U. S. AIR FORCE

PROJECT RAND

RESEARCH MEMORANDUM

THE SYSTEMS RESEARCH LABORATORY
AND ITS PROGRAM

R. L. Chapman,
with assistance from
W. C. Biel
J. L. Kennedy
A. Newell

RM-890

7 January 1952

Assigned to ____________________

This is a working paper. It may be expanded, modified, or withdrawn at any time. The views, conclusions, and recommendations expressed herein do not necessarily reflect the official views or policies of the United States Air Force.

The RAND Corporation
1700 MAIN ST. • SANTA MONICA • CALIFORNIA

TABLE OF CONTENTS

THE SYSTEMS RESEARCH LABORATORY AND ITS PROGRAM

Summary

This memorandum describes (1) a research program and (2) the laboratory in which the program is being pursued. The research is devoted to the study of man-machine systems, with particular emphasis on the men and their organizational properties. The laboratory has been constructed in such a way that a number of particular arrangements of men and machines may be studied. A generalized model, the Information Processing Center, for organizational research, is described in some detail.

I. INTRODUCTION

It would be redundant to cite the history of the growing concern for the performance of today's complex military systems. The event which prompts this addition to the verbage devoted to the subject is the birth of the Systems Research Laboratory (SRL), a new element of RAND.

How can the program of this new laboratory be distinguished from the research efforts which precede and parallel it?

The Systems Research Laboratory will be studying particular kinds of models -- models made of metal, flesh and blood. Many of the messy and illusive variables of human and hardware interactions will be put into the laboratory.

But if this is the object -- the study of living models, why not observe systems in real life, why put them in a laboratory? There are a number of reasons: (1) A range of variable values can be used. Constraints can be applied to maintain fixed values of certain variables, while freedom to manipulate others to the extent that would be impractical in real life can be achieved; (2) the scientist can build himself a good vantage point from which to watch a system operate; and (3) data can be collected efficiently there. In effect, the laboratory becomes a specialized computer which grinds out the consequences of humans interacting with hardware and each other only as they know how.

Once the model functions in the laboratory, it can be subjected to experimental assault. It can be made to perform

and perform. Does it do what it is supposed to do and how well? Why? The number of ways its internal and external behavior can be measured and these hard numbers sorted, analyzed and related is limited only by the talent and stamina of the researchers involved.

To study these models scientifically, experimental techniques will need to be stretched beyond their present development and a class of models will have to be chosen which the stretched techniques may possibly encompass. Such a research program takes cognizance of those scientists who like their concepts sharp and the manipulations rigorous as well as those scientists who like their concepts round and fully-packed and the manipulations unctuous.

While aspirations of this order are fine in principle, the degree to which scientific results are achieved is a question of practice. An essential ingredient of effective practice is an idea that can be translated into experimental form.

II. DESCRIPTION OF THE MODEL

As many approaches to systems problems can be found as scientists working on them. At this point, it would be difficult to prove that one approach is better than another. With due regard paid to other research efforts past and present, the SRL staff has evolved a series of research attitudes which it considers to be most fruitful. The specification of the model is a case in point.

The chosen model has these characteristics and can be called the Information Processing Center (IPC):

Human Elements:	The Information Processing Center has a number of human operators. There is division of labor among these humans and a command structure or hierarchy as well.
Function:	The IPC's task is to process one commodity -- information. By means of gathering, integrating, and disseminating information (processing), the IPC controls the total system of which it is a part.
Information Flow:	Processing implies flow. This flow of information is constrained by communication channels and is characterized by a good deal of feedback.
Continuous Input and Output:	The nature of the input and output is that a volume of information continues to flow. The IPC performs a series and variety of integrations, rather than one discrete operation.
Information Content:	The information that flows in, through, and out of the IPC is substantially objective in character. That is, it lends itself to scientific parameterization more so than so-called social communication.

The Information Processing Center is a small system common to a great many larger systems. Predicting its performance can be recognized as an important and also general question.

The term "information processing" rather than "decision-making" is chosen advisedly. Information processing is perhaps a more general term describing a class of behavior, of which decision-making may be a part.

How the choosing of models with these specification helps structure a research program can best be discussed once the research facilities have been described and some additional concepts have been explored.

III. THE RESEARCH FACILITIES

The Systems Research Laboratory has been built to handle Information Processing Centers of different sizes and shapes. Flexibility has been the keynote in its design.

There are two large areas in which IPCs can be accommodated. The larger is 35 by 31 feet with a 20-foot ceiling; the smaller is 18 by 28 feet with a 14-foot ceiling. Thus, amphitheater arrangements could be built and large vertical displays or screens for projection can be installed in both. The smaller of the two would be useful for briefing, de-briefing, and conferences.

A reception room separates the smaller laboratory area from the building entrance. A control room provides space for auxiliary equipment.

Observation facilities are an integral part of the layout. A mezzanine running the full width of the laboratory provides visual access to both experimental areas while a room adjoining the smaller lab can be used for observation as well.

Figure 1 shows the relation of these areas to each other, while Figure 2 pictures the vantage point provided by the observation deck.

