

AUGMENTING HUMAN INTELLECT : A Conceptual Framework. October 1962. By D. C. Engelbart



*STANFORD RESEARCH INSTITUTE, Menlo Park, California 94025 * USA.AFOSR-3233
Summary Report. Prepared for: DIRECTOR OF INFORMATION SCIENCES, AIR FORCE
OFFICE OF SCIENTIFIC RESEARCH, WASHINGTON 25, D.C. CONTRACT AF49(638)-1024.
SRI Project No. 3578 Approved: R. C. Amara, Manager System Engineering Department, J. D.
Noe, Director Engineering Science Division*

Abstract

This is an initial summary report of a project taking a new and systematic approach to improving the intellectual effectiveness of the individual human being. A detailed conceptual framework explores the nature of the system composed of the individual and the tools, concepts, and methods that match his basic capabilities to his problems. One of the tools that shows the greatest immediate promise is the computer, when it can be harnessed for direct on-line assistance, integrated with new concepts and methods.

Foreword

This report describes a study that was carried on at Stanford Research Institute under the joint sponsorship of the Institute and the Directorate of Information Sciences of the Air Force Office of Scientific Research [Contract AF 49(638)-1024]. Mrs. Rowena Swanson was the AFOSR Project Supervisor for this study.

INTRODUCTION

GENERAL

OBJECTIVE OF THE STUDY

CONCEPTUAL FRAMEWORK

GENERAL

THE BASIC PERSPECTIVE

DETAILED DISCUSSION OF THE H-LAM/T SYSTEM

The Source of Intelligence

Intelligence Amplification

Two-Domain System

Concepts, Symbols, and a Hypothesis

Capability Repertoire Hierarchy

Basic Capabilities

Structure Types

General

Mental Structuring

Concept Structuring

Symbol Structuring

Process Structuring

Physical Structuring

Interdependence and Regeneration

Roles and Levels

Model of Executive Superstructure

Flexibility in the Executive Role

Compound Effects

EXAMPLES AND DISCUSSION

BACKGROUND

What Vannevar Bush proposed in 1945

Comments Related to Bush's Article

Some Possibilities with Cards and Relatively Simple Equipment

An Existing Note and File System
Comments on the System
Associative-Linking Possibilities
An Experiment Illustrating Usage and Further System Possibilities
A Quick Summary of Relevant Computer Technology
Other Related Thought and Work

HYPOTHETICAL DESCRIPTION OF COMPUTER-BASED AUGMENTATION SYSTEM

Background

Single-Frame Composition

Single-Frame Manipulation

Structuring an Argument

General Symbol Structuring

Process Structuring

Team Cooperation

Miscellaneous Advanced Concepts

RESEARCH RECOMMENDATIONS

OBJECTIVES FOR A RESEARCH PROGRAM

BASIC RESEARCH CONDITIONS

WHOM TO AUGMENT FIRST

BASIC REGENERATIVE FEATURE

TOOLS DEVELOPED AND TOOLS USED

RESEARCH PLAN FOR ACTIVITY A 1

A Second Phase In The Research Program

SUMMARY

CONCLUSIONS

REFERENCES

INTRODUCTION

GENERAL

By “augmenting human intellect” we mean increasing the capability of a man to approach a complex problem situation, to gain comprehension to suit his particular needs, and to derive solutions to problems. Increased capability in this respect is taken to mean a mixture of the following: more-rapid comprehension, better comprehension, the possibility of gaining a useful degree of comprehension in a situation that previously was too complex, speedier solutions, better solutions, and the possibility of finding solutions to problems that before seemed insoluble. And by “complex situations” we include the professional problems of diplomats, executives, social scientists, life scientists, physical scientists, attorneys, designers -- whether the problem situation exists for twenty minutes or twenty years. We do not speak of isolated clever tricks that help in particular situations. We refer to a way of life in an integrated domain where hunches, cut-and-try, intangibles, and the human “feel for a situation” usefully coexist with powerful concepts, streamlined terminology and notation, sophisticated methods, and high-powered electronic aids.

Man's population and gross product are increasing at a considerable rate, but the complexity of his problems grows still faster, and the urgency with which solutions must be found becomes steadily greater in response to the increased rate of activity and the increasingly global nature of that activity. Augmenting man's intellect, in the sense defined above, would warrant full pursuit by an enlightened society if there could be shown a reasonable approach and some plausible benefits.

This report covers the first phase of a program aimed at developing means to augment the human intellect. These “means” can include many things -- all of which appear to be but extensions of means developed and used in the past to help man apply his native sensory, mental, and motor capabilities -- and we consider the whole system of a human and his augmentation means as a proper field of search for practical possibilities. It is a very important system to our society, and like most systems its performance can best be improved by considering the whole as a set of interacting components rather than by considering the components in isolation.

This kind of system approach to human intellectual effectiveness does not find a ready-made conceptual framework such as exists for established disciplines. Before a research program can be designed to pursue such an approach intelligently, so that practical benefits might be derived within a reasonable time while also producing results of long range significance, a conceptual framework must be searched out -- a framework that provides orientation as to the important factors of the system, the relationships among these factors, the types of change among the system factors that

offer likely improvements in performance, and the sort of research goals and methodology that seem promising.

In the first (search) phase of our program we have developed a conceptual framework that seems satisfactory for the current needs of designing a research phase. Section II contains the essence of this framework as derived from several different ways of looking at the system made up of a human and his intellect-augmentation means.

The process of developing this conceptual framework brought out a number of significant realizations: that the intellectual effectiveness exercised today by a given human has little likelihood of being intelligence limited -- that there are dozens of disciplines in engineering, mathematics, and the social, life, and physical sciences that can contribute improvements to the system of intellect-augmentation means; that any one such improvement can be expected to trigger a chain of coordinating improvements; that until every one of these disciplines comes to a standstill and we have exhausted all the improvement possibilities we could glean from it, we can expect to continue to develop improvements in this human-intellect system; that there is no particular reason not to expect gains in personal intellectual effectiveness from a concerted system oriented approach that compare to those made in personal geographic mobility since horseback and sailboat days.

The picture of how one can view the possibilities for a systematic approach to increasing human intellectual effectiveness, as put forth in Section II in the sober and general terms of an initial basic analysis, does not seem to convey all of the richness and promise that was stimulated by the development of that picture. Consequently, Section III is intended to present some definite images that illustrate meaningful possibilities derivable from the conceptual framework presented in Section II -- and in a rather marked deviation from ordinary technical writing, a good portion of Section III presents these images in a fiction-dialogue style as a mechanism for transmitting a feeling for the richness and promise of the possibilities in one region of the improvement space" that is roughly mapped in Section II.

The style of Section III seems to make for easier reading. If Section II begins to seem unrewardingly difficult, the reader may find it helpful to skip from Section II-B directly to Section III. If it serves its purpose well enough, Section III will provide a context within which the reader can go back and finish Section II with less effort.

In Section IV (Research Recommendations) we present a general strategy for pursuing research toward increasing human intellectual effectiveness. This strategy evolved directly from the concepts presented in Sections II and III; one of its important precepts is to pursue the quickest gains first, and use the increased intellectual effectiveness thus derived to help pursue successive gains. We see

the quickest gains emerging from (1) giving the human the minute-by-minute services of a digital computer equipped with computer-driven cathode-ray-tube display, and (2) developing the new methods of thinking and working that allow the human to capitalize upon the computer's help. By this same strategy, we recommend that an initial research effort develop a prototype system of this sort aimed at increasing human effectiveness in the task of computer programming.

To give the reader an initial orientation about what sort of thing this computer-aided working system might be, we include below a short description of a possible system of this sort. This illustrative example is not to be considered a description of the actual system that will emerge from the program. It is given only to show the general direction of the work, and is clothed in fiction only to make it easier to visualize.

Let us consider an augmented architect at work. He sits at a working station that has a visual display screen some three feet on a side; this is his working surface, and is controlled by a computer (his "clerk") with which he can communicate by means of a small keyboard and various other devices.

He is designing a building. He has already dreamed up several basic layouts and structural forms, and is trying them out on the screen. The surveying data for the layout he is working on now have already been entered, and he has just coaxed the clerk to show him a perspective view of the steep hillside building site with the roadway above, symbolic representations of the various trees that are to remain on the lot, and the service tie points for the different utilities. The view occupies the left two-thirds of the screen. With a "pointer," he indicates two points of interest, moves his left hand rapidly over the keyboard, and the distance and elevation between the points indicated appear on the right-hand third of the screen.

Now he enters a reference line with his pointer, and the keyboard. Gradually the screen begins to show the work he is doing -- a neat excavation appears in the hillside) revises itself slightly, and revises itself again. After a moment, the architect changes the scene on the screen to an overhead plan view of the site, still showing the excavation. A few minutes of study, and he enters on the keyboard a list of items, checking each one as it appears on the screen, to be studied later.

Ignoring the representation on the display, the architect next begins to enter a series of specifications and data -- a six-inch slab floor, twelve-inch concrete walls eight feet high within the excavation, and so on. When he has finished, the revised scene appears on the screen. A structure is taking shape. He examines it, adjusts it, pauses long enough to ask for handbook or catalog information from the clerk at various points, and re-adjusts accordingly. He often recalls from the "clerk" his working lists of specifications and considerations to refer to them, modify them, or add

to them. These lists grow into an evermore-detailed, interlinked structure, which represents the maturing thought behind the actual design.

Prescribing different planes here and there, curved surfaces occasionally, and moving the whole structure about five feet, he finally has the rough external form of the building balanced nicely with the setting and he is assured that this form is basically compatible with the materials to be used as well as with the function of the building.

