATIONS OF RESPONDING

5.1 Extension of Conventional Manual Control Theory To Supervisory Control

In this report we have defined supervisory control tasks in a very general way to mean primarily control of discontinuous multi-separate subtasks and include human operator planning, teaching (programming), monitoring (including minor on-line adjustments), intervention (major take-over and shut down modifications) and learning. If "supervisory control" is restricted to continuous control systems (or sampled more-or-less continuous systems) and only the minor parameter adjustment function of the operator is considered, the resulting system becomes amenable to modeling by direct extension of manual control theory.

Kok and Van Wijk (1977, 78) made an extensive study of a system (Figure 34) which can be considered an extension of the optimal control system, Figure 9. In their system the human operator, based on a threshold function of various observed state variables, the innovations (discrepancy) signal and costs, makes set point changes in the optimal controller. They ran experiments with operators supervising automatic control of a supertanker and got good predictions of when set point changes would be made.

White (1980) extended this model to human operator data, making similar adjustments to a chemical plant.

Muralidharan and Baron (1980), in a similar vein, used an intermittently adjusted optimal control model for human supervision of remotely piloted vehicles, where the operator must both decide when to intervene and what discrete corrective action to make.

These studies showed that observation noise, relative costs and the number of tasks which demand the sharing of the operator's attention all affect performance.

Sheridan (1976) proposed an expected-value maximization approach to model such systems, both for sensing and for controlling. Superficially such a model can be made to look like optimal control, though it is not.




Figure 34. The Kok and Van Wijk model for supervisory control of a supertanker

There should be continued progress on such semi-continuous system adjustment models. However the writer is doubtful that such models can cross the chasm of essential discontinuity to cope with the range of supervisory control tasks described at the beginning of this report. For example simple assembly or other manipulation tasks and most process control tasks, when looked at over a long enough time period, are simply not continuous tracking tasks.

5.2 New Opportunities and Problems with Computer-Aided Supervisory Command and Control

a) Teaching and Initialization

Sometimes the human supervisor must teach the computer ab initio by writing a fresh program to tell the HIS/TIS when to get what sensor data, when and how to interpret commands from the operator, how to use stored data for smoothing, where to branch depending on what contingencies, what control decisions to make based on all of the preceding and how the actuators should interpret those control signals.

Much more likely, for any given supervisory control system, this software is already in place, and the teaching task to be done is more one of "initializing" already written programs by specifying parameters and conditions to make those programs fit a given context and run to generate a display, to modify a control, etc.

The distinction between these two is not clean. Mostly supervisory controllers do not think of themselves as programmers, and mostly they do not do ab-initio-teaching. But as systems become more complex the imparting of new knowledge from the skilled operator to the computer in a form which the computer can understand will be a major responsibility of human operators.

b) Interoceptive and Exteroceptive Control

From the Latin, these mean respectively control by sensing (feedback) from inside the controlled vehicle or process (aircraft engine thrust, potentiometers in the joints of manipulator arm) and from outside the vehicle or process (distance to ground, visual picture of object to be manipulated).

The former is sometimes called "open loop" control, where e.g. a manipulator is blind to collisions with obstacles. Interoceptive sensing tends to be simple, reliable.

fast and cheap (except for inertial sensing) compared with exteroceptive sensing such as vision, touch, radar, etc. Interoceptive sensing may be necessary, whether or not exteroceptive sensing is used, to provide good stability and fast control of positioning movements. Exteroceptive sensing may be mandatory to align a vehicle or manipulator to external objects. Ultimately it is not usually a question of one or the other.

c) Analog vs Symbolic Control Devices and Codes

It is important that operation of hand, foot and voice controller devices be "natural". That may mean several things. First, it means that the control device may be correctly selected as an analogic or symbolic control. If commands are continuous in magnitude an analogic controller (knob, joystick, multi-degree-of-freedom master arm) should be used and its direction of motion should correspond both with the direction of motion of the response as seen in the corresponding display as well as with the population stereotype of the system response. That is, moving a speed control "up" should make the speed indicator go up and should increase speed. This three-way-geometric-semantic correspondence is called control-display compatibility. The ultimate analogic control device is a replica of the system, which the operator configures to be the way he wants the system to reconfigure itself (system state) - a map or three-dimensional model, which may be called a "local internal model".

If the command is a categorical selection, then a symbolic control (specialized pushbuttons or switches or general keyboard) is appropriate. The trend in supervisory control is toward symbolic control devices, since they are more suitable for knowledge-based behavior; analogic devices are more suitable for skill-based behavior (Sheridan and Verplank, 1978).

A full alphanumeric keyboard obviously allows a great range of symbolic statements, while isolated pushbuttons with special labels are not so suited. Ferrell (1978) showed that it is better in simple constrained tasks to use a simple constrained vocabulary. As the task becomes more complex the vocabulary, that is the software, should have a structure which serves as a mnemonic and a guide to the user. For example Chu et. al. (1980) demonstrated a manipulator command language structured in correspondence to the manipulation task.

d) Virtual Control Panels and Devices

The CRT is usually thought of as a display, but it can also be a control. If the picture or graphics generated on the CRT is that of a control panel, and if a transparent touch panel (measuring the x and y position of applied finger pressure) is overlaid on the CRT face, then the operator can operate the computer-generated control panel (a virtual control panel). For example, a computer-generated pocket calculator's display will indicate the result of a series of virtual button-presses on the CRT overlay (Figure 35). A light-pen or tablet or mouse can also be used to control a cursor on the virtual control, but these are slightly more cumbersome. It is likely that virtual control panels will come into increasingly widespread use.

The virtual control panel needs to give the operator some feedback that he has "made contact", some equivalent of the force detent on a push button. This can be provided either by a sound or by the display itself changing in some way.

Paging is useful with virtual control panels as well as with computer-generated displays. In some cases, continued display of the same control panel is appropriate, with only numbers or pointers or bars changing. In other cases, once the operator has entered certain required data using the proper control panel, it then would be appropriate that he go to a different control panel and interact. If the computer knows that it can anticipate his need and bring him the new control panel immediately on the same CRT. This is a form of "paging down", as discussed earlier in conjunction with information seeking, but in the present case it is a control procedure. In fact, if the CRT virtual control panel is combined with the CRT information display, in the process of "paging down" the distinction between display and control device essentially disappears.

Virtual control panels may be formatted to the operator's taste in the same way as information displays. Virtual control panels may soon be used in conjunction with computer recognition of voice commands. These technologies are new; insufficient experimental work has been done.

e) Authority vs Deference, Trading vs Sharing

Another problem which arises in the context of supervisory control is the degree to which the supervisor is authoritarian as opposed to being deferential toward the computer. At the extreme an authoritarian human supervisor always gives the




Figure 35. A virtual pocket calculator generated on a CRT and operable through a touch-panel overlay

orders to the computer, never takes a decision or cooperates with requests from it. In this case presumably the computer "reports back" for more orders either when it is finished with its assignment or it is in trouble and cannot finish. At the extreme of deference the supervisor in effect says to the computer "You do what you think is best and tell me what you want me to do".

Somewhat orthogonal to the authority-deference dimension is whether the operator's relation to the computer is one of "trading" (back and forth in alternating series) or "sharing" (operating in parallel). In trading the supervisor says to the computer, in effect, "you do what you can do, hand it back to me and I'll do what I can do, then I'll hand it back to you, etc." In sharing the supervisor says "You do some parts of the task and I'll do other parts at the same time." Authoritarian sharing is "I'll do what's easy for me, you do the rest" while deferential sharing is "You do what's easy for you, I'll do what's left".

We are gradually understanding these problems more clearly as we develop demonstration command language (see 5.4 through 5.10).

f) Zoom Control of Degree-of-Abstraction

Such zoom control for displays, for a given situation, lets the operator view the display in terms of narrowly focused skill-based considerations and events or, at other extreme, in broadly focused knowledge-based considerations and events. In like manner a zoom control for controls would adjust the hardware and software interface (control devices and associated command language) at one extreme to allow operator issuance of specific commands or detailed assignments, or at the other extreme to give broad policy directives, criteria or goal statements.

g) Prompts, Edits, Threats and Defaults

A sophisticated and cooperative computer will prompt its user, i.e. give reminders and suggestions as to what commands are appropriate next. But he may ignore or in some cases suppress this help if he is confident that he doesn't need it.

If the operator's command entries make some sense to the computer it can try to give him some editing feedback on the spelling of the proper alternative commands or the syntax, etc.

If the operator is too reluctant or recalcitrant the computer might indicate that it can only sit and wait for so long, or if clarification is not forthcoming it will return to an early point in a program, etc.