A major item of equipment for Information Processing Centers is the communication system. A custom-built telephone net is provided for this purpose. This telephone net has 100 circuits which terminate in jacks around the laboratory. By means of a switchboard, these circuits can be patched together in two-way conversations or in conferences. Handsets, headsets, or switch boxes can be plugged into the wall jacks to duplicate the

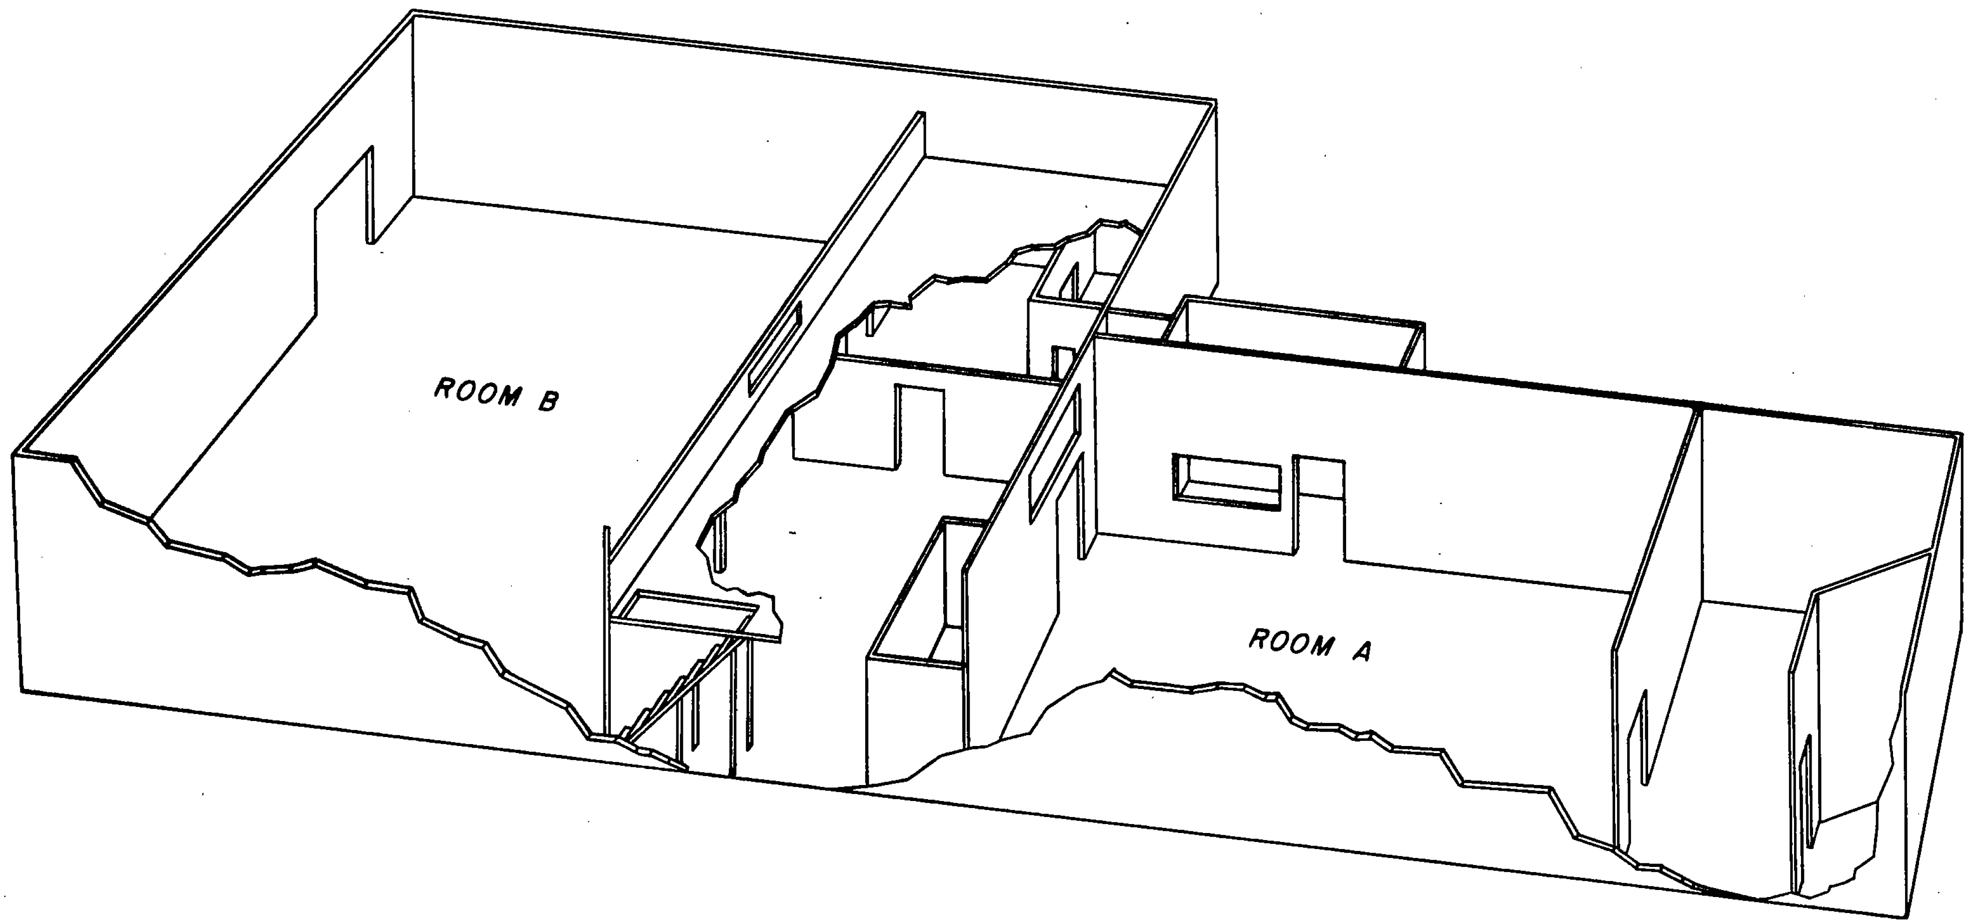

Fig .1— Plan of the systems research laboratory.