Now he begins to enter detailed information about the interior. Here the capability of the clerk to show him any view he wants to examine (a slice of the interior, or how the structure would look from the roadway above) is important. He enters particular fixture designs, and examines them in a particular room. He checks to make sure that sun glare from the windows will not blind a driver on the roadway, and the "clerk" computes the information that one window will reflect strongly onto the roadway between 6 and 6:30 on midsummer mornings.

Next he begins a functional analysis. He has a list of the people who will occupy this building, and the daily sequences of their activities. The "clerk" allows him to follow each in turn, examining how doors swing, where special lighting might be needed. Finally he has the "clerk" combine all of these sequences of activity to indicate spots where traffic is heavy in the building, or where congestion might occur, and to determine what the severest drain on the utilities is likely to be.

All of this information (the building design and its associated "thought structure") can be stored on a tape to represent the design manual for the building. Loading this tape into his own clerk, another architect, a builder, or the client can maneuver within this design manual to pursue whatever details or insights are of interest to him -- and can append special notes that are integrated into the design manual for his own or someone else's later benefit.

In such a future working relationship between human problem-solver and computer 'clerk,' the capability of the computer for executing mathematical processes would be used whenever it was needed. However, the computer has many other capabilities for manipulating and displaying information that can be of significant benefit to the human in non mathematical processes of planning, organizing, studying, etc. Every person who does his thinking with symbolized concepts (whether in the form of the English language, pictographs, formal logic, or mathematics) should be able to benefit significantly.

OBJECTIVE OF THE STUDY

The objective of this study is to develop a conceptual framework within which could grow a coordinated research and development program whose goals would be the following: (1) to find the

factors that limit the effectiveness of the individual's basic information-handling capabilities in meeting the various needs of society for problem solving in its most general sense; and (2) to develop new techniques, procedures, and systems that will better match these basic capabilities to the needs' problems, and progress of society. We have placed the following specifications on this framework:

That it provide perspective for both long-range basic research and research that will yield practical results soon.

That it indicate what this augmentation will actually involve in the way of changes in working environment, in thinking, in skills, and in methods of working.

That it be a basis for evaluating the possible relevance of work and knowledge from existing fields and for assimilating whatever is relevant.

That it reveal areas where research is possible and ways to assess the research, be a basis for choosing starting points, and indicate how to develop appropriate methodologies for the needed research.

Two points need emphasis here. First, although a conceptual framework has been constructed, it is still rudimentary. Further search, and actual research, are needed for the evolution of the framework. Second, even if our conceptual framework did provide an accurate and complete basic analysis of the system from which stems a human's intellectual effectiveness, the explicit nature of future improved systems would be highly affected by (expected) changes in our technology or in our understanding of the human being.

CONCEPTUAL FRAMEWORK

GENERAL

The conceptual framework we seek must orient us toward the real possibilities and problems associated with using modern technology to give direct aid to an individual in comprehending complex situations, isolating the significant factors, and solving problems. To gain this orientation, we examine how individuals achieve their present level of effectiveness, and expect that this examination will reveal possibilities for improvement.

The entire effect of an individual on the world stems essentially from what he can transmit to the world through his limited motor channels. This in turn is based on information received from the outside world through limited sensory channels; on information, drives, and needs generated within him; and on his processing of that information. His processing is of two kinds: that which he is generally conscious of (recognizing patterns, remembering, visualizing, abstracting, deducing, inducing, etc.), and that involving the unconscious processing and mediating of received and self-generated information, and the unconscious mediating of conscious processing itself.

The individual does not use this information and this processing to grapple directly with the sort of complex situation in which we seek to give him help. He uses his innate capabilities in a rather more indirect fashion, since the situation is generally too complex to yield directly to his motor actions, and always too complex to yield comprehensions and solutions from direct sensory inspection and use of basic cognitive capabilities. For instance, an aborigine who possesses all of our basic sensory-mental-motor capabilities, but does not possess our background of indirect knowledge and procedure, cannot organize the proper direct actions necessary to drive a car through traffic, request a book from the library, call a committee meeting to discuss a tentative plan, call someone on the telephone, or compose a letter on the typewriter.

Our culture has evolved means for us to organize the little things we can do with our basic capabilities so that we can derive comprehension from truly complex situations, and accomplish the processes of deriving and implementing problem solutions. The ways in which human capabilities are thus extended are here called augmentation means, and we define four basic classes of them:

Artifacts -- physical objects designed to provide for human comfort, for the manipulation of things or materials, and for the manipulation of symbols.

Language -- the way in which the individual parcels out the picture of his world into the concepts that his mind uses to model that world, and the symbols that he attaches to those concepts and uses in consciously manipulating the concepts ("thinking").

Methodology -- the methods, procedures, strategies, etc., with which an individual organizes his goal-centered (problem-solving) activity.

Training -- the conditioning needed by the human being to bring his skills in using Means 1, 2, and 3 to the point where they are operationally effective.

The system we want to improve can thus be visualized as a trained human being together with his artifacts, language, and methodology. The explicit new system we contemplate will involve as artifacts computers, and computer-controlled information-storage, information-handling, and

information-display devices. The aspects of the conceptual framework that are discussed here are primarily those relating to the human being's ability to make significant use of such equipment in an integrated system.

Pervading all of the augmentation means is a particular structure or organization. While an untrained aborigine cannot drive a car through traffic, because he cannot leap the gap between his cultural background and the kind of world that contains cars and traffic, it is possible to move step by step through an organized training program that will enable him to drive effectively and safely. In other words, the human mind neither learns nor acts by large leaps, but by steps organized or structured so that each one depends upon previous steps.

Although the size of the step a human being can take in comprehension, innovation, or execution is small in comparison to the over-all size of the step needed to solve a complex problem, human beings nevertheless do solve complex problems. It is the augmentation means that serve to break down a large problem in such a way that the human being can walk through it with his little steps, and it is the structure or organization of these little steps or actions that we discuss as process hierarchies.

Every process of thought or action is made up of sub-processes. Let us consider such examples as making a pencil stroke, writing a letter of the alphabet, or making a plan. Quite a few discrete muscle movements are organized into the making of a pencil stroke; similarly, making particular pencil strokes and making a plan for a letter are complex processes in themselves that become sub-processes to the over-all writing of an alphabetic character.

Although every sub-process is a process in its own right, in that it consists of further sub-processes, there seems to be no point here in looking for the ultimate bottom of the process-hierarchical structure. There seems to be no way of telling whether or not the apparent bottoms (processes that cannot be further subdivided) exist in the physical world or in the limitations of human understanding.

In any case, it is not necessary to begin from the bottom in discussing particular process hierarchies. No person uses a process that is completely unique every time he tackles something new. Instead, he begins from a group of basic sensory-mental-motor process capabilities, and adds to these certain of the process capabilities of his artifacts. There are only a finite number of such basic human and artifact capabilities from which to draw. Furthermore, even quite different higher order processes may have in common relatively high-order sub-processes".

When a man writes prose text (a reasonably high-order process), he makes use of many processes as sub-processes that are common to other high-order processes. For example, he makes use of planning, composing, dictating. The process of writing is utilized as a sub-process within many different processes of a still higher order, such as organizing a committee, changing a policy, and so on.

What happens, then, is that each individual develops a certain repertoire of process capabilities from which he selects and adapts those that will compose the processes that he executes. This repertoire is like a tool kit, and just as the mechanic must know what his tools can do and how to use them, so the intellectual worker must know the capabilities of his tools and have good methods, strategies, and rules of thumb for making use of them. All of the process capabilities in the individuals repertoire rest ultimately upon basic capabilities within him or his artifacts, and the entire repertoire represents an inter-knit, hierarchical structure (which we often call the repertoire hierarchy).

We find three general categories of process capabilities within a typical individuals repertoire. There are those that are executed completely within the human integument, which we call explicit-human process capabilities; there are those possessed by artifacts for executing processes without human intervention, which we call explicit-artifact process capabilities; and there are what we call the composite process capabilities, which are derived from hierarchies containing both of the other kinds.

We assume that it is our H-LAM/T system (Human using Language, Artifacts, Methodology, in which he is Trained) that has the capability and that performs the process in any instance of use of this repertoire. Let us look within the process structure for the LAM/T ingredients, to get a better "feel" for our models. Consider the process of writing an important memo. There is a particular concept associated with this process -- that of putting information into a formal package and distributing it to a set of people for a certain kind of consideration -- and the type of information package associated with this concept has been given the special name of memorandum. Already the system language shows the effect of this process -- i.e., a concept and its name.

The memo-writing process may be executed by using a set of process capabilities (in intermixed or repetitive form) such as the following planning, developing subject matter, composing text, producing hard copy, and distributing. There is a definite way in which these sub-processes will be organized that represents part of the system methodology. Each of these sub-processes represents a functional concept that must be a part of the system language if it is to be organized effectively into the human's way of doing things, and the symbolic portrayal of each concept must be such that the human can work with it and remember it.

If the memo is simple, a paragraph or so in length, then the first three processes may well be of the explicit-human type (i.e., it may be planned, developed) and composed within the mind) and the last two of the composite type. If it is a complex memo, involving a good deal of careful planning and development, then all of the sub-processes might well be of the composite type (e.g., at least including the use of pencil and paper artifacts)' and there might be many different applications of some of the process capabilities within the total process (i.e., successive drafts, revised plans).