Finally, the computer must have default values and subroutines which it uses when "better" (syntactically proper, sufficiently recent, etc.) entries are not made by the operator. Thresholds or criteria for using the defaults are then required.

h) Too Many Response Options

As with displays, there is the danger of confusing the operator and degrading system performance by having too many response options. If there is an easily remembered structure for response options, and if this structure is provided in hardcopy (labels, placard, etc.) on the control panel, a larger number of options are tolerable. The worst case is when a response option, appropriate to a situation which occurs rarely but then is very important, has a coding superficially similar to that of other commands, and does not fit into a neat and easily remembered structure, and has no available hard-copy explanation. In this case erroneous response is probable.

5.3 Special Teleoperator Control Problems

a) Absolute and Relative Coordinate Frames

It is convenient to program a manipulator arm in a coordinate frame which is fixed to the base of the manipulator. It would be nice if the positions of all objects in the environment were known in this same coordinate frame.

Unfortunately the base of the manipulator, for example when attached to a moving vehicle, may move relative to the environment. Also, and this is probably the more important problem for supervisory control, either or both initial positions and final (goal) positions of objects may change for a given subroutine. In both cases it is necessary to map one coordinate frame into another, i.e. specify any point, originally given in one set of coordinates, by values of another set of coordinates.

The latter case occurs where one wishes a programmed manipulator to "do the same thing" to another object in a new position, and/or to end up in a new location. Reprogramming the whole sequence of operations is not necessary. However the computer must be told (initialized) with the three position and three orientation vectors of the start and/or end locus of the new object relative to the base reference frame (Figure 36). This can be done in a number of ways, but an easy method is for the operator to do it with master-slave position control.

Getting or returning a tool or other object to the same location in earth coordinates, even though the manipulator base has moved relative to earth, is an example of the same thing.

Experiments in which a human supervisor initialized computer controlled manipulation subroutines are reported in 5.4 - 5.6.

b) Relative Motion Compensation

If an object to be manipulated is initially moving relative to the base coordinate frame, it is difficult for the human operator to initialize or to perform direct manual control on this "moving target" by joystick or master-slave control. If the object's motion is known or can be measured it is possible to have the computer make the manipulator move automatically to compensate for this relative motion. Then any joystick or master control inputs are added to this, as though the object were fixed. In 5.7 we report experiments which demonstrate a technique for effecting this computer aid.

c) Manipulation Accommodation

When a person puts a peg in a hole he "accommodates" in two ways. First, as he presses it into the hole it tends to move so as to align itself with the hole. If his grip on the peg is not perfectly aligned, the peg will tend to exert a force on his hand in a direction such as to move the hand and relieve the force. If the hand is relaxed, i.e. the hand or skin is elastic, it will be driven passively in that direction even if there is no sensation of the force and no active muscle involvement. This is "passive accommodation". If there is sensation and the person actively moves his hand so as to null out the force, this is "active accommodation". Theoretically the latter could occur even with a perfectly rigid peg, hole and hand. In reality, accommodation occurs in skilled neuromuscular behavior




Figure 36. Initializing new object reference frame relative to base reference frame

by a combination of passive and active response. This is true not only for putting a peg in a hole, but also for opening a door (accommodating to its circular arc without pulling it off its hinges) and other common manipulation tasks.

The need is to provide aiding for a teleoperator to do the same thing. Passive aiding can be achieved by making gripping surfaces and limbs sufficiently elastic, but this causes other problems, including positioning inaccuracies, vibration, etc. Active accommodation requires computation both for measurement and control.

C.S. Draper Laboratory (1977) has developed a unique passive accommodation device based on "remote center compliance", and shown how added strain gauges and active computer accommodation control can make manipulation performance even better.

d) Bandwidth, Time-Delay and Visco-Inertial Lag

We noted in 4.3a how bandwidth constraints of a communication channel between operator and teleoperator force degradation of a display. The degradation of resolution, grayscale and frame-rate of the display may produce degradation in decisions about what direct real-time control movements to command or what automatic sub-routines to call up.

But control is degraded in additional ways. A low bit-rate means that after a command is sent by the operator to the HIS (longer for a more complex command) it may take some time to get the whole message to the TIS. If there is a transmission time delay due to the communications medium itself (one second per 5000 feet for sound transmission in water, 1.5 seconds for electromagnetic transmission to the moon) the TIS will receive a control signal at least one time delay after the information is fresh, actually longer due to the low bit-rate added to the delay for completion of a message. In continuous direct control such a time delay is present going both ways, generally causing instability and forcing the operator into a very tedious "move-and-wait" mode of operation (Ferrell, 1965). Supervisory control relieves this problem, in the sense that only the supervisory commands are delayed, but these are not so dependent on immediate feedback, and the TIS control loop stability is not dependent on the supervisory commands. The tightly coupled control loops, those from which stability is determined, are closed locally to the TIS i.e. with no time delay.

Use of a predictor display, as described in 4.8, may be useful to the supervisor whether human control is supervisory or direct. Also, if there is some visco-inertial lag in the vehicle or manipulator being controlled, there may be some deleterious effect of the time-delay.

5.4 Implementation of a Supervisory Controlled Telemanipulation System

A supervisory control system for a manipulator called SUPERMAN was implemented in the laboratory by Brooks (1979). It used our Argonne E2 force reflecting master-slave manipulator. Both master and slave arms have all six degrees-of-freedom plus grip and are driven by AC 60-cycle servomotors. All fourteen motors have both potentiometers for position feedback and tachometers for rate feedback. The system was modified to accept commands from and to return position feedback to an Interdate Model 70 computer through an A/D converter. Details of the equipment are found in Brooks' report.

The aim was to provide a system which could be taught or programmed by any mix of analogic commands from the human operator (master positioning and force signals through a seven degree-of-freedom master, or a conventional spring loaded joystick) and symbolic commands (conventional alphanumeric keyboard, special push-button console), with various degrees of computer control.

For example some types of computer control which were "built in" to the system were:

- 1) terminal point control: operator or computer specifies final position of end effector, and computer determines trajectory;
- 2) path control: operator or computer can specify path constraints which determine trajectory;
- 3) resolved vector control: operator or computer specifies an end effector displacement along a coordinate fixed at the end-effector, and the computer determines joint angular displacements to achieve it;
- 4) resolved rate control: operator or computer specifies rate of end effector movement along a coordinate fixed at the end effector, and the computer determines the joint angular velocities required to achieve it;

- 5) force control: operator or computer specifies end effector force vector, and the computer determines required joint torques.

Figure 37 diagrams the SUPERMAN system. Figure 38 shows the hierarchy of computer modes and corresponding instruction codes. In STANDBY the computer may branch to EXECUTE to receive on-line action commands from the operator to execute what is in the task file, may give in to a TAKEOVER by the operator, may STOP, or may receive other specifications of how to control. In DEFINE the operator may enter commands through the specialized console DASI, or may go to EDIT to modify a previously defined command string. In some cases pseudo natural-language is used for communication: "IF FORCE > XXX, INCREMENT DOF XXX"; in other cases it is highly abstract. Brooks' report describes the command language in some detail and gives a step-by-step example of teaching the manipulator to take a nut off a bolt and put it in a box. The steps are shown in Figure 39 without explanation.

5.5 Experiments to Compare Supervisory Control of Telemanipulation with Direct Manual Control

Brooks conducted extensive laboratory experiments in remote manipulation with the SUPERMAN system to evaluate supervisory control in comparison to: (1) fixed rate control (positive-off-negative by means of a separate switch for each degree of freedom); (2) variable rate control (continuous adjustment for each degree-of-freedom by means of a joystick); (3) master-slave position control with no force feedback; and (4) master-slave position control with full force feedback. He used six tasks: (1) retrieving a tool from a rack; (2) returning a tool to a rack; (3) removing a large nut from its bolt; (4) picking up blocks and putting them into a bucket; (5) opening or closing a valve; (6) digging sand and filling a bucket. In each case the subject was instructed to perform the task as fast as possible without making errors. Both average time to complete the task and errors were recorded. The above two-way combinations of treatments were replicated for both mono and two-view video arrangements.

Three subjects were used for the first four tasks and only one for the last two tasks. Each three-way combination was repeated five times to obtain a mean and standard deviation. Unfortunately the subjects were not always the same and had differing amounts of training; but the relatively consistent results showed that this factor of variability made little difference in relative performance between control modes, and this was true across tasks.