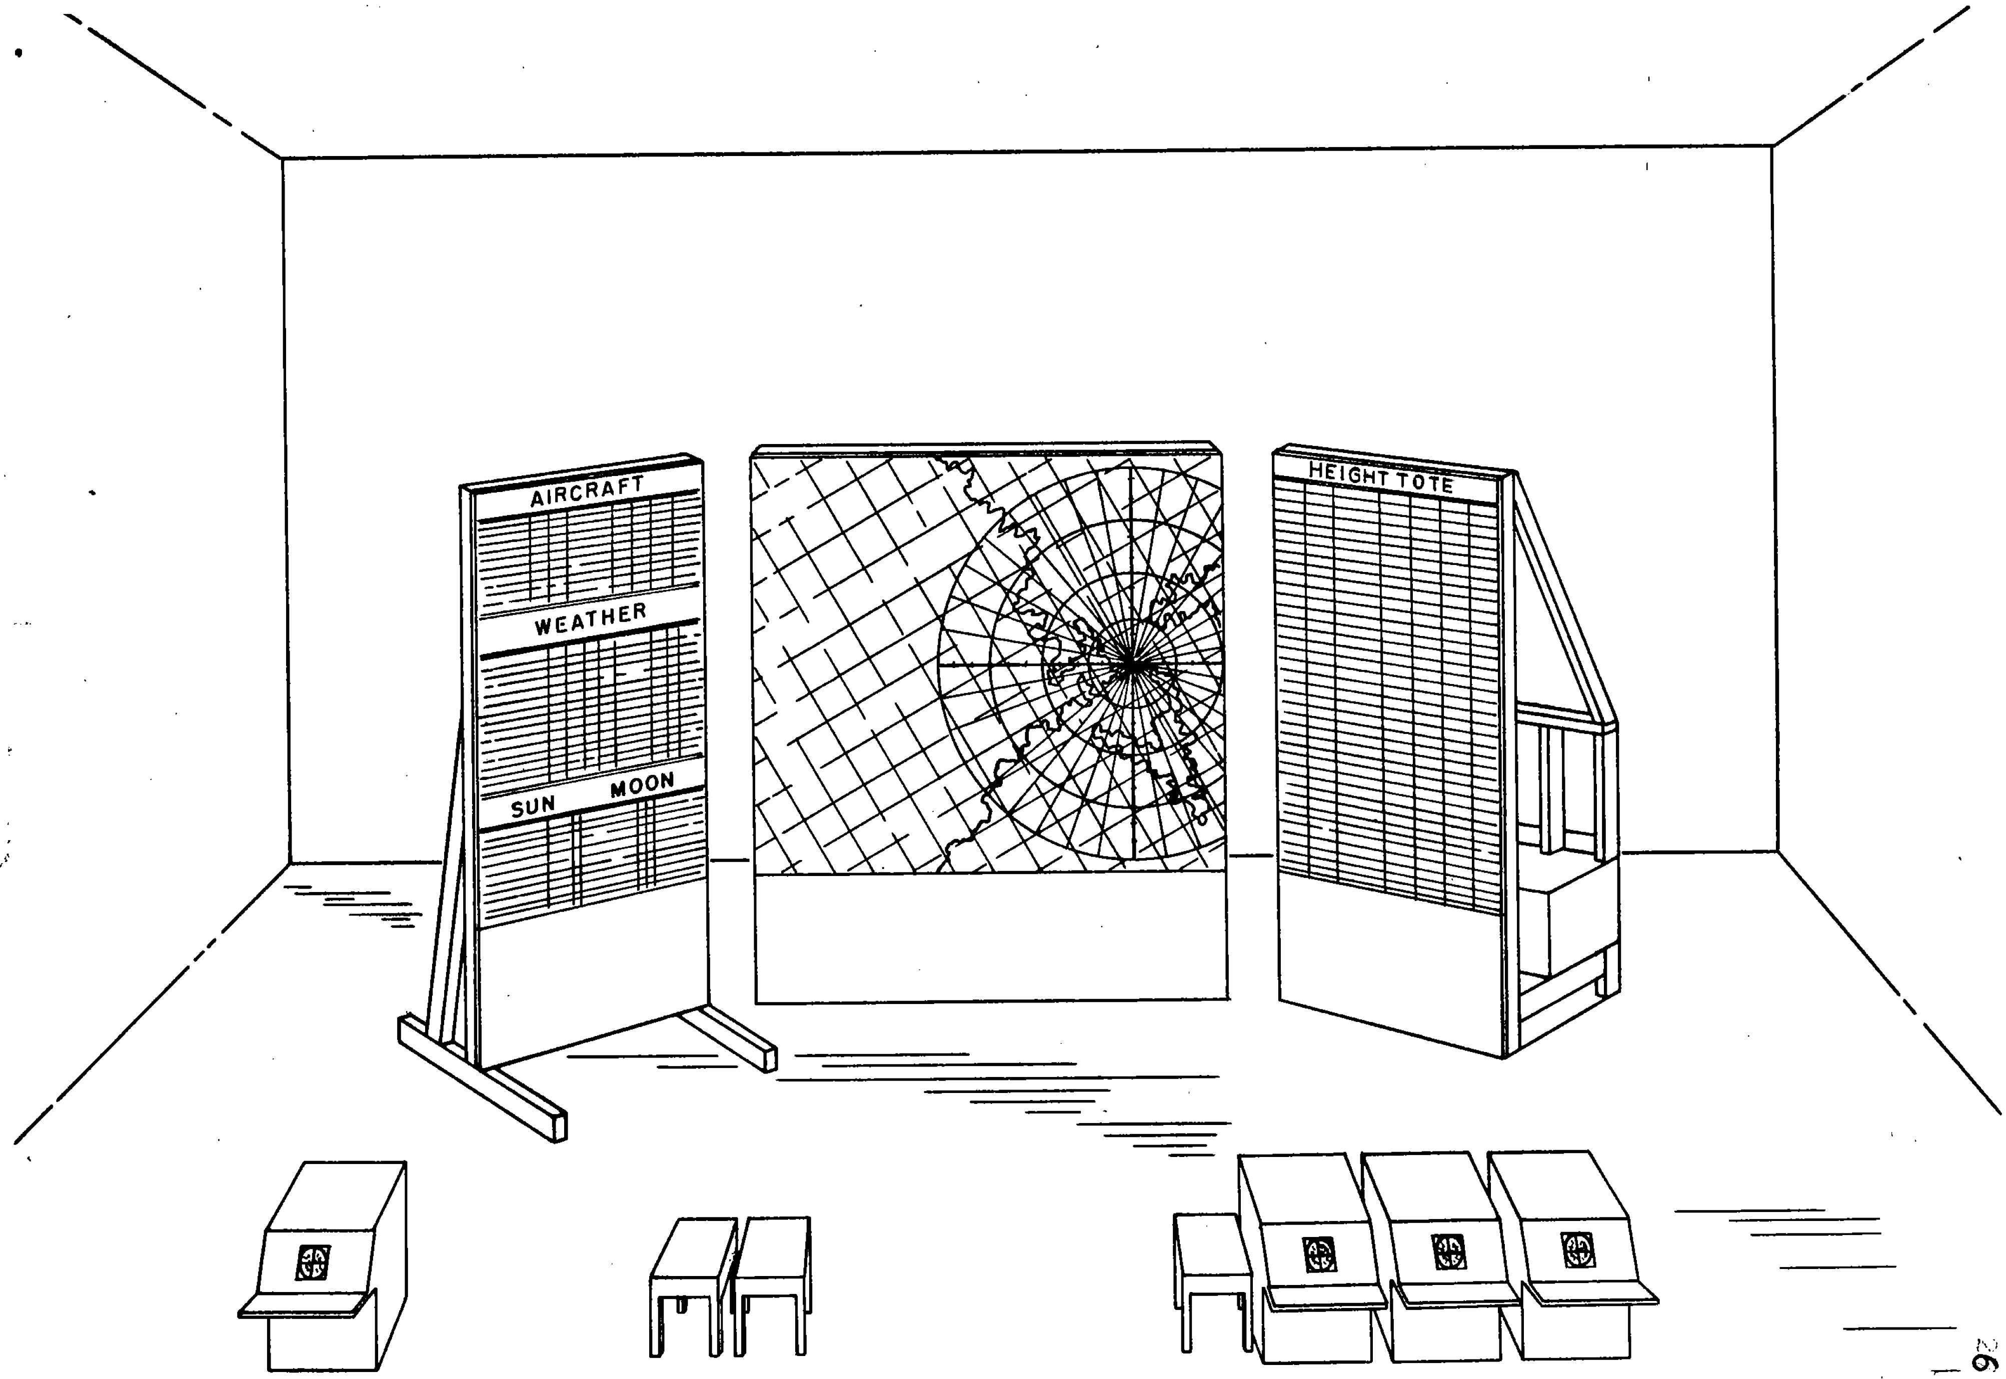

1166-2

Fig .2—View of experimental room B from observers deck.

arrangement of facilities in the IPC being studied. The switch boxes permit a number of lines to be brought in and served by one telephone instrument. The individual so accommodated can contact the persons to whom he is connected by means of the switches on his box.

A mike pickup system enables the experimenters to catch the conversations not directed into the telephone net. Nine microphones can be placed so that oral behavior of this sort can be monitored and/or recorded.

An experimenters' intercom completes the communication facilities. This intercom connects the lab spaces in which experimenters will be stationed and provides a means of addressing the subjects when starting and stopping experiments and giving instructions.

A monitoring station on the mezzanine enables the experimenters to hear as well as see what is going on. Each of the telephone circuits and mike outputs is brought to the monitor station so that any conversation can be tapped at will. A battery of fifteen recorders located in the control room can be cut into these lines selectively from this same point. A periodic time tone is generated and piped to each of the recorders so that the experimenters will be able to pinpoint the verbal events as to time of occurrence.

The inputs to IPCs can assume a number of forms. Written information such as might come off a teletype machine is one example. Graphic information of the sort that might appear on a radar scope is another. These, and others, can all be prepared and tabulated by IBM machinery. A presentation device has been

constructed to expose different amounts of printed information at different rates. This mechanism contains a carriage from an IBM tabulator by means of which a new 12 x 12 sheet could be shown every 10 seconds as easily as one additional typed line every two minutes. A drawing of the presentation instrument, of which there will be several, is contained in Figure 3.

IV. THE ORGANISM CONCEPT

What is a system? Webster says it is "an assemblage of objects united by some form of regular interaction or interdependence". Systems get built because there is a complicated job to be done requiring more than a single component.

Suppose a designer were to build a system, say an Information Processing Center, out of machines alone. He would have to find a definition of performance success, examine the environment in which the system would exist to parameterize that part of it relevant to task accomplishment, and break the job into pieces for which he could design black boxes. Each black box would have to have its input parameterized, its operation programmed, and its output parameterized in a form acceptable to the next black box in the system.