The set of sub-process capabilities discussed so far, if called upon in proper occasion and sequence, would indeed enable the execution of the memo-writing process. However, the very process of organizing and supervising the utilization of these sub-process capabilities is itself a most important sub-process of the memo-writing process. Hence, the sub- process capabilities as listed would not be complete without the addition of a seventh capability -- what we call the executive capability. This is the capability stemming from habit, strategy, rules of thumb, prejudice, learned method, intuition, unconscious dictates, or combinations thereof, to call upon the appropriate sub-process capabilities with a particular sequence and timing. An executive process (i.e., the exercise of an executive capability) involves such sub-processes as planning, selecting, and supervising, and it is really the executive processes that embody all of the methodology in the H-LAM/T system.

To illustrate the capability-hierarchy features of our conceptual framework, let us consider an artifact innovation appearing directly within the relatively low-order capability for composing and modifying written text, and see how this can affect a (or, for instance, your) hierarchy of capabilities. Suppose you had a new writing machine -- think of it as a high-speed electric typewriter with some special features. You could operate its keyboard to cause it to write text much as you could use a conventional typewriter. But the printing mechanism is more complicated; besides printing a visible character at every stroke, it adds special encoding features by means of invisible selective components in the ink and special shaping of the character.

As an auxiliary device, there is a gadget that is held like a pencil and, instead of a point, has a special sensing mechanism that you can pass over a line of the special printing from your writing machine (or one like it). The signals which this reading stylus sends through the flexible connecting wire to the writing machine are used to determine which characters are being sensed and thus to cause the automatic typing of a duplicate string of characters. An information-storage mechanism in the writing machine permits you to sweep the reading stylus over the characters much faster than the writer can type; the writer will catch up with you when you stop to think about what word or string of words should be duplicated next, or while you reposition the straightedge guide along which you run the stylus.

This writing machine would permit you to use a new process of composing text. For instance, trial drafts could rapidly be composed from re-arranged excerpts of old drafts, together with new words or passages which you stop to type in. Your first draft could represent a free outpouring of thoughts in any order, with the inspection of foregoing thoughts continuously stimulating new considerations and ideas to be entered. If the tangle of thoughts represented by the draft became too complex, you would compile a reordered draft quickly. It would be practical for you to accommodate more complexity in the trails of thought you might build in search of the path that suits your needs.

You can integrate your new ideas more easily, and thus harness your creativity more continuously, if you can quickly and flexibly change your working record. If it is easier to update any part of your working record to accommodate new developments in thought or circumstance, you will find it easier to incorporate more complex procedures in your way of doing things. This will probably allow you to accommodate the extra burden associated with, for instance, keeping and using special files whose contents are both contributed to and utilized by any current work in a flexible manner -- which in turn enables you to devise and use even-more complex procedures to better harness your talents in your particular working situation.

The important thing to appreciate here is that a direct new innovation in one particular capability can have far-reaching effects throughout the rest of your capability hierarchy. A change can propagate up through the capability hierarchy; higher-order capabilities that can utilize the initially changed capability can now reorganize to take special advantage of this change and of the intermediate higher-capability changes. A change can propagate down through the hierarchy as a result of new capabilities at the high level and modification possibilities latent in lower levels. These latent capabilities may previously have been unusable in the hierarchy and become usable because of the new capability at the higher level.

The writing machine and its flexible copying capability would occupy you for a long time if you tried to exhaust the reverberating chain of associated possibilities for making useful innovations within your capability hierarchy. This one innovation could trigger a rather extensive redesign of this hierarchy; your way of accomplishing many of your tasks would change considerably. Indeed this process characterizes the sort of evolution that our intellect-augmentation means have been undergoing since the first human brain appeared.

To our objective of deriving orientation about possibilities for actively pursuing an increase in human intellectual effectiveness, it is important to realize that we must be prepared to pursue such new- possibility chains throughout the entire capability hierarchy (calling for a system approach). It is also important to realize that we must be oriented to the synthesis of new capabilities from

reorganization of other capabilities, both old and new, that exist throughout the hierarchy (calling for a “system-engineering” approach).

THE BASIC PERSPECTIVE

Individuals who operate effectively in our culture have already been considerably “augmented.” Basic human capabilities for sensing stimuli, performing numerous mental operations, and for communicating with the outside world, are put to work in our society within a system -- an H-LAM/T system -- the individual augmented by the language, artifacts, and methodology in which he is trained. Furthermore, we suspect that improving the effectiveness of the individual as he operates in our society should be approached as a system-engineering problem -- that is, the H-LAM/T system should be studied as an interacting whole from a synthesis-oriented approach.

This view of the system as an interacting whole is strongly bolstered by considering the repertoire hierarchy of process capabilities that is structured from the basic ingredients within the H-LAM/T system. The realization that any potential change in language, artifact, or methodology has importance only relative to its use within a process' and that a new process capability appearing anywhere within that hierarchy can make practical a new consideration of latent change possibilities in many other parts of the hierarchy -- possibilities in either language, artifacts, or methodology -- brings out the strong interrelationship of these three augmentation means.

Increasing the effectiveness of the individual's use of his basic capabilities is a problem in redesigning the changeable parts of a system. The system is actively engaged in the continuous processes (among others) of developing comprehension within the individual and of solving problems; both processes are subject to human motivation, purpose, and will. To redesign the system's capability for performing these processes means redesigning all or part of the repertoire hierarchy. To redesign a structure, we must learn as much as we can of what is known about the basic materials and components as they are utilized within the structure; beyond that, we must learn how to view, to measure, to analyze, and to evaluate in terms of the functional whole and its purpose. In this particular case, no existing analytic theory is by itself adequate for the purpose of analyzing and evaluating over-all system performance; pursuit of an improved system thus demands the use of experimental methods.

It need not be just the very sophisticated or formal process capabilities that are added or modified in this redesign. Essentially any of the processes utilized by a representative human today -- the processes that he thinks of when he looks ahead to his day's work -- are composite processes of the sort that involve external composing and manipulating of symbols (text, sketches, diagrams, lists, etc.). Many of the external composing and manipulating (modifying, rearranging) processes serve

such characteristically “human” activities as playing with forms and relationships to ask what develops, cut- and-try multiple-pass development of an idea, or listing items to reflect on and then rearranging and extending them as thoughts develop.

Existing, or near-future, technology could certainly provide our professional problem-solvers with the artifacts they need to have for duplicating and rearranging text before their eyes, quickly and with a minimum of human effort. Even an apparently minor an advance could yield total changes in an individual's repertoire hierarchy that would represent a great increase in over-all effectiveness. Normally the necessary equipment would enter the market slowly; changes from the expected would be small, people would change their ways of doing things a little at a time, and only gradually would their accumulated changes create markets for more radical versions of the equipment. Such an evolutionary process has been typical of the way our repertoire hierarchies have grown and formed.

But an active research effort, aimed at exploring and evaluating possible integrated changes throughout the repertoire hierarchy, could greatly accelerate this evolutionary process. The research effort could guide the product development of new artifacts toward taking long-range meaningful steps; simultaneously competitively minded individuals who would respond to demonstrated methods for achieving greater personal effectiveness would create a market for the more radical equipment innovations. The guided evolutionary process could be expected to be considerably more rapid than the traditional one.

The category of “more radical innovations” includes the digital computer as a tool for the personal use of an individual. Here there is not only promise of great flexibility in the composing and rearranging of text and diagrams before the individual's eyes but also promise of many other process capabilities that can be integrated into the H-LAM/T system's repertoire hierarchy.

DETAILED DISCUSSION OF THE H-LAM/T SYSTEM

The Source of Intelligence

When one looks at a computer system that is doing a very complex job, he sees on the surface a machine that can execute some extremely sophisticated processes. If he is a layman, his concept of what provides this sophisticated capability may endow the machine with a mysterious power to sweep information through perceptive and intelligent synthetic thinking devices. Actually, this sophisticated capability results from a very clever organizational hierarchy so that pursuit of the source of intelligence within this system would take one down through layers of functional and physical organization that become successively more primitive.

To be more specific, we can begin at the top and list the major levels down through which we would pass if we successively decomposed the functional elements of each level in search of the “source of intelligence.” A programmer could take us down through perhaps three levels (depending upon the sophistication of the total process being executed by the computer) perhaps depicting the organization at each level with a flow chart. The first level down would organize functions corresponding to statements in a problem-oriented language (e.g., ALGOL or COBOL), to achieve the desired over-all process. The second level down would organize lesser functions into the processes represented by first-level statements. The third level would perhaps show how the basic machine commands (or rather the processes which they represent) were organized to achieve each of the functions of the second level.

Then a machine designer could take over, and with a block diagram of the computer's organization he could show us (Level 4) how the different hardware units (e.g., random-access storage, arithmetic registers, adder, arithmetic control) are organized to provide the capability of executing sequences of the commands used in Level 3. The logic designer could then give us a tour of Level 5, also using block diagrams, to show us how such hardware elements as pulse gates, flip-flops' and AND, OR, and NOT circuits can be organized into networks giving the functions utilized at Level 4. For Level 6 a circuit engineer could show us diagrams revealing how components such as transistors, resistors, capacitors, and diodes can be organized into modular networks that provide the functions needed for the elements of Level 5.

Device engineers and physicists of different kinds could take us down through more layers. But rather soon we have crossed the boundary between what is man-organized and what is nature-organized, and are ultimately discussing the way in which a given physical phenomenon is derived from the intrinsic organization of sub-atomic particles, with our ability to explain succeeding layers blocked by the exhaustion of our present human comprehension.

If we then ask ourselves where that intelligence is embodied, we are forced to concede that it is elusively distributed throughout a hierarchy of functional processes -- a hierarchy whose foundation extends down into natural processes below the depth of our comprehension. If there is any one thing upon which this 'intelligence depends' it would seem to be organization. The biologists and physiologists use a term “synergism” to designate (from Webster's Unabridged Dictionary, Second Edition) the “...cooperative action of discrete agencies such that the total effect is greater than the sum of the two effects taken independently...” This term seems directly applicable here, where we could say that synergism is our most likely candidate for representing the actual source of intelligence.