Figure 37. Diagram of Brooks' SUPERMAN supervisory control manipulator system




Figure 38. Hierarchy of computer modes and codes for SUPERMAN




Figure 39. Example of Brooks' supervisory commands to take a nut off a bolt




Figure 39 (continued). Example of Brooks' supervisory commands to take a nut off a bolt.

Figure 40 illustrates the comparison of manual control in each of the four modes (with both one and two video views) with each other and with supervisory control. Note that in this experiment supervisory control consists of two phases. The first is an initialization phase where the operator gives the master manipulator arm its initial position conditions and operates the correct push buttons on the specialized console. In the second phase the computer executes the operation independently of the operator. Note that in the digging task the initialization phase is much longer than in the nut removal task, and hence in some cases master-slave control is better.

Table 4 gives the data for all tasks and subjects separately, the numbers being ratios of completion time for manual control in the given mode to completion time for supervisory control. Note that the subjects, even though of varying levels of training, are not so different, and that there is hardly any difference between one and two video views. It is evident that master - slave was sometimes faster than supervisory control, but neither joystick nor switch control was ever faster.

What is remarkable about these results is how good even this first implementation of supervisory control in telemanipulation proved to be for these ideal conditions for manual control, i.e. no time delay, no intermittent display, good video feedback in other respects as well, no requirement that the operator time-share his attention with other tasks. Were any of these factors significantly less than ideal it is fairly clear that supervisory control would show much greater advantage over direct manual control.

Brooks also noted that subjects were considerably more fatigued after work in the manual control mode than in supervisory control.

5.6 Refinement of Command Language for Supervisory Telemanipulation

Yoerger (1982) has extended the work of Brooks by developing a much more robust and general-purpose command language and associated software structure. This work has proceeded at two locations simultaneously. The first is at the MIT Man-Machine Systems Laboratory, using the same Argonne E2 manipulator that Brooks used. (However in this case the manipulator has been interfaced to a PDP 11/34 computer and an AN 5400 A/D converter by Tani (1980)). The other is the EAVE-WEST submersible






Figure 40. Brooks' experimental results comparing direct manual control with supervisory control (sum of manual "location" initializing plus "computer" times). MS means master-slave, FFB means force-feedback, JVRC means joystick-variable rate-control, SVRC means switch-variable rate-control. Data are shown for three different tasks and for one and two video views.

	1-VIEW				2-VIEW				
	MS	MS NO FFB	JVRC	SVRC	MS	MS NO FFB	JVRC	SVRC	
VALVE 1-DOF	1.1	1.3	1.2	2.0	0.8	0.9	1.1	2.0	E1
NUT-OFF 2-DOF	1.0	1.2	1.6	2.0	0.8	1.1	1.2	2.1	T1
	0.9	1.2	1.4	1.5	0.8	0.9	1.1	1.9	E1
	1.5	1.9	*	2.1	1.1	1.5	*	2.0	U2
SAMPLER 3-DOF	0.7	0.7	1.2	1.8	0.8	0.8	1.3	1.9	T2
	0.7	0.7	1.2	1.6	0.7	0.7	1.3	1.7	T3
	0.6	0.6	1.2	1.7	0.7	0.7	1.3	1.8	E1
SCOOPER 7-DOF	0.5	0.6	1.9	4.8	0.5	0.6	1.9	4.9	E1
RETURN-TOOL 6-DOF	0.9	1.0	3.2	9.4	0.9	1.0	3.0	10.6	T4
	1.1	1.5	2.8	12.8	0.9	1.1	2.3	11.3	E1
	1.9	2.6	*	16.9	1.7	2.4	*	16.3	U1
GET-TOOL 6-DOF	0.8	0.9	3.2	10.8	1.0	1.2	3.2	11.7	T4
	1.0	1.3	2.3	12.5	0.9	1.1	2.4	12.4	E1
	1.8	2.6	*	15.3	1.4	1.7	*	16.0	U1

Table 4. Brooks' experimental results: ratios of time for manual control to time for supervisory control

vehicle in Heckman's laboratory at the Naval Ocean Systems Center, San Diego. The latter has a five-degree-of-freedom arm-hand (plus grip) and is driven by DC torque motors through harmonic drives, with position feedback.

Yoerger's system draws heavily on and combines the contributions of others:

- movement descriptions which are 4×4 transforms (Uicker, Pieper, Brooks)
- both relative and absolute movements (Brooks)
- both analogic and symbolic input (Verplank, Brooks, Gossard)
- task-oriented supervisory command language (Chu, Crooks, Freedy)
- structural programming constructs (Dykstra)
- binary sensing routines for manipulator-environment interaction (Grossman)
- use of touch sensor (Flyer, Wood, others)
- FORTH computer language extensibility (Moore)
- supervisory control (Sheridan)

The primary design goal is that this system is to be used by people, unlike other systems based primarily on compatibility with machine vision devices, CAD/CAM data bases, etc. It uses an English-like command syntax. It has an interactive and non-mathematical symbolic interface to the Cartesian mathematics. While the interface is less general than in some other systems, it is easier to understand. It is oriented toward computer-literate people, but little knowledge of manipulator computation is required.

The use of analogic, symbolic relative and absolute movements may be mixed freely. There are three types of analogically defined movements:

- 1) POSITION. It is defined by moving the arm there. When a position command is encountered the arm is servoed there, using an internal 4×4 transform.
- 2) MOVEMENT. Both initial and final positions are given to indicate a relative movement to the computer. Subsequently the computer executes the movement with the most direct trajectory relative to whatever its current position is.
- 3) PATH. This is both defined and executed as a series of positions.

Movements may also be defined symbolically, by use of the keyboard, where distances can be specified. Commands exist for relative translations, such as UP, RIGHT, FORWARD, etc. Commands for relative rotation are PITCH, YAW and ROLL. As the NOSC arm has only five degrees-of-freedom, the third rotation cannot be specified. OPEN and CLOSE commands refer to the manipulator grippers.

A "task level" controller allows the operator to function at a still higher level. The simplest form of task level control consists of grouping movement commands to make a procedure, for example

: TURN CLOSE 90 ROLL OPEN -90 ROLL:

defines a procedure named TURN which closes the grippers, turns the hand 90 degrees, opens the gripper, turns the hand -90 degrees. Such a procedure can be executed 10 times by

: OPEN - VALVE 10 TIMES DO TURN LOOP:

Another example is the definition of SEARCH to move the arm forward in 1 cm. increments until the touch sensor reports that contact has been made.

Another principal contribution of Yoerger is the use of the manipulator with touch sensor to actively model the environment and use that model to execute a task. The needs for undersea operations are seen as different from those developed at artificial intelligence laboratories where movements of the manipulator are implied by stating a desired relationship between two relatively known objects. In the undersea situation the environment is essentially unknown, and one must start with first-order assumptions such as that a surface has been contacted which is either flat or curved. One might like to determine precisely what and where it is so as to perform a scanning or other operation on it.

For example, if the operator decides he is in the vicinity of a planar wall he types PLANE WALL and the manipulator searches, i.e. touches the surface at three points and computes the coordinates of the plane so defined. The operator can adjust the size and shape of the triangle used to define the plane. The operator can later return to the nearest point on that plane by the command WALL MOVE, or perform a patterned raster or other scan across that plane by the command WALL SCAN.

If the operator decides a curved surface would be a better model for present purposes he may command SURFACE CURVED, which identifies the surface as a series of planar surfaces, the coarseness of which may be adjusted by the operator. CURVED MOVE, CURVED TRAVERSE, INSPECT and PLACE SENSORS will perform the expected operations relative to this curved surface.

Experiments with human subjects using this command system on various manipulation tasks are in progress at this writing. Task completion time and quality of performance are being studied as a function of control mode, syntax, terms, etc. as well as constraints on the video bandwidth.

5.7 Experiments with Relative Motion Compensation in Telemanipulation

One problem encountered in remote manipulation of objects undersea is that the manipulator base may move relative to the object being manipulated, and this makes either direct manual control or supervisory control difficult. This relative motion occurs either because a manipulator is being supported by a vehicle which is hard to hold steady against ocean currents or other disturbances, or because the object being manipulated is being so buffeted, or both.

A means to overcome this is to make some measurement of the relative changes in displacement and orientation between manipulator base and object, either by optical, sonic or mechanical means, then to compensate for these changes by added motion of the end effector. In particular, the use of a mechanical "measurement arm" is discussed by Brooks.

Hirabayashi (1981) implemented such a scheme experimentally. He constructed a six degree-of-freedom (all angular movement) measurement arm which was lightweight and flaccid (offered little restraint). A six-degree-of-freedom Jacobian matrix transformation then allowed determination of the relative displacement of any object to which the measurement arm was attached.