It is often too difficult, uneconomical, or cumbersome at the present state of the art to design a system completely in the metal. And so human beings are inserted at certain points in the system. Man-machine systems will, in many cases, provide acceptable performance part of the time. Not only does man prove to be more economical and less cumbersome on some

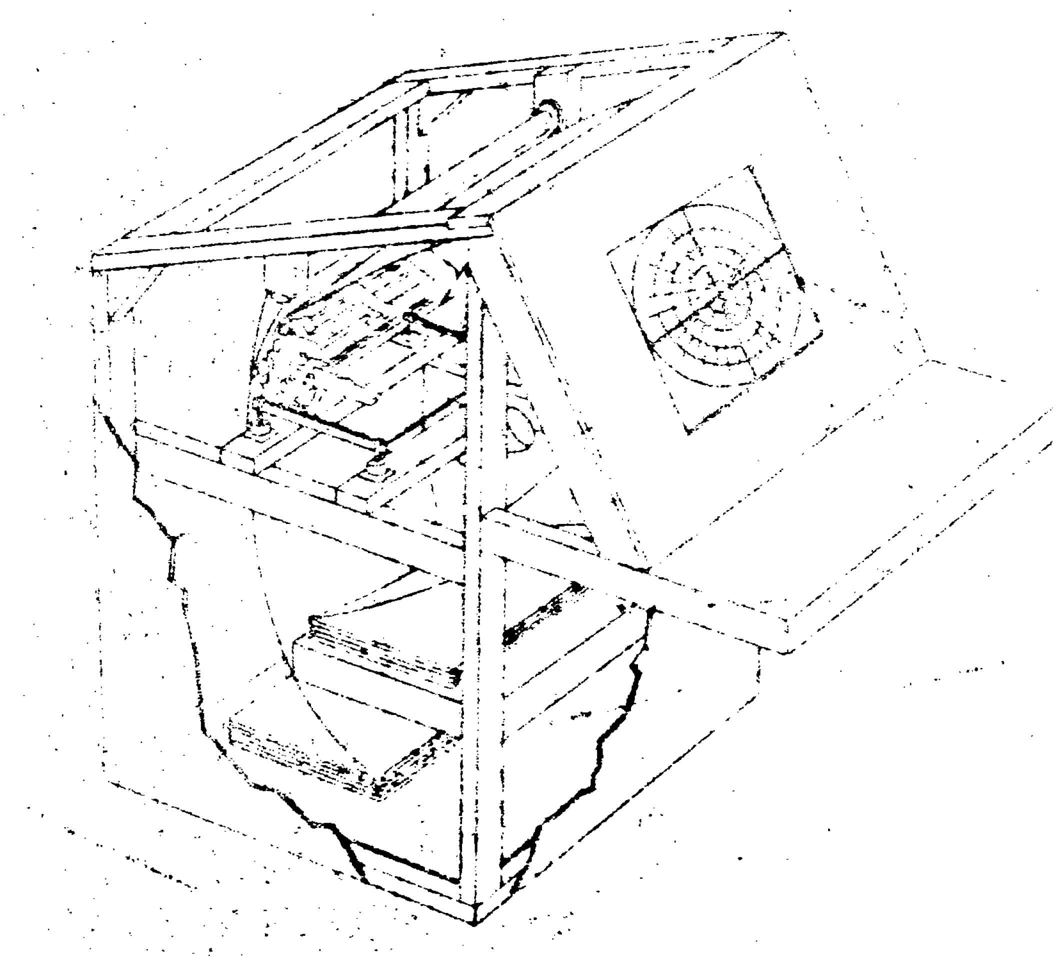

Fig. 3 — Input presentation device

occasions, but, it seems, he can work out a program for himself. He can learn.

What is involved in the learning process? There must, in the first place, be motivation to learn. (This also goes by the name of drive and need; having a criterion of performance is somewhat analagous.) The operator must distinguish the relevant parameters in the environment and select appropriately from his population of responses. Making the proper response at the right time is learning when it satisfies the operator's dominant motivation. When this happens, the operator is said to be "reinforced". Once reinforced, a response is more likely to be repeated. For example, a rifleman makes a number of responses in trying to hit the target. His motivation (trying to hit the target) is satisfied when he does.

Once learned is not permanently learned, however. The human will not repeat responses ad infinitum without further reinforcement. He will "extinguish", or fail to respond in that fashion, if he gets no more reinforcement.

It would seem that once a human operator had learned to do a task successfully, that either he or a skilled observer could describe the environmental parameters to which the operator attended and the operations that he performed. The difficulty posed for the observer is apparent -- many of the operations that take place are not exposed to view. They take place in the operator's "head". Why, then, does not the operator describe this phenomenon? The deplorable fact is that a human can learn without knowing what he learned or even that he had learned.

The individual operator often works at a task where he can learn effective performance by experiencing success or failure solely through his own efforts. This is not the case for the human element in a system. The system member contributes only partially to a goal achieved by the group as a whole. The desired dominant motivation, or criterion of performance, is system success, which does not necessarily follow effective component performance. Thus, the essential learning may not take place.

This discussion has alluded to two issues: (1) that the interactions between components confound the process of simple addition of elements to get a system; and (2) that human elements can learn to parameterize the environment and program their operation. The distinction between man and machine in the ability to learn is a very fuzzy separation indeed. Machines have been built that learn to some extent. Further, it has been found that even metal components have the distressing habit of compensating for each other's excesses and limitations. This is adaptive behavior of a sort, much akin to learning. Although adaptive behavior in the wider sense may yet remain the province of the human, the moral of the story is the same. There is an assemblage of components united by some form of regular interaction or interdependence whose performance can be studied only as a unit.

This is the organism concept. An organism is a highly complex structure with parts so integrated that their relation to one another is governed by their relation to the whole.

By means of this concept, the experimenter can delimit his problem, including in his model those elements that interact and are interdependent. A model containing fewer than all the interacting parts would leave out something essential.

Inasmuch as organism is used here in the philosophic rather than the biological sense, "the whole" is recognized as the performance criterion for which the system was designed. A clear-cut operational definition of system success is essential for system study. Where humans are components of a system, their learning capacity is not realized without explicit identification of their dominant motivation with system purpose and reinforcement of proper response with system success.

V. THE RESEARCH PROGRAM IN BRIEF

It would be well to state the research program in broad outline before proceeding with more intimate details. Getting the model into the laboratory, training the organism, and experimentation are its successive phases.