Actually, each of the social, life, or physical phenomena we observe about us would seem to derive from a supporting hierarchy of organized functions (or processes), in which the synergistic principle gives increased phenomenological sophistication to each succeeding higher level of organization. In particular, the intelligence of a human being, derived ultimately from the characteristics of individual nerve cells, undoubtedly results from synergism.

Intelligence Amplification

It has been jokingly suggested several times during the course of this study that what we are seeking is an “intelligence amplifier.” (The term is attributed originally to W. Ross Ashby. At first this term was rejected on the grounds that in our view one's only hope was to make a better match between existing human intelligence and the problems to be tackled, rather than in making man more intelligent. But deriving the concepts brought out in the preceding section has shown us that indeed this term does seem applicable to our objective.

Accepting the term “intelligence amplification” does not imply any attempt to increase native human intelligence. The term “intelligence amplification” seems applicable to our goal of augmenting the human intellect in that the entity to be produced will exhibit more of what can be called intelligence than an unaided human could; we will have amplified the intelligence of the human by organizing his intellectual capabilities into higher levels of synergistic structuring. What possesses the amplified intelligence is the resulting H-LAM/T system, in which the LAM/T augmentation means represent the amplifier of the human's intelligence.

In amplifying our intelligence, we are applying the principle of synergistic structuring that was followed by natural evolution in developing the basic human capabilities. What we have done in the development of our augmentation means is to construct a superstructure that is a synthetic extension of the natural structure upon which it is built. In a very real sense, as represented by the steady evolution of our augmentation means, the development of “artificial intelligence” has been going on for centuries.

Two-Domain System

The human and the artifacts are the only physical components in the H-LAM/T system. It is upon their capabilities that the ultimate capability of the system will depend. This was implied in the earlier statement that every composite process of the system decomposes ultimately into explicit-human and explicit-artifact processes. There are thus two separate domains of activity within the H-LAM/T system: that represented by the human, in which all explicit-human processes occur; and that represented by the artifacts, in which all explicit-artifact processes occur. In any composite process, there is cooperative interaction between the two domains, requiring interchange of energy

(much of it for information exchange purposes only). Figure 1 depicts this two domain concept and embodies other concepts discussed below.

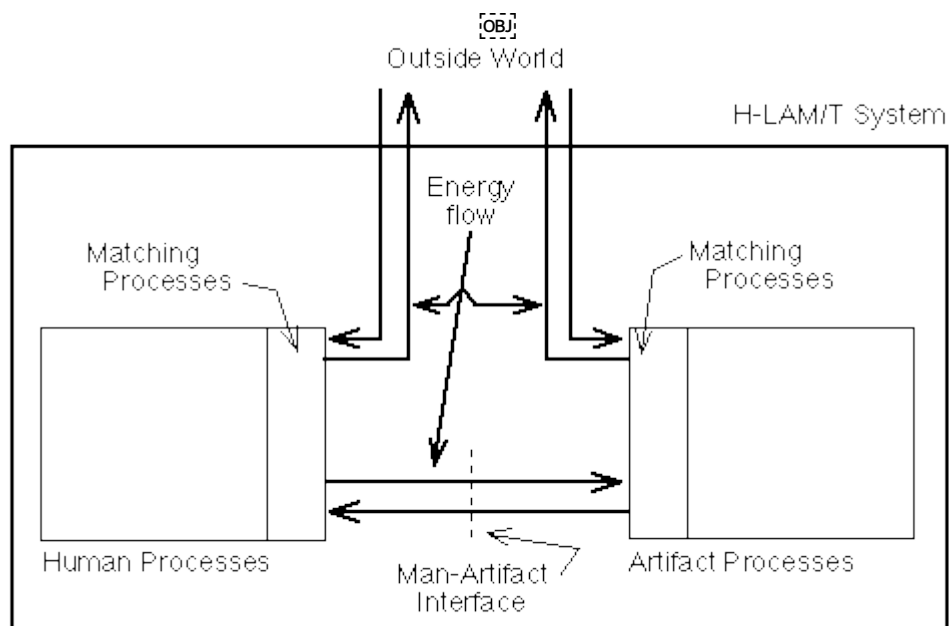


Fig.1: The Two Sides of the H-LAM/T System

Where a complex machine represents the principal artifact with which a human being cooperates, the term “man-machine interface” has been used for some years to represent the boundary across which energy is exchanged between the two domains. However, the “man-artifact interface” has existed for centuries, ever since humans began using artifacts and executing composite processes.

Exchange across this “interface” occurs when an explicit-human process is coupled to an explicit-artifact process. Quite often these coupled processes are designed for just this exchange purpose, to provide a functional match between other explicit-human and explicit-artifact processes buried within their respective domains that do the more significant things. For instance, the finger and hand motions (explicit human processes) activate key-linkage motions in the typewriter (couple to explicit-artifact processes). But these are only part of the matching processes between the deeper human processes that direct a given word to be typed and the deeper artifact processes that actually imprint the ink marks on the paper.

The outside world interacts with our H-LAM/T system by the exchange of energy with either the individual or his artifact. Again, special processes are often designed to accommodate this exchange. However, the direct concern of our present study lies within the system, with the internal processes that are and can be significantly involved in the effectiveness of the system in developing the human's comprehension and pursuing the human's goals.

Concepts, Symbols, and a Hypothesis

Before we pursue further direct discussion of the H-LAM/T system, let us examine some background material. Consider the following historical progression in the development of our intellectual capabilities:

Concept Manipulation -- Humans rose above the lower forms of life by evolving the biological capability for developing abstractions and concepts. They could manipulate these concepts within their minds to a certain extent, and think about situations in the abstract. Their mental capabilities allowed them to develop general concepts from specific instances, predict specific instances from general concepts, associate concepts, remember them, etc. We speak here of concepts in their raw, un-verbalized form. For example, a person letting a door swing shut behind him suddenly visualizes the person who follows him carrying a cup of hot coffee and some sticky pastries. Of all the aspects of the pending event, the spilling of the coffee and the squashing of the pastry somehow are abstracted immediately, and associated with a concept of personal responsibility and a dislike for these consequences. But a solution comes to mind immediately as an image of a quick stop and an arm stab back toward the door, with motion and timing that could prevent the collision, and the solution is accepted and enacted. With only non-symbolic concept manipulation, we could probably build primitive shelter, evolve strategies of war and hunt, play games, and make practical jokes. But further powers of intellectual effectiveness are implicit in this stage of biological evolution (the same stage we are in today).

Symbol Manipulation -- Humans made another great step forward when they learned to represent particular concepts in their minds with specific symbols. Here we temporarily disregard communicative speech and writing, and consider only the direct value to the individual of being able to do his heavy thinking by mentally manipulating symbols instead of the more unwieldy concepts which they represent. Consider, for instance, the mental difficulty involved in herding twenty- seven sheep if, instead of remembering one cardinal number and occasionally counting, we had to remember what each sheep looked like, so that if the flock seemed too small we could visualize each one and check whether or not it was there.

Manual, External, Symbol Manipulation -- Another significant step toward harnessing the biologically evolved mental capabilities in pursuit of comprehension and problem solutions came with the development of the means for externalizing some of the symbol-manipulation activity, particularly in graphical representation. This supplemented the individual's memory and ability to visualize. (We are not concerned here with the value derived from human cooperation made possible by speech and writing, both forms of external symbol manipulation. We speak of the manual means of making graphical representations of symbols -- a stick and sand, pencil and paper

and eraser, straight edge or compass, and so on.) It is principally this kind of means for external symbol manipulation that has been associated with the evolution of the individual's present way of doing his concept manipulation (thinking).

It is undoubtedly true that concepts which people found useful ended up being symbolized in their language, and hence that the evolution of language was affected by the concepts the people developed and used. However, Korzybski and Whorf (among others) have argued that the language we use affects our thinking to a considerable extent. They say that a lack of words for some types of concepts makes it hard to express those concepts, and thus decreases the likelihood that we will learn much about them. If this is so, then once a language has begun to grow and be used, it would seem reasonable to suspect that the language also affects the evolution of the new concepts to be expressed in that language.

Apparently there are counter-arguments to this; e.g., if a concept needs to be used often but its expression is difficult, then the language will evolve to ease the situation. However, the studies of the past decade into what are called "self-organizing" systems seem to be revealing that subtle relationships among its interacting elements can significantly influence the course of evolution of such a system. If this is true, and if language is (as it seems to be) a part of a self organizing system, then it seems probable that the state of a language at a given time strongly affects its own evolution to a succeeding state.

For our conceptual framework, we tend to favor the view that a language does exert a force in its own evolution. We observe that the shift over the last few centuries in matters that are of daily concern to the individual has necessarily been forced into the framework of the language existing at the time, with alterations generally limited to new uses for old words, or the coining of new words. The English language since Shakespeare has undergone no alteration comparable to the alteration in the cultural environment; if it had, Shakespeare would no longer be accessible to us. Under such evolutionary conditions, it would seem unlikely that the language we now use provides the best possible service to our minds in pursuing comprehension and solving problems. It seems very likely that a more useful language form can be devised.

The Whorfian hypothesis states that the world view of a culture is limited by the structure of the language which that culture uses. But there seems to be another factor to consider in the evolution of language and human reasoning ability. We offer the following hypothesis, which is related to the Whorfian hypothesis: Both the language used by a culture, and the capability for effective intellectual activity are directly affected during their evolution by the means by which individuals control the external manipulation of symbols. (For identification, we will refer to this as the Neo-Whorfian hypothesis.)