Using a task-board with holes into which pegs were to be inserted, Hirabayashi drove the task board with a continuous random positioning device (three degrees of freedom, roughly 0.2 hz bandwidth, 6 inches root-mean-square amplitude.) He

then attached the measurement arm to this task board, and used the resulting measurement of displacement to produce a compensatory displacement bias between the master and slave.

When the arm was under computer control it compensated to within 0.2 inches, even with a crude three foot long measurement arm. When the arm was used in direct master-slave control it was found to be much easier with the compensation than without it to put pegs into the holes in the moving task board.

5.8 Simulation Experiments with Supervisory Control of Undersea Vehicles

At the time of this writing various simulation experiments are in progress with supervisory control of vehicles, but little experimental data are available as yet.

We have built a tethered three-wheeled laboratory vehicle shown in Figure 41 to simulate for an operator the real-time control of an undersea vehicle. Its wheels turn and drive in synchrony. A video camera is mounted on a tiltable arm which moves up and down on a column, which in turn can rotate on the wheeled base. All five degrees-of-freedom are driven by stepping motors and serve to displace the video camera in five degrees-of-freedom, the missing degree-of-freedom being rotation of the camera around its own axis. The operator sits at a console and sees on a monitor what the camera sees as it maneuvers relative to a scale model of an environment. The operator gives supervisory commands through a standard keyboard to an 11/34 computer which in turn gives rate signals over a wire to an on-board Z-80 computer, which in turn outputs pulses which drive the stepping motor translators.

Preliminary tests by Messner (1982) compared supervisory position control (where the video camera is commanded to move approximately to a certain location and orientation in five degrees of freedom) to direct manual (switch) rate control. These showed supervisory position control to be much easier to use for translation and rate control easier to use for orientation of the vehicle.

The computer-graphic display of a vehicle with manipulator arm attached (shown in Figure 30) has been matched by Kazerooni (1982) to a real-time hydrodynamic model of a vehicle, its thrusters, its sonar and its control system. The parameters of this model have been adjusted to match a variety of real vehicles including Woods Hole's Alvin and Hydroproducts' RCV-150. Then supervisory control schemes can be evaluated by being "test flown" on this simulator in real-time by a human operator.




Figure 41. Wheeled laboratory vehicle drives video camera in five degrees-of-freedom relative to a scale model of an environment. Operator, viewing the video scene, commands a computer simulation of (any) vehicle, which in turn commands the camera motion.

Such a computer model can also be inserted between the human operator and the five-degree-of-freedom video camera-mover to create a realistic simulator similar to those used for aircraft.

5.9 Emergency Control and Ultimate Authority

Emergencies in large, complex, capital-intensive, risk-producing man-machine systems pose very special problems of authority and style in man-computer cooperation.

One philosophy is that when the emergency is sufficiently critical and if that criticality can be measured, an automatic system should seize control from the human operator and take immediate action. In a supervisory system, where it may take some time for the operator to figure out what is going wrong and what to do about it, and since an operator under stress is likely to be less reliable than otherwise, such automatic take-over is a good idea. This is the philosophy of the "safety system" required by law in all western nations' commercial nuclear reactors to drop in control rods and perform other emergency control functions.

Another philosophy is that if the state of the system is sufficiently evident and if the operator is normally in the control loop or capable of getting back in the loop quickly, as in piloting a fighter aircraft, let him trigger the automatic system. The ejection seat is much a system.

Should the operator be able to take control back from the automatic emergency control system in case he sees that it is not doing the correct thing? Again the prevailing attitudes are colored by circumstance. In the nuclear power plant the attitude is pretty much "no", since automatic control events occur too quickly for him to comprehend. Exceptions are made for corrections to longer time-constant events which unfold more slowly following the first automatic response. For the aircraft pilot, the feeling is "yes" exemplified by the fact that while making an automatic landing the pilot can immediately recover control either by pushing a yellow button or by jerking the control column.

Thus, in the supervisory control systems of the type we are concerned with the trend seems to be for the supervisor to defer to the automatic system, much as a good manager would defer to a skilled technician, stepping in when both (1) there is ample evidence that the technician has gone awry and (2) there is time to find another approach. Ultimately the supervisor takes the position toward the computer "You may complicate life for me, hurry me, take over control from me, and even make me look silly, but ultimately I can pull your plug!"

Operators should be trained in how to deal with the automation during emergencies, and the simulator is the best means of doing so. Simulator exercises for this purpose should not be of expected, standard scenarios, but situations the operators have never seen before. Operators should be taught to allow the automatic systems to operate until certain conditions obtain, based on knowledge of how they work. He should also be provided with tools and procedures and criteria for orderly take-over and, if necessary, implementation of "fail-soft" abort.

5.10 Telling a Computer How to Decide

The ultimate human response problem in supervisory control is not the planning or programming of the detailed task procedures, the adjustments or trimming during monitoring, or the more substantive interventions in case of emergency. As we understand better how to program computers to make them smarter they surely will be able to recognize patterns and make quicker and more reliable on-the-spot control decisions than people can. Inescapably the most difficult problem will be to endow the computers with the most profound aspect of human intelligence - our values, mores, or utilities, i.e. our bases for deciding what is good and what is bad and how to weigh the relative worths of different events based upon what is knowable and measurable.

Explicating "values" is not a new problem; it is ancient one. But the potential for intelligent computers to be our servants and colleagues in system control tasks forces a return to this fundamental problem. We include under the rubric "values" the concepts of utility or relative worth, probability or expectation, and relation (causation, correlation and membership).

Methods of elicitation and analysis of subjective judgements of such variables has long been the business of psychophysics, but the computerization of both the elicitation and the application of such data is relatively new. Yntema, Torgerson (1961) and Yntema and Klem (1965) have proposed useful techniques for combining single-attribute worth judgements in multi-attribute predictions of worth. Keeney (1962), making some additional assumptions to make the mathematics rigorous, has extended these ideas to "multi-attribute utility theory". Sicherman (1975) developed a practical man-computer program for evolving a person's multi-attribute utility function.

Utility methods require quantitative attributes to be understood by the judges from the outset. Multi-dimensional scaling (Shepard, Romney and Nerlove, 1972), starting with "dissimilarity" judgements for all pairings of a set of objects or events, determines that set of attributes which best discriminate the data. Computer programs are available which map the set of objects or events in a space of up to four or five dimensions.

A new approach for telling a computer how to decide is the theory of "fuzzy sets". Each input or output variable (from or to the environment) is characterized by a membership function which specifies its degree of association with a linguistic term or "fuzzy variable". For example, in driving, when the car ahead is 50 ft. away it may be considered 0.7 "close", when 20 ft. away it may be considered 0.95 "close". Then a set of statements is elicited from a judge about a task. For example, when the car in front is "close" and "stops quickly", then his judgement is "brake hard"; when he is farther way" and "stops quickly or gently", then he judges "brake gently". Fuzzy logic algorithms then are applied to such statements and generate a truth table. This provides a quantitative input-output "action matrix" which the computer then can use for executing control automatically. A human operator can continue to "tune" the computer's judgement, especially for rare events (Buharali, 1982).

6. MATCHING CAPABILITIES TO TASKS AND EVALUATING PERFORMANCE

6.1 Tasks and Tools

Supervisory control is emerging because it is a better way of doing certain tasks. As a means to combine people, computers, and artificial sensors and effectors, one is interested to compare it with other (non-supervisory) ways of doing the same tasks. These might include: human only (with essentially no mechanical aids); human augmented by artificial sensors, effectors, displays or controls (but no computer); further augmentation by significant computation but with no control loops closed through the computer; and automation only (with no human control), as in Figure 1.

The planning, intervention and learning modes of supervisory control can be implemented in any of the first three ways, i.e. a person acting with or without help from artificial devices. During monitoring the human operator can turn his attention away so that system (5) in Figure 1 (complete automatic control) exists momentarily. Thus in considering "how much" supervisory control to employ it is worth asking (1) what type and what extent of sensory or cognitive augmentation (S,C or R) is appropriate, (2) for what length of time, (3) in which of the five supervisory functions this aid would be useful.

The various sensory-computer-motor aids or augmentations can be thought of as "tools" given to the human operator by which he accomplishes the tasks he is assigned. But the tool must first be matched to the user, just as a carpenter's hand tool must be matched to the size and strength of his hand, before he can think of its effectiveness in a task (really the effectiveness of carpenter-plus-tool). This is represented in Figure 42. The operator's capabilities (strength, speed, discrimination, etc.) must be matched to the requirements (force, accuracy, etc.) of the tool or instrument, while the operator's requirements (for sensory feedback, avoidance of over-loading) must be matched to the way the tool behaves. Once this match has been made one can consider matching the capabilities of the operator-tool combination (e.g. in information sensing or response) to task requirements, and the capabilities of the task to provide information and disturbance to the requirements of the man and tool.