Once a real-life Information Processing Center had been selected for study, these steps will be taken to get it into the laboratory. The inputs to the actual IPC will be examined and parameterized. The forms that these inputs took will, together with the parameters, govern the preparation of problems on IBM machinery.

The structure of the real-life IPC will be investigated with the view to duplicating the communication net and defining the functions of each member of the team. Team members will

be obtained, either from the actual IPCs or from other sources, invoking the selecting criteria desired.

After a check of the laboratory representation by members of the actual system, training sessions will be started. The experimenters will endeavor to influence individual members' dominant motivation in the direction of the explicitly stated system purpose. Performance will be reinforced with knowledge of group performance. Organized group conferences will be held to promote group interaction and to accelerate the learning process.

When performance levels off, experimentation will begin. The substance of the research program is here in performance. Performance and more performance.

The Systems Research Laboratory has been referred to as a specialized computer, grinding out the consequences of human and hardware interaction. The hard core of the system is the IPC, not the laboratory or the experimental group. The experimenters have the task of recording the answers and comprehending the nature of the processes in force. To that end, advantage will be taken of all mechanical and administrative means to free the experimenters from routine. Inputs will be prepared and presented by machine. Performance will be recorded. Transcription of content will be organized for clerical accomplishment. These data will be marked on IBM cards and analyzed by machine.

The SRL Staff will meet the scientific problem of describing the phenomenon under observation with a balanced approach. Some of the staff is addicted to initial model building and

subsequent verification or modification; others feel that behavior can be translated into a variety of measures and that concepts can be found from analytic ordering of these data. This issue will fortunately be pursued in an ideal environment, in the laboratory, where fact is the great leveler.

Research can be conducted to serve several objects: one is to comprehend the phenomenon, and another is to influence the outcome of a process. These are not necessarily synonomous aims. The disparity between the two is not marked when a precise criterion of performance is available. Thus, the program of the Systems Research Laboratory will work towards both goals.

VI. THE ENVIRONMENT, THE STIMULUS, AND THE INPUT

The classic learning paradigm is stimulus → organism → response. To study an organism's response, a stimulus must be presented to it. What serves as a stimulus for that organism is determined by its environment but influenced by its structure and sensory capacity. The stimulus need not represent a one to one correspondence with reality.

For a more explicit formulation of the problem, consider Figure 4. The environment can be abstracted by use of the parameters--e_1, e_2 ....e_n -- which may be related to each other. The organism has a number of receptor elements--a, b, c--which in man-machine systems are often hardware elements. These receptors are stimulated by certain qualities -- s_1, s_2, s_3 ....s_n --which may be functions of a number of environment parameters. These sensory elements may perceive on a selective basis, either systematic

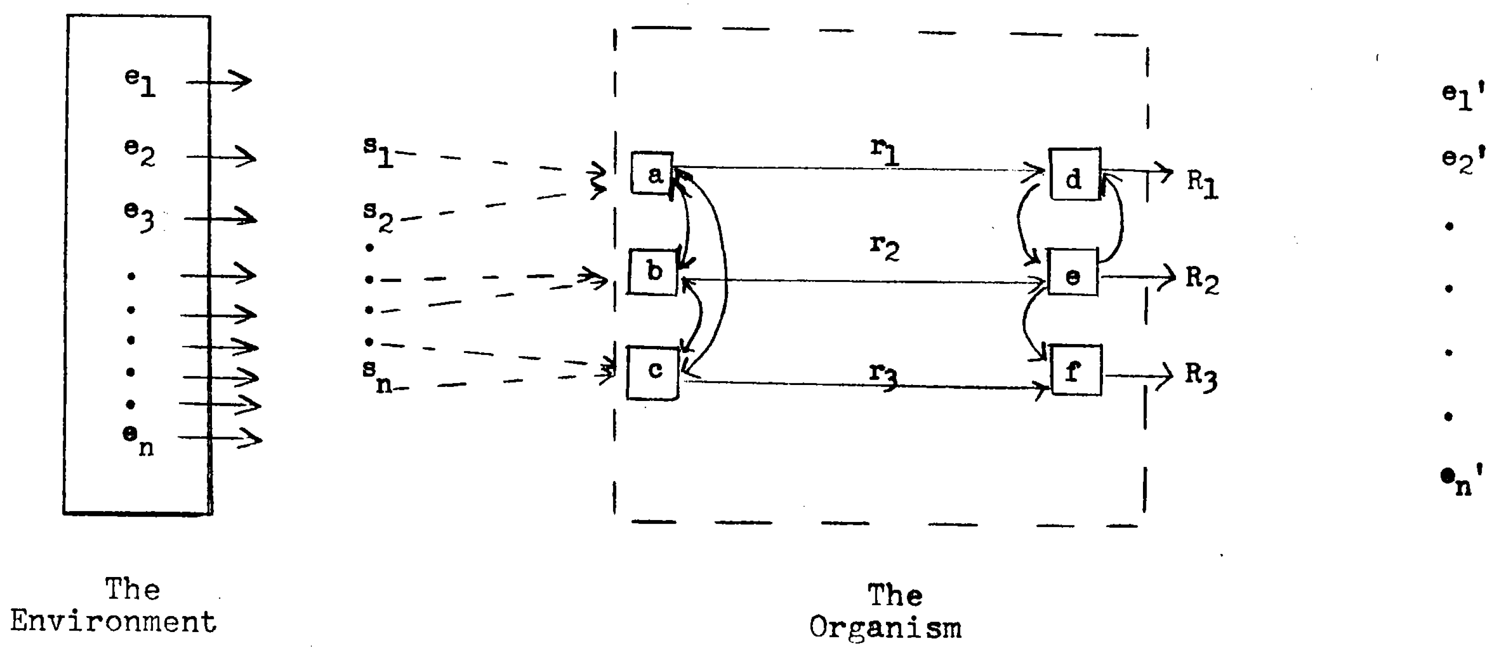

Fig. 4 -- Schematic Diagram of the Environment, the Stimulus, the Organism and Its Response.

or random. The stimuli prompt behavior of the organism--r_1, r_2, r_3, etc-- and a subsequent set of responses--R_1, R_2, R_3-- which in turn either change the environment parameters into the form e'_1, e'_2, e'_3,....e'_n or adjust the organism to e_1, e_2, e_3....e_n.