If the Neo-Whorfian hypothesis could be proved readily, and if we could see how our means of externally manipulating symbols influence both our language and our way of thinking, then we would have a valuable instrument for studying human-augmentation possibilities. For the sake of discussion, let us assume the Neo-Whorfian hypothesis to be true, and see what relevant deductions can be made.

If the means evolved for an individual's external manipulation of his thinking-aid symbols indeed directly affect the way in which he thinks, then the original Whorfian hypothesis would offer an added effect. The direct effect of the external-symbol-manipulation means upon language would produce an indirect effect upon the way of thinking via the Whorfian-hypothesis linkage. There would then be two ways for the manner in which our external symbol manipulation was done to affect our thinking.

One way of viewing the H-LAM/T system changes that we contemplate -- specifically, integrating the capabilities of a digital computer into the intellectual activity of individual humans -- is that we are introducing new and extremely advanced means for externally manipulating symbols. We then want to determine the useful modifications in the language and in the way of thinking that could result. This suggests a fourth stage to the evolution of our individual-human intellectual capability:

Automated external symbol manipulation -- In this stage, symbols with which the human represents the concepts he is manipulating can be arranged before his eyes, moved, stored, recalled, operated upon according to extremely complex rules -- all in very rapid response to a minimum amount of information supplied by the human, by means of special cooperative technological devices. In the limit of what we might now imagine, this could be a computer, with which we could communicate rapidly and easily, coupled to a three-dimensional color display within which it could construct extremely sophisticated images -- with the computer being able to execute a wide variety of processes upon parts or all of these images in automatic response to human direction. The displays and processes could provide helpful services -- we could imagine both simple and exotic varieties -- and could involve concepts that we have never yet imagined (as the pre-graphic thinker of Stage 2 would be unable to predict the bar graph, the process of long division, or a card file system).

These hypotheses imply great richness in the new evolutionary spaces opened by progressing from Stage 3 to Stage 4. We would like to study the hypotheses further, examining their possible manifestations in our experience, ways of demonstrating their validity, and possible deductions relative to going to Stage 4.

In search of some simple ways to determine what the Neo-Whorfian hypothesis might imply, we could imagine some relatively straightforward means of increasing our external symbol-

manipulation capability and try to picture the consequent changes that could evolve in our language and methods of thinking. Actually, it turned out to be simpler to invert the problem and consider a change that would reduce our capability for external symbol manipulation. This allowed an empirical approach which proved both simple and effective. We thus performed the following experiment.

Brains of power equal to ours could have evolved in an environment where the combination of artifact materials and muscle strengths were so scaled that the neatest scribing tool (equivalent to a pencil, possibly had a shape and mass as manageable as a brick would be to us—assuming that our muscles were not specially conditioned to deal with it. We fastened a pencil to a brick and experimented. Figure 2 shows the results, compared with typewriting and ordinary pencil writing. With the brick pencil, we are slower and less precise. If we want to hurry the writing, we have to make it larger. Also, writing the passage twice with the brick-pencil tires the untrained hand and arm.

How would our civilization have matured if this had been the only manual means for us to use in graphical manipulation of symbols? For one thing, the record keeping that enables the organization of commerce and government would probably have taken a form so different from what we know that our social structure would undoubtedly have evolved differently. Also, the effort in doing calculations and writing down extensive and carefully reasoned argument would dampen individual experimentation with sophisticated new concepts, to lower the rate of learning and the rate of useful output, and perhaps to discourage a good many people from even working at extending understanding. The concepts that would evolve within our culture would thus be different, and very likely the symbology to represent them would be different -- much more economical of motion in their writing. It thus seems very likely that our thoughts and our language would be rather directly affected by the particular means used by our culture for externally manipulating symbols, which gives little intuitive substantiation to our Neo-Whorfian hypothesis.

Augmentation is fundamentally a matter of organization.
(typewriter, 7 seconds).

Augmentation is fundamentally a matter of
organization. (cursive script, 20 seconds).

Augmentation is fundamentally a matter
of organization. ("de-augmented" cursive script, 65 seconds).

Augmentation
is fundamentally
a matter of

[de-augmented cursive script, large size--42 seconds to complete whole passage (completed on separate sheet)].

Fig. 2: Experimental Results of Tying a Brick to a Pencil to "De-Augment" the Individual

To reflect further upon the implications of this hypothesis, the following hypothetical artifact development can be considered, representing a different type of external symbol manipulation that could have had considerable effect. Suppose that our young technology of a few generations ago

had developed an artifact that was essentially a high speed, semi-automatic table-lookup device -- cheap enough for almost everyone to afford and small and light enough to be carried on the person. Assume that the individual cartridges sold by manufacturers (publishers) contained the look-up information, that one cartridge could hold the equivalent of an unabridged dictionary, and that a one-paragraph definition could always be located and displayed on the face of the device by the average practised individual in less than three seconds. The fortunes of technological invention, commercial interest, and public acceptance just might have evolved something like this.

If it were so very easy to look things up, how would our vocabulary develop, how would our habits of exploring the intellectual domains of others shift, how might the sophistication of practical organization mature (if each person can so quickly and easily look up applicable rules), how would our education system change to take advantage of this new external symbol-manipulation capability of students and teachers (and administrators)?

The significance to our study of the discussion in this section lies in the perspective it gives to the ways in which human intellectual effectiveness can be affected by the particular means used by individuals for their external symbol manipulation. It seems reasonable to consider the development of automated external symbol manipulation means as a next stage in the evolution of our intellectual power.

Capability Repertoire Hierarchy

The concept of our H-LAM/T system possessing a repertoire of capabilities that is structured in the form of a hierarchy is most useful in our study. We shall use it in the following to tie together a number of considerations and concepts. There are two points of focus in considering the design of new repertoire hierarchies: the materials with which we have to work, and the principles by which new capability is constructed from these basic materials.

Basic Capabilities

”Materials” in this context are those capabilities in the human and in the artifact domains from which all other capabilities in the repertoire hierarchy must be constructed. Each such basic capability represents a type of functional component with which the system can be built, and a thorough job of redesigning the system calls for making an inventory of the basic capabilities available. Because we are exploring for perspective, and not yet recommending research activities, we are free to discuss and define in more detail what we mean by “basic capability”, without regard to the amount of research involved in making an actual inventory.

The two domains, human and artifact, can be explored separately for their basic capabilities, In each we can isolate two classes of basic capability; these classes are distinguished according to whether or not the capability has been put to use within out augmentation means. The first class (those in use) can be found in a methodical manner by analyzing present capability hierarchies. For example, select a given capability, at any level in the hierarchy, and ask yourself if it can be usefully changed by any means that can be given consideration in the augmentation research contemplated, If it can, then it is not basic but it can be decomposed into an eventual set of basic capabilities. As you proceed down through the hierarchy, you will begin to encounter capabilities that cannot be usefully changed, and these will make up your inventory of basic capabilities. Ultimately, every such recursive decomposition of a given capability in the hierarchy will find every one of its branching paths terminated by basic capabilities. Beginning such decomposition search with different capabilities in the hierarchy will eventually uncover all of those basic capabilities used within that hierarchy or augmentation system. Many of the branching paths in the decomposition of a given higher-order capability will terminate in the same basic capability, since a given basic capability will often be used within many different higher-order capabilities.

Determining the class of basic capabilities not already utilized within existing augmentation systems requires a different exploration method. Examples of this method occur in technological research, where analytically oriented researchers search for new understandings of phenomena that can add to the research engineers list of things to be used in the synthesis of better artifacts.

Before this inventorying task can be pursued in any specific instance, some criteria must be established as to what possible changes within the H-LAM/T system can be given serious consideration. For instance, some research situations might have to disallow changes which require extensive retraining, or which require undignified behavior by the human. Other situations might admit changes requiring years of special training, very expensive equipment, or the use of special drugs.

The capability for performing a certain finger action, for example, may not be basic in our sense of the word. Being able to extend the finger a certain distance would be basic but the strength and speed of a particular finger motion and its coordination with higher actions generally are usefully changeable and therefore do not represent basic capabilities. What would be basic in this case would perhaps be the processes whereby strength could be increased and coordinated movement patterns learned, as well as the basic movement range established by the mechanical-limit loci of the muscle-tendon-bone system. Similar capability breakdowns will occur for sensory and cognitive capabilities.

Structure Types

General

The fundamental principle used in building sophisticated capabilities from the basic capabilities is structuring -- the special type of structuring (which we have termed synergetic) in which the organization of a group of elements produces an effect greater than the mere addition of their individual effects. Perhaps "purposeful" structuring (or organization) would serve us as well, but since we aren't sure yet how the structuring concept must mature for our needs, we shall tentatively stick with the special modifier, "synergetic." We are developing a growing awareness of the significant and pervasive nature of such structure within every physical and conceptual thing we inspect, where the hierarchical form seems almost universally present as stemming from successive levels of such organization.

The fundamental entities that are being structured in each and every case seems to be what we could call processes, where the most basic of physical processes (involving fields, charges, and momenta associated with the dynamics of fundamental particles) appear to be the hierarchical base. There are dynamic electro-optical-mechanical processes associated with the function of our artifacts (as well as metabolic, sensory, motor) and cognitive processes of the human, which we find to be relatively fundamental components within the structure of our H-LAM/T system -- and each of these seems truly to be ultimately based (to our degree of understanding) upon the above mentioned basic physical processes. The elements that are organized to give fixed structural form to our physical objects -- e.g., the "element" of tensile strength of a material-are also derived from what we could call synergetic structuring of the most basic physical processes.