Figure 42. Matching capabilities (C) and requirements (R) of operator to supervisory tool or instrument, and in turn matching these in combination to the task

Another useful metaphor of the supervisor-computer (plus sensor-effector) relationship is that of the human staff-aid or consultant. The employer and employee each want to know what the capabilities and requirements of the other are so they can work together. Since now we assume intelligence on the part of the instrument we might add goals and intentions to the list. In fact, as previously noted, we carry in our heads mental models of and norms for our employees or employers. Sophisticated, cooperating employees and employers may, on this basis, be likely to communicate to each other their requirements and capabilities.

Generally as tasks become more complex (informational or thermodynamic entropy can be used to characterize complexity) the degree of automation decreases. This is because we just don't know how to explicate such tasks: we do them intuitively or artfully. Figure 43 suggests this relationship. The capability of any supervisory control system may be plotted qualitatively on these coordinates, technological perfection being the upper right-hand corner.

6.2 Criteria for Matching Computer Augmentations to Supervisory Operator

Ten criteria are listed in Table 5, with ordinal "applicability ratings" estimated for each supervisory function. Speed of use is very important during emergency intervention, important during teaching and monitoring which are routine, less important during planning and learning which are off-line. The next four items, while qualitatively similar to each other, have differing degrees of importance for the different functions. Transparency of system state is important primarily with respect to system operation. Display integration is inherent in planning and learning from experience. Compatability with internal modeling must be strong to have planning and learning take place, while in other functions it is not the most essential. Naturalness is most important in teaching and intervention, when control-display compatability is critical. Mental workload and stress are obviously most relevant during functions of operation, and especially emergency interventions. Mental workload is discussed further in 6.5. All supervisory functions must be somewhat precise and accurate as they affect the others. Learning from experience obviously requires the most memory to record experience, while teaching takes the next most to allow command flexibility. Intervention poses the greatest risk since by definition it happens when the system is at risk. Monitoring probably deserves a bit more attention with respect to economic cost since that is where the time is spent and failures are detected.

AUTOMATION vs. ENTROPY




Figure 43. Qualitative relationship between task entropy (complexity) and degree of automation. The diagonal from upper left to lower right is a "frontier" of development.

Table 5. Applicability of various criteria to evaluating the augmentation/aiding of supervising functions. (0 is least applicability, 3 is greatest).

CRITERION	SUPERVISORY FUNCTION					Σ
	PLAN	TEACH	MONITOR	INTERVENE	LEARN	
1) speed of use	1	2	2	3	1	9
2) transparency, ability to see through to system state	0	0	3	3	1	7
3) integration of variables into meaningful picture	3	1	2	2	3	11
4) compatability with internal modeling	3	2	2	2	3	12
5) naturalness of operation, display-control compatability	1	3	2	3	1	10
6) mental workload, stress	1	2	2	3	1	9
7) precision or accuracy of operation	2	2	2	1	1	8
8) computer memory allocation	2	1	1	1	3	8
9) reliability, risk	1	2	2	3	1	9
10) economic cost	1	1	2	1	1	6
TOTAL Σ	15	16	20	22	16	89

Other criteria apply and are likely to be used in due course. This table is intended primarily as an example of how criteria can be developed.

6.3 Some Methods for Analyzing Supervisory Control

The number of analytical methods which can be applied to supervisory control is surely a function of the number of interested analysts. We will discuss only two candidates.

a) Time-Line Allocation Ratios for Activities


Figure 44 shows a time-line for various activities the supervisor might engage in. Three bands for S(sensing), C(cognition and R(responding) indicate broad categories of supervisor behavior. Numbers on blocks of time identify particular subcategories (e.g. observing a display of a particular system variable, modifying or initializing a particular program). D indicates time duration, U indicates time between updates. We could have a separate time-line for different computer activities (these can be simultaneous with human supervisory activities) but for simplicity we show only the supervisor. We also assume that any time not observing displays or modifying programs (both of which the computer can measure) is "thinking".

Various sums or ratios of these times can be determined to serve as indicators of efficiency, for example:

- the total fraction of time spent engaging in a particular activity
- the total time spent observing variables related to a particular response divided by the time spent programming that response
- the time spent engaging in a particular activity divided by the time between updates for that activity

b) State-Dependent Transition Probabilities for Activities

Figure 5 showed a matrix of activity categories among which the supervisor may allocate his attention. Dimensions were: tasks; the five supervisory functions; knowledge, rule and skill-based behavior; SCR categories and activities within each of S, C and R. For the HIS and TIS computers there were similar but reduced matrices, lacking the supervisory function breakdown.



- S sensing activity
- C cognitive activity
- R responding activity
- DS duration of sensing activity
- DR duration of responding activity
- US update interval for sensing activity
- UR update interval for responding activity

Figure 44. Hypothetical time-line for different supervisory activities

The state-dependent transition probability is the probability that, given that attention is allocated to one cell in this matrix, the next activity will be some other cell. This may be represented by a two-dimensional matrix, where rows are all possible "present" activities and columns are all possible "next" activities. The cells allow for all possible transitions from any present to any future activity. The diagonal cells have no meaning in this case.

To measure such transitions "within" the human supervisor would be difficult, since it would require some observation or verbal protocol technique. It would be easier to measure activity transitions within the HIS, since the HIS itself could do the bookkeeping. This would be an indication of how the supervisor was interacting with the HIS, but it would not tell the full story of human fine-grain attention switches, say between S,C and R categories. Possibly the HIS automated bookkeeping could be combined with verbal protocol or real-time category judgements by a human observer to infer transitions in supervisory activity.

When transition is defined as a change to a different activity, i.e. a different cell in the matrix, all record of time is lost. An alternative procedure is to define a transition by clock intervals, and allow "transitions" to the same activity (diagonal cells) as well as to different activities. Also, by adding frequencies or probabilities over any row or column one can determine fractions of total time spent in that activity.

6.4 Supervisory Errors and System Reliability

Human errors may be classified in many ways. They may be classified according to whether they are errors of omission or commission or out-of-sequence or too-late, or according to whether they are associated with sensing, motor or cognitive function, or within cognitive function whether memory or deduction. A recent distinction made by Norman (1981) is between "mistakes" (errors of intention as to what should happen) and "slips" (errors in executing the intention).

All of these errors occur in supervisory control. In any of the five supervisory functions the wrong mental model can lead to incorrect prediction, a misdetermined intention and eventual error. This is a "mistake". In programming and monitoring, procedures may occur, for example, where a relatively infrequent sequence of steps includes part of another more familiar sequence. Here the operator can easily "get off on the wrong track". This is a "slip".

Syntax errors can occur in programming, such as inappropriate inference in monitoring and deciding when to stop taking data and start taking corrective action. Stress can certainly produce errors for any of the functions, but particularly during intervention. In the learning function the wrong inferences can be made from collected data.

There is a serious definition problem with human error, in that both the concept of "error" and the elaborate mathematical apparatus called "reliability theory" that goes with it assume a binary criterion of human performance. That is, the assumption is made that behavior is either satisfactory or it is not. Electronic and mechanical devices often fail in fairly discrete ways, but sometimes they simply become "off calibration" or "wear" so that their performance is not neatly assignable to "normal" or "abnormal" categories. With human behavior the situation is likely to be more difficult to assign on a binary basis. In performing tasks people take more time or less, are more accurate or less, follow procedures more or less. It is rare that they simply fall apart or stop functioning.

Errors of omission or commission in throwing switches are straightforward enough. But errors in observing or thinking surely are not, and we do not have good means for measuring these outside the laboratory (through we try our best with verbal protocols and related techniques). Further, and unlike machines, people somehow become aware that they have made errors and often are able to correct them before their effects are felt by the system.

Both the behaviorist and the feedback control purist might contend that people err only because they have not received adequate feedback from their environment, i.e. the rewards and punishments are not sufficiently frequent or strong (in control language the feedback loop has too large a time delay or too small a gain). Thus, in the chain of events consisting of thinking the right thoughts, moving the head and body in the right directions, observing and interpreting the right displays, grasping and correctly actuating the right controls, and correctly confirming that the system responded in the appropriate way, there are many feedback signals. If sufficient feedback signals are not present behavior tends to go awry and error occurs.