The experimenter may choose to control as experimental variables either the environment or stimulus parameters, depending on the question he asks. In choosing to control the environment parameters, he addresses himself to the relation between aspects of real life environment and the organism's behavior under certain conditions of hardware performance. When controlling the stimulus parameters, the experimenter has more direct access to stressing certain elements of the organism (not only the receptors--a, b and c-- but also such elements as d which is prodded by r_1, a response of a to s_1 and s_2) under certain constraints on its operation.

Scientists have frequently examined an organism in its environment and chosen as the experimental variable the obvious stimulus prompting behavior. When a complex organism exists in an equally complicated environment, however, the stimulus conditions are not so obvious. This is expecially true if the organism has any latitude in adjusting its structure to meet the situation.

Parameterizing the environment of an organism necessarily implies consideration of the organism and its function. But such a process can be pursued with more freedom than in the case of parameterizing the stimulus where the perspective is dictated

by assumptions regarding its detailed operating characteristics.

The input to the laboratory organism will always be the stimulus whether experimental control is affected there or in the prior domain, environment. The preparation of the input will, however, depend on the choice made. When the control point is in the environment, the receptor hardware design and performance characteristics determine the way the inputs are selected from the environment which is constructed according to certain parameter values. On the other hand, when the stimulus parameters are controlled, the input is constructed directly. Reference is then made to the environmental concommitants to see whether reality has been violated. These are reverse computation processes. In one, the population is defined and a sample drawn from it; in the other, the sample is defined and the population derived from it and checked against real life.

Choosing between these two ways of selecting experimental variables is a very sticky problem. Most inputs are undoubtedly built with regard paid to both points of view.

Perhaps, however, there is a time and place for emphasis in each direction. In the initial stage of a research program, the "safe bet" would seem to be to parameterize the environment. Thus, a control base is used which operates irrespective of the adjustment the organism makes to it. Further, receptor performance can be included as an organism variable so that the magnitude of that effect can be determined in comparison with others.

When optimum operating procedures have been adopted and the nature of the stimulus parameters more precisely ascertained, the

more immediate leverage of stimulus variables can be utilized with greater confidence.

VII. THE DESIGNER, THE EXECUTIVE, AND THE MANAGER OF THE ORGANISM

Typically, the organism is designed to have a considerable number of both black boxes and humans. When humans are involved, the executive joins the designer in planning the system and defining its purpose. Both are immediately concerned with two organism variables: structure and element capability. Division of labor among black boxes and individuals, command structure or hierarchy, and the communication system are all aspects of structure. The individual elements--whether machines or humans--have certain capabilities. There are obvious differences in capability between the two classes of elements, men and machines; the differences within each class may be equally large. The designer and executive work these problems out between them: the designer builds machines to do certain jobs; the executive selects individuals for particular positions on the basis of experience, past performance, and, perhaps, their talents as measured by psychological tests.

When these elements are combined into a structure, most particularly if humans are involved, the organism may or may not perform satisfactorily. As has been stated previously, the interactions between elements mitigate against prediction of performance by the simple process of adding element capabilities. For example, some experts have maintained for the past several years that the Red Sox are the best team in baseball--on paper. There is some

additional quality which contributes to successful operation. Supplying this quality is the problem passed to the manager.

The manager is concerned not with whether the organism should or should not perform as designed but with whether it does or does not. He tries to supply motivation to the humans--sell the team's purpose to them; he modifies the structure and placement of talent within a limited range; he tries to maintain an optimum learning atmosphere by giving the individuals freedom to develop more effective techniques through participation; and he supplies a necessary part of the learning process, rewarding and punishing the individuals. The successful response can be learned and it will be remembered, only when it is reinforced.

The manager's function may be spread out over a number of individuals. Traditionally, the "boss" exercises this function. However, as a matter of fact, whether because of design, default, or malfunction, certain aspects of managerial function may come to rest, even in an effective organism, far from the titular head of the group.

The principal unique talent that man has, which justifies his inclusion in the organism, is his ability to learn. The manager's job is to mobilize and administrate this resource. Under his influence, learning is facilitated, performance maintained once learned, and relearning carried out when conditions change.

The distinction between the executive and the manager is made to point up the difference between the "boss's" continuous and discontinuous functions. (The same man may fill both jobs on occasion.) The executive, together with the designer, explores

the technical, economic, and political limits on further improvement of the structure and capabilities of the elements. Planning the organism is done with symbols and abstractions. On the other hand, the manager lives where life is real and life is earnest. The environment is there, not a set of parameters. The interactions between elements occur even if the manager does not have a set of symbols to characterize them. And successes and failures are experienced whether they can be explained or not. The manager has fewer degrees of freedom than the executive to to modify the structure and capacity of the elements; for the most part, he has only the organism's ability to learn to work with.

The Systems Research Laboratory is concerned with the managerial function. Most Information Processing Centers have their structure and machine elements specified. Theoretically, they should work better than they do. Can the element interactions administered to by the manager be optimized so that the IPCs will work well? If so, can the internal behavior of a good organism be described and distinguished from that of a bad organism?

There is the question whether organism learning, or optimization of interactions, can be experienced vicariously. Where can managerial function, which has such a tremendous influence, be better studied than in a flesh and blood model?