But at the level of the capability hierarchy where we wish to work, it seems useful to us to distinguish several different types of structuring -- even though each type is fundamentally a structuring of the basic physical processes. Tentatively we have isolated five such types -- although we are not sure how many we shall ultimately want to use in considering the problem of augmenting the human intellect, nor how we might divide and subdivide these different manifestations of physical-process structuring. We use the terms "mental structuring", "concept structuring", "symbol structuring", "process structuring," and "physical structuring."

Mental Structuring

Mental structuring is what we call the internal organization of conscious and unconscious mental images, associations, or concepts (or whatever it is that is organized within the human mind) that somehow manages to provide the human with understanding and the basis for such as judgment, intuition, inference, and meaningful action with respect to his environment. There is a term used in psychology, cognitive structure, which so far seems to represent just what we want for our concept

of mental structure, but we will not adopt it until we become more sure of what the accepted psychological meaning is and of what we want for our conceptual framework.

For our present purpose, it is irrelevant to worry over what the fundamental mental “things” being structured are, or what mechanisms are accomplishing the structuring or making use of what has been structured. We feel reasonably safe in assuming that learning involves some kind of meaningful organization within the brain, and that whatever is so organized or structured represents the operating model of the individual's universe to the mental mechanisms that derive his behavior. And further, our assumption is that when the human in our H/LAM system makes the key decision or action that leads to the solution of a complex problem, it will stem from the state of his mental structure at that time. In this view then, the basic purpose of the system's activity on that problem up to that point has been to develop his mental structure to the state from which the mental mechanisms could derive the key action.

Our school systems attest that there are specific experiences that can be given to a human that will result in development of his mental structure to the point where the behavior derived there from by his mental mechanisms shows us that he has gained new comprehension -- in other words, we can do a certain amount from outside the human toward developing his mental structure. Independent students and researchers also attest that internally directed behavior on the part of an individual can directly aid his structure-building process.

We don't know whether a mental structure is developed in a manner analogous to (a) development of a garden, where one provides a good environment, plants the seeds, keeps competing weeds and injurious pests out, but otherwise has to let natural processes take their course, or to (b) development of a basketball team, where much exercise of skills, patterns, and strategies must be provided so that natural processes can slowly knit together an integration, or to (c) development of a machine, where carefully formed elements are assembled in a precise, planned manner so that natural phenomena can immediately yield planned function. We don't know the processes, but we can and have developed empirical relationships between the experiences given a human and the associated manifestations of developing comprehension and capability, and we see the near-future course of the research toward augmenting the human's intellect as depending entirely upon empirical findings (past and future) for the development of better means to serve the development and use of mental structuring in the human.

We don't mean to imply by this that we renounce theories of mental processes. What we mean to emphasize is that pursuit of our objective need not wait upon the understanding of the mental processes that accomplish (what we call) mental structuring and that derive behavior therefrom. It would be to ignore the emphases of our own conceptual framework not to make fullest use of any

theory that provided a working explanation for a group of empirical data. What's more, our entire conceptual framework represents the first pass at a “theoretical model with which to organize our thinking and action.”

Concept Structuring

Within our framework we have developed the working assumption that the manner in which we seem to be able to provide experiences that favor the development of our mental structures is based upon concepts as a “medium of exchange.” We view a concept as a tool that can be grasped and used by the mental mechanisms, that can be composed, interpreted, and used by the natural mental substances and processes. The grasping and handling done by these mechanisms can often be facilitated if the concept is given an explicit “handle” in the form of a representative symbol. Somehow the mental mechanisms can learn to manipulate images (or something) of symbols in a meaningful way and remain calmly confident that the associated conceptual manipulations are within call.

Concepts seem to be structurable, in that a new concept can be composed of an organization of established concepts. For present purposes, we can view a concept structure as something which we might try to develop on paper for ourselves or work with by conscious thought processes, or as something which we try to communicate to one another in serious discussion. We assume that, for a given unit of comprehension to be imparted, there is a concept structure (which can be consciously developed and displayed) that can be presented to an individual in such a way that it is mapped into a corresponding mental structure which provides the basis for that individual's “comprehending” behavior. Our working assumption also considers that some concept structures would be better for this purpose than others, in that they would be more easily mapped by the individual into workable mental structures, or in that the resulting mental structures enable a higher degree of comprehension and better solutions to problems, or both.

A concept structure often grows as part of a cultural evolution -- either on a large scale within a large segment of society, or on a small scale within the activity domain of an individual. But it is also something that can be directly designed or modified, and a basic hypothesis of our study is that better concept structures can be developed -- structures that when mapped into a human's mental structure will significantly improve his capability to comprehend and to find solutions within his complex-problem situations.

A natural language provides its user with a ready made structure of concepts that establishes a basic mental structure, and that allows relatively flexible, general-purpose concept structuring. Our concept of language as one of the basic means for augmenting the human intellect embraces all of the concept structuring which the human may make use of.

Symbol Structuring

The other important part of our “language” is the way in which concepts are represented -- the symbols and symbol structures. Words structured into phrases, sentences, paragraphs, monographs - - charts, lists, diagrams, tables, etc. A given structure of concepts can be represented by any of an infinite number of different symbol structures, some of which would be much better than others for enabling the human perceptual and cognitive apparatus to search out and comprehend the conceptual matter of significance and/or interest to the human. For instance, a concept structure involving many numerical data would generally be much better represented with Arabic rather than Roman numerals and quite likely a graphic structure would be better than a tabular structure.

But it is not only the form of a symbol structure that is important. A problem solver is involved in a stream of conceptual activity whose course serves his mental needs of the moment. The sequence and nature of these needs are quite variable, and yet for each need he may benefit significantly from a form of symbol structuring that is uniquely efficient for that need.

Therefore, besides the forms of symbol structures that can be constructed and portrayed, we are very much concerned with the speed and flexibility with which one form can be transformed into another, and with which new material can be located and portrayed.

We are generally used to thinking of our symbol structures as a pattern of marks on a sheet of paper. When we want a different symbol-structure view, we think of shifting our point of attention on the sheet, or moving a new sheet into position. But another kind of view might be obtained by extracting and ordering all statements in the local text that bear upon consideration A of the argument -- or by replacing all occurrences of specified esoteric words by one's own definitions. This sort of “view generation” becomes quite feasible with a computer-controlled display system, and represents a very significant capability to build upon.

With a computer manipulating our symbols and generating their portrayals to us on a display, we no longer need think of our looking at the symbol structure which is stored -- as we think of looking at the symbol structures stored in notebooks, memos, and books. What the computer actually stores need be none of our concern, assuming that it can portray symbol structures to us that are consistent with the form in which we think our information is structured.

A given concept structure can be represented with a symbol structure that is completely compatible with the computer's internal way of handling symbols, with all sorts of characteristics and relationships given explicit identifications that the user may never directly see. In fact, this structuring has immensely greater potential for accurately mapping a complex concept structure than does a structure an individual would find it practical to construct or use on paper.

The computer can transform back and forth between the two-dimensional portrayal on the screen, of some limited view of the total structure, and the aspect of the n-dimensional internal image that represents this “view”. If the human adds to or modifies such a “view,” the computer integrates the change into the internal-image symbol structure (in terms of the computer's favored symbols and structuring) and thereby automatically detects a certain proportion of his possible conceptual inconsistencies.

Thus, inside this instrument (the computer) there is an internal-image, computer-symbol structure whose convolutions and multi-dimensionality we can learn to shape to represent to hitherto unattainable accuracy the concept structure we might be building or working with. This internal structure may have a form that is nearly incomprehensible to the direct inspection of a human (except in minute chunks).

But let the human specify to the instrument his particular conceptual need of the moment, relative to this internal image. Without disrupting its own internal reference structure in the slightest, the computer will effectively stretch, bend, fold, extract, and cut as it may need in order to assemble an internal substructure that is its response, structured in its own internal way. With the set of standard translation rules appropriate to the situation, it portrays to the human via its display a symbol structure designed for his quick and accurate perception and comprehension of the conceptual matter pertinent to this internally composed substructure.

No longer does the human work on stiff and limited symbol structures, where much of the conceptual content can only be implicitly designated in an indirect and distributed fashion. These new ways of working are basically available with today's technology -- we have but to free ourselves from some of our limiting views and begin experimenting with compatible sets of structure forms and processes for human concepts, human symbols, and machine symbols.

Process Structuring

Essentially everything that goes on within the H-LAM/T system and that is of direct interest here involves the manipulation of concept and symbol structures in service to the mental structure. Therefore, the processes within the H-LAM/T system that we are most interested in developing are those that provide for the manipulation of all three types of structure. This brings us to the fourth category of structuring, process structuring.

As we are currently using it, the term includes the organization, study, modification, and execution of processes and process structures. Whereas concept structuring and symbol structuring together represent the language component of our augmentation means, process structuring represents the methodology component (plus a little more, actually). There has been enough previous discussion

of process structures that we need not describe the notion here, beyond perhaps an example or two. The individual processes (or actions) of my hands and fingers have to be cooperatively organized if the typewriter is to do my bidding. My successive actions throughout my working day are meant to cooperate toward a certain over-all professional goal.

Many of the process structures are applied to the task of organizing, executing, supervising, and evaluating other process structures. Many of them are applied to the formation and manipulation of symbol structures (the purpose of which will often be to support the conceptual labor involved in process structuring).

Physical Structuring

Physical structuring, the last of the five types which we currently use in our conceptual framework, is nearly self-explanatory. It pretty well represents the artifact component of our augmentation means, insofar as their actual physical construction is concerned.