Currently the most accepted model of human reliability is THERP (technique for human error rate prediction) developed by Swain (1980) and his colleagues. Presently it is being applied to commercial nuclear power plants. Its elements are:

- 1) baseline human error probabilities (HEP's) tabulated for a set of behavioral elements, such as failure to follow written procedure, errors of omission in otherwise proper procedure for using check-off lists, failure to note correct status of an indicator lamp, selection of wrong control switch, etc.;
- 2) multiplication of these basic HEP's by "performance shaping (correction) factors" related to situational elements such as expectancies, stress, training, environmental conditions, etc.;
- 3) modification of the resulting HEP's according to dependence upon prior tasks, where no modification is appropriate if there is no dependence, and probability of failure, given failure on the prior task, approaches one if there is high dependence;
- 4) further modification according to the likelihood of no recovery from an error before that error has effect;
- 5) multiplication of task sequential net probabilities of correct action, i.e.;

$$\Pi_i (1 - \text{net HEP}_i)$$

to determine reliability of "chunks" of human behavior which are separable from chunks of machine behavior;

- 6) integration of these numbers into full system reliability analysis using fault trees, event trees, etc.

In supervisory control, the operator is to a large extent the programmer and goal-giver for a subordinate intelligent system; he is not following a set procedure given him from above. Therefore it may be asserted that he is more accountable for the success or failure of his system. On the other hand, his own errors in behavior are less easily defineable and measurable than in simple manual tasks; his interaction with the computer and its sensors and effectors is close and collegial. He may more easily blame design, maintenance or management, and outside investigators will find it difficult to identify error causality to be otherwise.

6.5 Mental Workload in Supervisory Control

Mental workload has become a very important topic recently, and the motivation for this is closely associated with the trend toward supervisory control. Before the industrial revolution physical workload, not mental, was the prime concern. Then, as the industrial automation became more widespread, emphasis switched to mental

workload, usually concerned with how much production activity could be sustained continuously, as measured by quantity and quality of product coming off the line. As the human operator increasingly becomes a supervisor, mental work load can vary from very passive routine monitoring to very urgent and stressful activity to diagnose and recover from failure. Such events may occur as a sudden transient from a passive monitoring state. This makes mental workload greater, since the operator may not have kept up with recent changes in state variables, and may have to go through additional steps to access required computer-based information.

Mental workload is a construct like "intelligence" in that it cannot be observed directly; it must be inferred. Yet there is almost universal consensus that mental workload is experienced, and that somehow it can predict when performance is likely to deteriorate, before a measure of performance itself will indicate deterioration.

Some would define and measure mental workload in terms of the task to be done within a given time. The more extensive the task the greater the workload. This is how aircraft manufacturers define and measure mental workload. Most researchers in the field reject this approach, claiming that by this measure it would not matter whether a robot or a person did the task; the mental workload would be the same, and would not be particularly "mental". Assuming that mental workload is a result of the task, not the same as the task, there are three types of definition/measure. The first is the so-called secondary task, an additional task which the subject operator is asked to perform when he has spare time. This is usually a simple cognitive verbal or skill task. By definition the better he scores on this secondary task, the more spare capacity he has from the primary task, and thus the less the mental workload of the primary task. The criticism of this technique is that it tends to interfere with the primary task; indeed a very cooperative subject may try hard to do his best on the secondary task, at the cost of significantly reduced attention to the primary task. In real systems, on the other hand, operators may refuse to cooperate at all to engage in secondary tasks.

A second broad class of techniques includes physiological measures of such phenomena as heart rate variability, galvanic skin response, pupillary diameter, spectral changes in the voice, chemical changes in the blood or urine, and changes in the on-going electroencephalogram or the "evoked response potential" of the brain - particularly the P3 (300 msec) characteristic wave.

The third technique is the subjective rating scale. One form of this is a single category scale similar to the Cooper-Harper scale now commonly used by test pilots for rating handling qualities of aircraft. Another is a three-attribute rating scale, there being some consensus that "fraction of time busy", "cognitive complexity" and "emotional stress" are quasi-independent components of mental workload, and may or may not be present in any situation (Sheridan and Simpson, 1979). This scale in modified form is now being used by the Air Force, FAA and Airbus Industrie in France. A criticism of subjective scales is that operators are sometimes overconfident of their own ability to perform and therefore underrate their own mental workload. Then too, some feel that "objective" measures are inherently better than subjective ones. Nevertheless it has been common practice to validate or calibrate other indices of mental workload against this one truly "mental" measure - the subjective judgement.

Moray (1982) reviewed the literature on subjective measurement of mental workload as part of our research program.

6.6 Simulation, Test-Retest Reliability and Validity

Because emergency or high-stress conditions do not occur very often, and because doing "fire drills" on actual systems may be too inconvenient and too costly, simulators are used to assess human response capability in such situations. The simulator can be made to record the operator's behavior, both correct and erroneous.

Full scale or "hi-fi" simulators, though they may have realism, are sometimes not economical to use. Some supervisory control investigations may be carried out with "part-task" or very much simplified simulators without changing much the basic relationships between system design parameters and human/system response. In fact, experimental control may be better with the simpler or part-task simulator. This is not to say that observation and after-the-fact reporting of errors and critical incidents in real systems should be discouraged. It is just that such reporting in real systems is likely to be biased by many factors.

In performing such an assessment it is useful to consider that there are five categories of variability which affect the final results:

- 1) decision by experimenters as to what constitutes acceptable behavior, or how good are different behaviors relative to one another;
- 2) set up of simulator, selection and instruction of operator subjects;
- 3) decision by operator subjects as to what they intend to do (which may differ from (1) due to incomplete communication or other reasons);
- 4) what operator subjects actually do (which may differ from 1 and 2);
- 5) subjective and objective reduction of data.

The sheer complexity of supervisory control systems suggests that test-retest reliability for any given measure may not be high. Reliable measures are more likely if experienced operators are used repeatedly to serve as their own controls. The number of runs is usually limited by economics.

Because of complexity, validity (whether one is making a meaningful measure and/or measuring what he thinks he is measuring) is also a difficult problem. There is always the charge that the simulator "doesn't feel like" or "doesn't pose the threat" of the "real thing". Aircraft simulators seem to have succeeded in providing sufficient realism, however. Since the supervisory operator is increasingly removed from direct observation of the task, having to observe the process through instruments, the costly "out the window" visual realism is less likely to be a factor.

There is a tendency on the part of training personnel to "standardize" all emergencies. It is useful to repeat and have the trainees anticipate some emergencies (such as stall or engine fire in aircraft, large-break loss of cooling in nuclear plants) so that they can respond quickly and "buy time". However, there is insufficient emphasis on responding to brand new types of failure, or combination failures, never seen or heard of before, not in the rule book. This is the essence of supervisory control knowledge-based behavior. Simulators allow for such training, but relatively little creative use has been made of simulators for this purpose.

6.7 How Far to Go in Automation and When to Stop

The driving force for supervisory control is the same as the driving force for automation: new technology in the form of computers and sensors and communications and robotic devices, together with associated software, makes many new forms of automation possible. The new technology does seem to carry its own imperative to be used.

We generally affirm this trend because we anticipate improved system performance, greater reliability and better economy. We also anticipate that people can be cast in supervisory and therefore higher (more knowledge-based, as compared to rule and skill-based) roles. Presumably this enhances the dignity of the human operator and causes a more aesthetically satisfying match of capabilities between man and machine. These are our hopes, our aspirations and to a large extent our expectations. We are empowering the human operator by promoting him to supervisor of a semi-intelligent, semi-autonomous machine. Or are we?

Upon thoughtful reflection we may be in for some new problems which tend to offset our gains. As the degree and sophistication of supervisory control increases the supervisory operator becomes more separated physically from the actual process he is controlling, his own (HIS) sensing and motor activities become desynchronized with the (TIS) control loop, and the coding of the signals in these two loops become more different from each other. The operator thereby may suffer a loss of empathy for the system and perhaps then lose his sense of responsibility. An operator who formerly found his dignity in being an expert at some manual or visual skill may become "deskilled". He may become a supervisor who no longer performs that direct manual skill and perhaps could not perform very well if the situation called upon him to do so. He thereby loses dignity. The supervisory operator may come to accept being mystified to some extent, giving up on truly understanding the system he controls, instead having to resort to somewhat blind faith. Thereby he may suffer a general technical insult that maybe the machine is better than he is.

For these reasons, as well as the possibility that perhaps the capability or reliability of the automation is not so great as initially advertised, it may be useful to consider how far to go on one or another scale of "degree of automation". There are a number of ways such a scale could be constituted, e.g. the amount or cost of physical equipment, the complexity or sophistication in terms of variables interacting and under control, or the speed, power or precision of control. Another type of scale is a continuum from pure passivity (available advice) at one end of the scale to total autonomy at the other. Table 6 is such a scale. It is clear that supervisory control as has been described in this report stops short of going to the limit.