What leverage does the researcher have to utilize managerial function as an experimental variable? For one thing, if human elements homogeneous with respect to their talents were chosen and they were rotated through the positions in the IPC, the patterns of information flow might change. That these differences

were due to managerial exercise rather than to the capabilities of particular individuals performing in certain positions would be a question of fact not an assumption. A concentrated effort would be made to restrict the range of talents of the humans. This effort would yield values for team member capabilities which could be correlated with changes in information flow patterns and organism successes and failues. If these correlations were significant, there would be reason to believe that individual differences rather than managerial exercise were causing performance variation.

By choosing an organism made up of a number of men and machines rather than an individual, the experimenter has, in effect, magnified the operation he wishes to study. The flow of information is now over a set of communication channels rather than in an individual's nerve net. A good deal of the interaction is brought out into the open where it can be observed.

The requirement has been stated that the information flowing in an IPC should be substantially objective in nature. The implication is that a task should be chosen whose accomplishment does not depend strongly upon individual's social history, social perceptions, or set of personal values. A series and variety of integrations rather than a discrete response has further qualified the nature of the task. Thus, the human members have a chance to settle down to their task; their performance can be judged on the basis of a number of acts rather than on a single complex response to a set of stimuli.

The concern that a good deal of the information flowing in

the IPC be objective exists for another reason as well. The experimenters need leverage to measure information content. Consideration will be given to the kinds of information that must necessarily flow in any organism in the following section of this report. It will suffice to say at this point that a base of readily parameterizable information will be needed. Then, at a minimum, hard numbers may be assigned to other kinds of information in terms of the effect they have on the flow of task performance information.

VIII. MEASURING THE ORGANISM'S BEHAVIOR AND RESPONSE

As the title of this section has indicated, there are two measurement problems: (1) finding out how well the organism does what it is supposed to do; and (2) measuring its internal behavior so that its successes and failures can be accounted for. The first is a straight-forward problem because the organism chosen for study, the IPC, was selected because performance criteria were available. The second problem, however, involves the codification of a rather large body of behavior. The experimenter needs to find a limited number of measures which describe what is happening with both parsimony and completeness.

In what ways may the behavior of an Information Processing Center be recorded? Time studies and/or micromotion studies could be made of the activities of the people involved; motion picture photographs could be taken of any large displays that are maintained; any forms that are completed in the course of the organism's operation could be collected for examination; the

affective reactions of the individuals could be noted; or the post-problem responses of the operators could be obtained by means of questionnaires, interviews or "de-briefing" group discussion.

The IPC processes only one commodity, information; a good deal of it will flow in verbal form. The Systems Research Laboratory is well set up to record this verbal behavior, both that which goes into the communication net and that which does not. Humans are addicted, unfortunately, to the use of a set of symbols, language, which is very hard to code. How, then, are these recordings to be reduced to a set of numbers for analysis?

A simple example may be helpful in considering the kinds of information that might flow in an organism. The following is a direct transcription of conversation between two operators, A and B in an IPC:

A: Talk real loud, now. (S)

B: Gosh, I had a target. I should have reported it. (T)
I have a target, 20 West 65 North. Another target-- (P)

A: Wait a minute (T)

B: OK (T)

A: Assign Track 1 (P) and give compass direction first. (S)

B: What's the compass direction (T)

A: Like West or North (S)

B: All right, I have a target at West 40 North 60 (P)

A: Assign Track 2 (P)

B: I have a target at West 25 North 60 (P)

A: Assign Track 3 (P)

Two kinds of information seem to be necessary in an organism: That which has to do with establishing the routine for problem solution, and that which carries out the operation once the routine is determined. The first of these might be called set information and the second, performance information. This distinction is apparent in many learning situations; how to do the problem must first be determined, then a variety of problems can be handled. This is commonly referred to as learning the "set".

In many IPCs, routine is still an open question. Perhaps the environment of a particular IPC can be met with a single routine. That routine must then be found through learning and maintained. But the learning process requires that another kind of information be considered. Somehow, each individual must find out how he is doing in relation to the system criterion if he is to learn to perform properly. This third category of information might be termed, for want of a better word, therapeutic. In the conversation above, examples of set information are labelled (S), performance information (P) and therapeutic information (T).

At this point, it is possible to recognize that only the performance information could be objective in nature. That part of the information flow that relates to actual task accomplishment, not to how the task is to be done, is required by definition to be objective. That it would contain explicit content such as numbers would be preferred. There are, of course, many organisms which meet this specification.

Of the three kinds of information suggested--performance.

set, and therapeutic--the performance information will be easiest to measure because it can be specified as objective. The set and therapeutic information flows to promote the processing of performance information. There are certainly different ways to process information; perhaps these processes can be described in the following way:

Translation:	Where information is transferred from graphic to verbal form, or vice versa.
Transmission:	Where information is simply passed on directly.
Dissemination:	Where items are selected from a population of information and sent to several of many possible addresses.
Verification:	Where information is somehow checked against another source.
First Derivative:	Where a single fact is inferred from a succession of bits of information.
Second Derivative:	Where changes in the first derivative are inferred.

Responses have also been measured, traditionally, as to their amount and accuracy. The latency of a response--that is, the delay between the presentation of the stimulus and the response--is a familiar measure as well. These three ways of measuring information and its flow, in addition to the previously mentioned classifications, will be related to organism performance against criteria to find optimum organism behavior.

The point of these descriptive concepts is to show that there are ways to measure information content and flow while avoiding reference to a particular task. If prediction of the behavior of a class of models is to be made, concepts which

can be generalized throughout the class must be developed. Many other measurement ideas will be generated as the research program continues. As in other scientific endeavors, the adequacy of the measurement techniques can only be judged in terms of the predictions that can be made about organism behavior.