Interdependence and Regeneration

A very important feature to be noted from the discussion in this section bears upon the interdependence among the various types of structuring which are involved in the H-LAM/T system, where the capability for doing each type of structuring is dependent upon the capability for doing one or more of the other types of structuring. (Assuming that the physical structuring of the system remains basically unchanged during the system's operation, we exclude its dependence upon other factors in this discussion.)

This interdependence actually has a cyclic, regenerative nature to it which is very significant to us. We have seen how the capability for mental structuring is finally dependent, down the chain, upon the process structuring (human, artifact, composite) that enables symbol-structure manipulation. But it also is evident that the process structuring is dependent not only upon basic human and artifact process capabilities, but upon the ability of the human to learn how to execute processes -- and no less important, upon the ability of the human to select, organize, and modify processes from his repertoire to structure a higher-order process that he can execute. Thus, a capability for structuring and executing processes is partially dependent upon the human's mental structuring, which in turn is partially dependent upon his process structuring (through concept and symbol structuring), which is partially dependent upon his mental structuring, etc.

All of this means that a significant improvement in symbol-structure manipulation through better process structuring (initially perhaps through much better artifacts) should enable us to develop improvements in concept and mental-structure manipulations that can in turn enable us to organize and execute symbol-manipulation processes of increased power. To most people who initially

consider the possibilities for computer-like devices augmenting the human intellect, it is only the one-pass improvement that comes to mind, which presents a picture that is relatively barren compared to that which emerges when one considers this regenerative interaction.

We can confidently expect the development of much more powerful concepts pertaining to the manner in which symbol structures can be manipulated and portrayed, and correspondingly more complex manipulation processes that in the first pass would have been beyond the human's power to organize and execute without the better symbol, concept, and mental structuring which his augmented system provided him. These new concepts and processes, beyond our present capabilities to use and thus never developed, will provide a tremendous increased-capability payoff in the future development of our augmentation means.

Roles and Levels

In the repertoire hierarchy of capabilities possessed by the H-LAM/T system, the human contributes many types of capability that represent a wide variety of roles. At one time or another he will be the policy maker, the goal setter, the performance supervisor, the work scheduler, the professional specialist, the clerk, the janitor, the entrepreneur, and the proprietor (or at least a major stockholder) of the system. In the midst of some complex process, in fact, he may well be playing several roles concurrently -- or at least have the responsibility of the roles. For instance, usually he must be aware of his progress toward a goal (supervisor), he must be alert to the possibilities for changing the goal (policy maker, planner), and he must keep records for these and other roles (clerk).

Consider a given capability (Capability 1) at some level in the repertoire hierarchy. There seems to be a sort of standard grouping of lower-order capabilities from which this is composed, and these exist in two classes -- what we might call the executive class and what we might call the direct-contributive class. In the executive class of capabilities we find those used for comprehending, planning, and executing the process represented by Capability 1. In the direct-contributive class we find the capabilities organized by the executive class toward the direct realization of Capability 1. For example, when my telephone rings, I execute the direct-contributive processes of picking up the receiver and saying hello. It was the executive processes that comprehended the situation, directed a lower-order executive-process that the receiver be picked up and, when the receiver was in place (first process accomplished), directed the next process, the saying hello. That represents the composition of my capability for answering the phone.

For a low-level capability, such as that of writing a word with a pencil, both the executive and the direct-contributive subprocesses during actual execution would be automatic. This type of automatic capability need only be summoned by a higher executive process in order for trained automatic responses to execute it.

At a little higher level of capability, more of the conscious conceptual and executive capabilities become involved. To call someone on the telephone, I must consciously comprehend the need for this process and how I can execute it, I must consciously pick up the directory and search for the name and telephone number, and I must consciously direct the dialing of the number.

At a still higher level of capability, the executive capabilities must have a degree of power that unaided mental capabilities cannot provide. In such a case, one might make a list of steps and check each item off as it is executed. For an even more complex process, comprehending the particular situation in which it is to be executed, even before beginning to plan the execution, may take months of labor and a very complex organization of the system's capabilities.

Imagining a process as complex as the last example brings us to the realization that at any particular moment the H-LAM/T system may be in the middle of executing a great number of processes. Assume that the human is in the middle of the process of making a telephone call. That telephone call is a subprocess in the middle of the process of calling a committee meeting. But calling a committee meeting is a subprocess in the middle of the process of determining a budgetary policy, which is in turn but a subprocess in the middle of the process of estimating manpower needs, and so on.

Not only does the human need to play various roles (sometimes concurrently) in the execution of any given process, but he is playing these roles for the many concurrent processes that are being executed at different levels. This situation is typical for any of us engaged in reasonably demanding types of professional pursuits, and yet we have never received explicit training in optimum ways of carrying out any but a very few of the roles at a very few of the levels. A well-designed H-LAM/T system would provide explicit and effective concepts, terms, equipment, and methods for all these roles, and for their dynamic coordination.

Model of Executive Superstructure

It is the repertoire hierarchy of process capabilities upon which the ultimate capability of the H-LAM/T system rests. This repertoire hierarchy is rather like a mountain of white-collar talent that sits atop and controls the talents of the workers. We can illustrate this executive superstructure by considering it as though it were a network of contractors and subcontractors in which each capability in the repertoire hierarchy is represented by an independent contractor whose mode of operation is to do the planning, make up specifications, subcontract the actual work, and supervise the performance of his subcontractors. This means that each subcontractor does the same thing in his turn. At the bottom of this hierarchy are those independent contractors who do actual “production work.”

If by some magical process the production workers could still know just what to do and when to do it even though the superstructure of contractors was removed from above them, no one would know the difference. The executive superstructure is necessary because humans do not operate by magic, but even a necessary superstructure is a burden. We can readily recognize that there are many ways to organize and manage such a superstructure, resulting in vastly different degrees of efficiency in the application of the workers' talents.

Suppose that the activity of the production workers was of the same nature as the activity of the different contractors, and that this activity consisted of gaining comprehension and solving problems. And suppose that there was only so much applicable talent available to the total system. The question now becomes how to distribute that talent between superstructure and workers to get the most total production. The efficiency of organization within the superstructure is now doubly important so that a minimum of talent in the superstructure produces a maximum of organizational efficiency in directing the productivity of the remaining talent.

In the situation where talent is limited, we find a close parallel to our H-LAM/T system in its pursuit of comprehension and problem solutions. We obtain an even closer parallel if we say that the thinking, planning, supervising, record keeping, etc., for each contractor is actually done by a single individual for the whole superstructure, time-sharing his attention and talents over these many tasks. Today this individual cannot be depended upon to have any special training for many of these roles; he is likely to have learned them by cut and try and by indirect imitation.

A complex process is often executed by the H-LAM/T system in a multi-pass fashion (i.e., cut and try). In really complex situations, comprehension and problem solutions do not stand waiting at the end of a straightforward path; instead, possibilities open up and plans shift as comprehension grows. In the model using a network of contractors, this type of procedure would entail a great deal of extra work within the superstructure -- each contractor involved in the process would have the specifications upon which he bid continually changed, and would continually have to respond to the changes by restudying the situation, changing his plans, changing the specifications to his subcontractors, and changing his records. This is a terrific additional burden, but it allows a freedom of action that has tremendous importance to the effectiveness the system exhibits to the outside world.

We could expect significant gains from automating the H-LAM/T system if a computer could do nothing more than increase the effectiveness of the executive processes. More human time, energy, and productive thought could be allocated to direct-contributive processes, which would be coordinated in a more sophisticated, flexible and efficient manner. But there is every reason to believe that the possibilities for much-improved symbol and process structuring that would stem

from this automation will directly provide improvements in both the executive and direct-contributive processes in the system.

Flexibility in the Executive Role

The executive superstructure is a necessary component in the H-LAM/T system, and there is finite human capability which must be divided between executive and direct-contributive activities. An important aspect of the multi-role activity of the human in the system is the development and manipulation of the symbol structures associated with both his direct-contributive roles and his executive roles.

When the system encounters a complex situation in which comprehension and problem solutions are being pursued, the direct-contributive roles require the development of symbol structures that portray the concepts involved within the situation. But executive roles in a complex problem situation also require conceptual activity -- e.g., comprehension, selection, supervision -- that can benefit from well-designed symbol structures and fast, flexible means for manipulating and displaying them. For complex processes, the executive problem posed to the human (of gaining the necessary comprehension and making a good plan) may be tougher than the problem he faced in the role of direct-contributive worker. If the flexibility desired for the process hierarchies (to make room for human cut-and-try methods) is not to be degraded or abandoned, the executive activity will have to be provided with fast and flexible symbol-structuring techniques.

The means available to humans today for developing and manipulating these symbol structures are both laborious and inflexible. It is hard enough to develop an initial structure of diagrams and text, but the amount of effort required to make changes is often prohibitively great; one settles for inflexibility. Also, the kind of generous flexibility that would be truly helpful calls for added symbol structuring just to keep track of the trials, branches, and reasoning thereto that are involved in the development of the subject structure; our present symbol-manipulation means would very soon bog down completely among the complexities that are involved in being more than just a little bit flexible.

We find that the humans in our H-LAM/T systems are essential working continuously within a symbol structure of some sort, shifting their attention from one structure to another as they guide and execute the processes that ultimately provide them with the comprehension and the problem solutions that they seek. This view increases our respect for the essential importance of the basic capability of composing and modifying efficient symbol structures. Such a capability depends heavily upon the particular concepts that are isolated and manipulated as entities, upon the symbology used to represent them, upon the artifacts that help to manipulate and display the symbols, and upon the methodology for developing and using symbol structures. In other words,

this capability depends heavily upon proper language, artifacts, and methodology, our basic augmentation means.