Table 6. A scale of degrees of automation (from pure passivity by the computer to total autonomy).

- 1) Computer offers no assistance
- 2) Computer offers complete set of alternatives
- 3) Computer narrows selection to restricted set
- 4) Computer suggests one alternative
- 5) Computer executes that suggestion if human approves
- 6) Computer allows human to veto prior to automatic execution
- 7) Computer necessarily informs human after automatic execution
- 8) Computer informs human after automatic execution only if he asks
- 9) Computer informs human only if computer decides to do so
- 10) Computer decides what to do, does it, and does not tell human

7. CONCLUSIONS

1. There is a general trend toward supervisory control, where a human operator supervises a computer and the computer controls a dynamic process. This has great promise for remote undersea systems as well as other complex electro-mechanical control systems. The human operator interacts with the computer-based displays and controls to intermittently allocate his attention among various different tasks and to perform the supervisory functions of: planning; teaching or programming the computer; monitoring its automatic operation; intervening to adjust it or stop it or seize control; and learning from experience. The human-interactive computer may in turn communicate with multiple smaller computers which are specialized to and associated physically with various sensors, manipulators, vehicles, etc. which are remote from the human operator.
2. Modeling and experimenting with supervisory control is currently a very active field. It necessarily involves many considerations of both human and machine sensing, communicating, deciding and controlling, at different levels of knowledge, rule and skill-based behavior. Straight-forward extension of modern control theory has not proven sufficient.
3. The concept of the "internal model" of the external dynamic process seems to be useful for both human and machine components. The consistency of internal models in the heads of various operators with each other and with the actual process is a criterion for error or failure detection. Display, control-console and computer all can aid the operator in his internal modeling and decision behavior.
4. Supervisory control requires displays which become less like pilots' flight-directors and more like reference librarians or staff assistants. The computer and associated graphic display technology have permitted virtually infinite possibilities for display format, trending and predicting, dynamic simulation, information access, etc. In terms of human factors some fundamental questions about "transparency" and "perceptual overload" have become critical.

5. Supervisory control also requires a new look at man-computer language in the most general sense: (a) the proper mix of analogic and symbolic coding for giving commands (as well as for thinking about how to give commands); (b) the proper variety of choice/response alternatives; (c) the depth of hierarchy in coding and recoding ("chunking") commands; and (d) how best to tell a computer the criteria for deciding. Supervisory control has raised anew the problem of ultimate authority: when should man be able to take over from machine and vice-versa.

6. Various new supervisory control systems are being designed and evaluated in specific task contexts. Tradeoffs need to be studied with respect to general versus special-purpose designs, and between higher degrees of automation versus greater simplicity. Though obviously we want high performance, high reliability, safety and low cost, techniques for evaluation of supervisory control systems demand much further work.

REFERENCES

- Bainbridge, L. (1974), "Analysis of Verbal Protocols from a Process Control Task" in Edwards, E. and Lees, F., The Human Operator in Process Control, Taylor and Francis, London.
- Banks, W.W., Gertman, D.I. and Petersen, R.J. (1982), Human Engineering Design Considerations for Cathode Ray Tube Generated Displays, U.S. Dept. Energy NUREG/CR-2496.
- Barrett, D.S. (1981), A Computer Graphic Method for Visual Search Through Multi-Attribute Data, MIT Man-Machine Systems Laboratory Report.
- *Brooks, T L. (1979), Superman: A System for Supervisory Manipulation and the Study of Human/Computer Interactions, MIT Man-Machine Systems Laboratory Report.
- Buharali, A. (1982), Fuzzy Sets and Other Methods for Telling a Computer How to Decide, MIT Man-Machine Systems Laboratory Report.
- C.S. Draper Laboratory (1977), Exploratory Research in Industrial Assembly, Annual Report, see also preceding years.
- Chu, Y.Y., Crooks, W.H., Freedy, A. (1980), Man-Machine Communication in Remote Manipulation: Task Oriented Supervisory Command Language (TOSC), Perceptronics Report TFRT-1034-80-3
- *Deghues, B.J. (1980), Operator-Adjustable Frame-Rate, Resolution and Gray Scale Tradeoff in Fixed-Bandwidth Remote Manipulation Control, MIT Man-Machine Systems Laboratory Report.
- Ferrell, W.R. (1965), "Remote Manipulation with Transmission Delay", IEEE Trans. on Human Factors in Electronics, Vol. HFE-6, No. 1.
- Ferrell, W.R. (1973), "Command Language for Supervisory Control of Remote Manipulation in Heer, E. (ed.) Remotely Manned Systems, California Institute of Technology.
- *Fyler, D.C. (1981), Computer Graphic Representation of Remote Environment Using Position Tactile Sensors, MIT Man-Machine Systems Laboratory.
- Goodstein, L. and Rasmussen, J. (1981), "The Use of Man-Machine System Design Criteria in Computerized Control Rooms", in Aune, A.B. and Vlietstra, eds. Automation for Safety in Shipping and Offshore Petroleum Operations, North-Holland, New York.
- *Hirabayashi, H. (1981), Supervisory Manipulation for Assembling Mechanical Parts While Compensating for Relative Motion, MIT Man-Machine Systems Laboratory Report.
- Kazerooni, H. (1982), Digital Dynamic Simulation and Control of Underwater Vehicles, MIT Sea Grant Program report.
- Keeney, R.L. (1968), "Evaluating Multi-dimensional Situations using a Quasi-Separable Utility Function", IEEE Trans. on Man-Machine Systems, Vol. MMS-9, No. 2.
- Kelley, C.R. (1968). Manual and Automatic Control, Wiley, New York.

Knepp, L.A. (1981), Computer Graphics in Multi-Attribute Data-Base Searches: Implementation and Evaluation, MIT Man-Machine Systems Laboratory Report.

Kok, J. and Van Wijk, R (1978), Evaluation of Models Describing Human Operator Control of Slowly Responding Systems, Ph.D. thesis, Technical Univ. of Delft, Netherlands.

*Messner, S. (1982) Strategies for Control of a Remote Vehicle under Limited Bandwidth, MIT Man-Machine Systems Lab. Report.

Minsky, Marvin (1975), "A Framework for Representing Knowledge". In the Psychology of Computer Vision, edited by P.H. Winston, McGraw Hill, New York.

*Moray, N. (1980), Human Information Processing and Supervisory Control, MIT Man-Machine Systems Laboratory Report.

Moray, N. (1981). Personal Communication.

*Moray, N. (1982), "Subjective Mental Workload", Human Factors, Vol. 24, No. 1.

Muralidharan and Baron, S. (1980), "Demon: a Human Operator Model for Decision-Making, Monitoring and Control." J. Cybernetics and Information Science, Vol. 3, No. 1-4.

Norman, D.A. (1981), "Categorization of Action Slips", Psychological Review, 88, 1-15.

*Ranadive, V. (1979), Video Resolution, Frame Rate and Grayscale Tradeoffs under Limited Bandwidth for Undersea Teleoperation, MIT Man-Machine Systems Laboratory Report.

Rasmussen, J. (1976), "Outlines of a Hybrid Model of the Process Plant Operator", in Sheridan, T. and Johanssen, G., eds. Monitoring Behavior and Supervisory Control, Plenum, New York.

Rouse, W.B. (1980), Systems Engineering Models of Human-Machine Interaction, North-Holland, New York.

Shepard, R., Romney, A.K., and Nerlove S.B, eds. (1972), Multidimensional Scaling. Seminar Press, New York.

Sheridan, T. (1970), "On How Often the Supervisor Should Sample," IEEE Trans. Systems Man and Cybernetics, Vol. SSC-6, No. 2.

Sheridan, T.B. and Ferrell, W.R. (1974), Man Machine Systems, MIT Press, Boston.

Sheridan, T.B. (1976), "Toward a General Model of Supervisory Control", in Sheridan, T. and Johanssen, G., Monitoring Behavior and Supervisory Control, Plenum, New York.

*Sheridan, T.B. and Verplank, W.L. (1978), Human and Computer Control of Undersea Teleoperators, MIT Man-Machine System Laboratory Report.

Sheridan, T.B. and Simpson, R.W. (1979), Toward the Definition and Measurement of Mental Workload of Transport Pilots, MIT Dept. of Aeronautics and Astronautics Report.

Sheridan, T.B. (1980), "Understanding Human Error and Aiding Human Diagnostic Behavior in Nuclear Power Plants, in Rassmussen, J. and Rouse, W.B. Human Detection and Diagnosis of System Failures, Plenum Press, New York.

Sicherman, A. (1975), An Interactive Computer Program for Assessing and Using Multi-attribute Utility Functions, Tech. Rep. III, MIT Operations Research Center.

Snyder, H.L. (1980), Human Visual Performance and Plot Panel Display Image Quality, Virginia Polytechnic Institute and State Univ. Report HFL-80-1.

Swain, A.D. and Guttman, H.E. (1980), Handbook of Human Reliability Analysis with Emphasis in Nuclear Power Plant Applications, Sandia Laboratory, NUREG/CR-1278.

*Tani, K. (1980), Supervisory Control of Remote Manipulation with Compensation for Moving Target, MIT Man-Machine Systems Laboratory Report.

*Tsach, U. (1982), MIT Ph.D. thesis, in progress. See also Tsach, U., Sheridan, T.B. and Tzelgov, J. (1982), A New Method for Failure Detection and Location in Complex Dynamic Systems, Proc. 1982 American Control Conf., June 14-16, Arlington, Virginia.

Tulga, M.K. and Sheridan, T.B. (1980), "Dynamic Decisions and Work Load in Multi-task Supervisory Control", IEEE Trans. on Systems, Man and Cybernetics, Vol. SMC-10, No. 5.

Vermeulen, B.M. (1981), Computer-Aided Visualization for Design Decisions, MIT Man-Machine Systems Laboratory Report.

White, T.N. (1980), "Supervision of Dynamic Systems: Monitoring, Decision-Making and Control", Proc. 16th Annual Conf. on Manual Control, MIT, Cambridge, Massachusetts.

*Winey, C.M. (1981), Computer Simulated Visual And Tactile Feedback as an Aid to Manipulator and Vehicle Control, MIT Man-Machine Systems Laboratory Report.

*Wood, W. (1982), Use of Machine Aids in Dynamic Multi-Task Environments: A Comparison of an Optimal Model to Human Behavior, MIT Man-Machine Systems Laboratory Report.

Yntema, D.B. and Torgerson, W.S. (1961), "Man-Computer Cooperation in Decisions Requiring Common Sense", IEEE Trans. on Human Factors in Electronics, Vol. HFE 2.

Yntema, D.B. and Klem (1965), Telling a Computer How to Evaluate Multi-dimensional Situations", IEEE Trans. on Human Factors in Electronics, Vol. HFE-6, No. 1.

*Yoerger, D. (1982), MIT PH.D. thesis, in progress.

*Asterisks indicate reports of research done on this ONR contract.

DISTRIBUTION LIST

CDR Paul R. Chatelier
Office of the Deputy Under
Secretary of Defense
OUSDRE (E&LS)
Pentagon, Room 3D129
Washington, D.C. 20301

Engineering Psychology Programs
Code 442
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217 (5 cys)

Undersea Technology Programs
Code 220
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Dr. Eugene Silva
Ocean Sciences, Code 421
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Communication & Computer Technology
Programs
Code 240
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Manpower, Personnel, and Training
Programs
Code 270
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Information Sciences Division
Code 433
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Psychology & Neuro Biology Programs
Code 441B
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Special Assistant for Marine
Corps Matters
Code 100M
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Commanding Officer
ONR Eastern/Central Regional Office
ATTN: Dr. J. Lester
Building 114, Section D
666 Summer Street
Boston, MA 02210

Commanding Officer
ONR Western Regional Office
ATTN: Mr. R. Lawson
1030 East Green Street
Pasadena, CA 91106

Dr. Perry Alers
Marine Systems Branch, Code 5821
Naval Research Laboratory
Washington, D.C. 20375

Director
Naval Research Laboratory
Technical Information Division
Code 2627
Washington, D.C. 20375

Dr. Michael Melich
Communications Sciences Division
Code 7500
Naval Research Laboratory
Washington, D.C. 20375

Office of the Chief of Naval Operations
Deep Submergence Systems Division
OP-231
Washington, D.C. 20350

Dr. Robert G. Smith
Office of the Chief of Naval
Operations, OP 987H
Personnel Logistics Plans
WASHINGTON, D.C. 20350

CDR G. Worthington
Office of the Chief of Naval
Operations, OP 372G
Washington, D.C. 20350

Dr. W. Mehuron
Office of the Chief of Naval Operations,
OP 987
Washington, D.C. 20350

Dr. Andreas B. Rechnitzer
Office of the Chief of Naval
Operations, OP 952F
Naval Oceanography Division
Washington, D.C. 20350

Dr. Jerry C. Lamb
Combat Control Systems
Naval Underwater Systems Center
Newport, RI 02840

Human Factors Department
Code N215
Naval Training Equipment Center
Orlando, FL 32813

Mr. John Freund
NAVSEA 05R2
Naval Sea Systems Command
Washington, D.C. 20362

Dr. Gary Poock
Operations Research Department
Naval Postgraduate School
Monterey, CA 93940

Mr. H. Talkington
Ocean Engineering Department
Naval Ocean Systems Center
San Diego, CA 92152

Mr. Paul Heckman
Naval Ocean Systems Center
San Diego, CA 92152

Mr. Warren Lewis
Human Engineering Branch
Code 8231
Naval Ocean Systems Center
San Diego, CA 92152

Dr. Ross L. Pepper
Naval Ocean Systems Center
Hawaii Laboratory
P. O. Box 997
Kailua, HI 96734

Dr. A. L. Slafkosky
Scientific Advisor
Commandant of the Marine Corps
Code RD-1
Washington, D.C. 20380

Naval Material Command
NAVMAT 0722 - Room 508
800 North Quincy Street
Arlington, VA 22217

Mr. Phillip Andrews
Naval Sea Systems Command
NAVSEA 0341
Washington, D.C. 20362

Mr. Milon Essoglou
Naval Facilities Engineering Command
R&D Plans and Programs
Code 03T
Hoffman Building II
Alexandria, VA 22332

CDR Robert Biersner
Naval Medical R&D Command
Code 44
Naval Medical Center
Bethesda, MD 20014

Dr. Arthur Bachrach
Behavioral Sciences Department
Naval Medical Research Institute
Bethesda, MD 20014

CDR Thomas Berghage
Naval Health Research Center
San Diego, CA 92152

Dr. George Moeller
Human Factors Engineering Branch
Submarine Medical Research Lab
Naval Submarine Base
Groton, CT 06340

Dr. Robert Blanchard
Navy Personnel Research and
Development Center
Command and Support Systems
San Diego, CA 92152

Mr. J. Williams
Department of Environmental Sciences
U.S. Naval Academy
Annapolis, MD 21402

Human Factor Engineering Branch
Naval Ship Research and Development
Center, Annapolis Division
Annapolis, MD 21402

Dr. James McGrath
CINCLANT FLT HQS
Code 04E1
Norfolk, VA 23511

Mr. John Quirk
Naval Coastal Systems Laboratory
Code 712
Panama City, FL 32401

CDR W. Moroney
Code 55MP
Naval Postgraduate School
Monterey, CA 93940

Dr. Joseph Zeidner
Technical Director
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Director, Organizations and Systems
Research Laboratory
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Technical Director
U.S. Army Human Engineering Labs
Aberdeen Proving Ground, MD 21005

U.S. Air Force Office of Scientific
Research
Life Sciences Directorate, NL
Bolling Air Force Base
Washington, D.C. 20332

Chief, Systems Engineering Branch
Human Engineering Division
USAF AMRL/HES
Wright-Patterson AFB, OH 45433

Director, Human Factors Wing
Defence & Civil Institute of
Environmental Medicine
Post Office Box 2000
Downsview, Ontario M3M 3B9
CANADA

Defense Technical Information Center
Cameron Station, Bldg. 5
Alexandria, VA 22314 (12 cys)

Dr. Craig Fields
Director, System Sciences Office
Defense Advanced Research Projects
Agency
1400 Wilson Boulevard
Arlington, VA 22209

Dr. M. Montemerlo
Human Factors & Simulation
Technology, RTE-6
NASA HQS
Washington, D.C. 20546

Dr. J. Miller
Florida Institute of Oceanography
University of South Florida
St. Petersburg, FL 33701

Dr. Robert R. Mackie
Human Factors Research, Inc.
5775 Dawson Avenue
Goleta, CA 93017

Dr. H. McI. Parsons
Human Resources Research Office
300 N. Washington Street
Alexandria, VA 22314

Dr. Jesse Orlansky
Institute for Defense Analyses
400 Army-Navy Drive
Arlington, VA 22202