When the course of action must respond to new comprehension, new insights and new intuitive flashes of possible explanations or solutions, it will not be an orderly process. Existing means of composing and working with symbol structures penalize disorderly processes very heavily, and it is part of the real promise in the automated H-LAM/T systems of tomorrow that the human can have the freedom and power of disorderly processes.

Compound Effects

Since many processes in many levels of the hierarchy are involved in the execution of a single higher-level process of the system, any factor that influences process execution in general will have a highly compounded total effect upon the system's performance. There are several such factors which merit special attention.

Basic human cognitive powers, such as memory intelligence, or pattern perception can have such a compounded effect. The augmentation means employed today have generally evolved among large statistical populations, and no attempt has been made to fit them to individual needs and abilities. Each individual tends to evolve his own variations, but there is not enough mutation and selection activity, nor enough selection feedback, to permit very significant changes. A good, automated H-LAM/T system should provide the opportunity for a significant adaptation of the augmentation means to individual characteristics. The compounding effect of fundamental human cognitive powers suggests further that systems designed for maximum effectiveness would require that these powers be developed as fully as possible -- by training, special mental tricks, improved language, new methodology.

In the automated system that we contemplate, the human should be able to draw on explicit-artifact process capability at many levels in the repertoire hierarchy; today, artifacts are involved explicitly in only the lower-order capabilities. In the future systems, for instance, it should be possible to have computer processes provide direct and significant help in his processes at many levels. We thus expect the effect of the computer in the system to be very much compounded. A great deal of richness in the future possibilities for automated H-LAM/T systems is implied here -- considerably more than many people realize who would picture the computer as just helping them do the things they do now. This type of compounding is related to the reverberating waves of change discussed in Section II-A.

Another factor can exert this type of compound effect upon over-all system performance: the human's unconscious processes. Clinical psychology seems to provide clear evidence that a large

proportion of a human's everyday activity is significantly mediated or basically prompted by unconscious mental processes that, although “natura” in a functional sense, are not rational. The observable mechanisms of these processes (observable by another, trained person) includes masking of the irrationality of the human's actions which are so affected, so that few of us will admit that our actions might be irrational, and most of us can construct satisfying rationales for any action that may be challenged.

Anything that might have so general an effect upon our mental actions as is implied here, is certainly a candidate for ultimate consideration in the continuing development of our intellectual effectiveness. It may be that the first stages of research on augmenting the human intellect will have to proceed without being able to do anything about this problem except accommodate to it as well as possible. This may be one of the very significant problems whose solution awaits our development of increased intellectual effectiveness.

EXAMPLES AND DISCUSSION

BACKGROUND

The conceptual structure which we have evolved to orient and guide the pursuit of increasing man's intellectual effectiveness has been described in the foregoing sections in a rather general and abstract fashion. In this section we shall try to develop more concrete images of these concepts, of some of the future possibilities for augmentation, and of the relationship between these different concepts and possibilities.

It must be borne in mind that a great deal of study and invention is yet to be done in developing the improved augmentation means that are bound to come, and that the examples which we present in this report are intended only to show what is meant by the generalizations which we use, and to provide a feeling on the part of the reader for the richness and power of the improvements we can likely develop in our augmentation means. Many of the examples are realizable today (in fact, some have been realized) and most of the rest are reasonably straight forward extrapolations into the near future. We predict that what actually develops in the new augmentation means will be consistent with our conceptual framework, but that the particulars will be full of surprises.

Each of the examples will show a facet of how the little steps that the human can take with his sensory-mental-motor apparatus can be organized cooperatively with the capabilities of artifacts to accomplish significant things in the way of achieving comprehension and solving problems. This organization, as we have shown in Section II, can be viewed as the five different types of structuring which we outlined, where much of the structuring that goes on in the human's total problem solving activity is for the purpose of building a mental structure which in a way "puts the human up where he can see what is going on and can point the direction to move next."

An early paper, offering suggestions toward augmenting the human intellect, that fits well and significantly within the framework which we have developed was written by Vannevar Bush (6) in 1945. Indeed, it fits so well, and states its points so nicely, that it was deemed appropriate to our purpose here to summarize it in detail and to quote from it at considerable length.

What Vannevar Bush proposed in 1945

He wrote as World War II was coming to an end, and his principal purpose seemed to be to offer new professional objectives to those scientists who were soon to be freed from war-motivated research and development. It would seem that he also wished to induce a general recognition of a growing problem -- storage, retrieval, and manipulation of information for and by intellectual workers -- and to show the possibilities he foresaw for scientific development of equipment which could significantly aid such workers in facing this problem. He summarized the situation: "...There is a growing mountain of research...The investigator is staggered by the findings and conclusions of thousands of other workers. Professionally our methods of transmitting and reviewing the results of research are generations old...truly significant attainments become lost in the mass of the inconsequential...The summation of human experience is being expanded at a prodigious rate, and the means we use for threading through the consequent maze to the momentarily important item is the same as was used in the days of square-rigged ships."

Then he brought out some general considerations for hope: "... But there are signs of a change as new and powerful instrumentalities come into use...Photocells...advanced photography...thermionic tubes... cathode ray tubes...relay combinations...there are plenty of mechanical aids with which to effect a transformation in scientific records." And he points out that devices which we commonly use today -- e.g., a calculating machine or an automobile -- would have been impossibly expensive to produce in earlier eras of our technological development. "...The world has arrived at an age of cheap complex devices of great reliability and something is bound to come of it."

In six and a half pages crammed full of well-based speculations, Bush proceeds to outline enough plausible artifact and methodology developments to make a very convincing case for the

augmentation of the individual intellectual worker. Extension of existing photographic techniques to give each individual a continuously available miniature camera for recording anything in view and of interest, and to realize a high-quality 100:1 linear reduction ratio for micro-record files for these photographs and published material; voice-recognition equipment (perhaps requiring a special language) to ease the process of entering new self-generated material into the written record -- these are to provide the individual with information-generating aid.

For the detailed manipulation of mathematical and logical expressions, Bush projects computing aids (which have been surpassed by subsequent development) that allow the individual to exercise a greater proportion of his time and talents in the tasks of selecting data and the appropriate transformations and processes which are to be executed, leaving to the machinery the subsequent execution. He suggests that new notation for our verbal symbols (perhaps binary) could allow character recognition devices to help even further in the information-manipulation area, and also points out that poor symbolism ("...the exceedingly crude way in which mathematicians express their relationships. They employ a symbolism which grew like Topsy and has little consistency; a strange fact in that most logical field.") stands in the way of full realization of machine help for the manipulations associated with the human's real time process of mathematical work. And "...Then, on beyond the strict logic of the mathematician, lies the application of logic in everyday affairs. We may some day click off arguments on a machine with the same assurance that we now enter sales on a cash register."

Then "...So much for the manipulation of ideas and their insertion into the record. Thus far we seem to be worse off than before -- for we can enormously extend the record; yet even in its present bulk we can hardly consult it. This is a much larger matter than merely the extraction of data for the purposes of scientific research; it involves the entire process by which man profits by his inheritance of acquired knowledge. The prime action of use is selection, and here we are halting indeed. There may be millions of fine thoughts, and the account of the experience on which they are based, all encased within stone walls of acceptable architectural form; but if the scholar can get at only one a week by diligent search, his syntheses are not likely to keep up with the current scene." He goes on to discuss possible developments that could allow very rapid (in the human's time frame) selection of unit records from a very large file -- where the records could be dry-process photographic micro-images upon which the user could add information at will.

Bush goes on to say, "The real heart of the matter of selection, however, goes deeper than a lag in the adoption of mechanisms...Our ineptitude in getting at the record is largely caused by the artificiality of systems of indexing." He observes the power of the associative recall which human memory exhibits, and proposes that a mechanization of selection by association could be realized to considerable advantage. He spends the last two pages (a quarter of his article) describing a device

embodying this capability, and points out some features of its use and of its likely effect. This material is so relevant and so well put that I quote it in its entirety:

”Consider a future device for individual use, which is a sort of mechanized private file and library. It needs a name, and to coin one at random, “memex” will do. A memex is a device in which an individual stores all his books, records, and communications, and which is mechanized so that it may be consulted with exceeding speed and flexibility. It is an enlarged intimate supplement to his memory.

”It consists of a desk, and while it can presumably be operated from a distance, it is primarily the piece of furniture at which he works. On the top are slanting translucent screens, on which material can be projected for convenient reading. There is a keyboard, and sets of buttons and levers. Otherwise it looks like an ordinary desk.

”In one end is the stored material. The matter of bulk is well taken care of by improved microfilm. Only a small part of the interior of the memex is devoted to storage, the rest to mechanism. Yet if the user inserted 5000 pages of material a day it would take him hundreds of years to fill the repository, so he can be profligate and enter material freely.

”Most of the memex contents are purchased on microfilm ready for insertion. Books of all sorts, pictures, current periodicals, newspapers, are thus obtained and dropped into place. Business correspondence takes the same path. And there is provision for direct entry. On the top of the memex is a transparent platen. On this are placed longhand notes, photographs, memoranda, all sort of things. When one is in place, the depression of a lever causes it to be photographed onto the next blank space in a section of the memex film, dry photography being employed.

”There is, of course, provision for consultation of the record by the usual scheme of indexing. If the user wishes to consult a certain book, he taps its code on the keyboard, and the title page of the book promptly appears before him, projected onto one of his viewing positions. Frequently-used codes are mnemonic, so that he seldom consults his code book; but when he does, a single tap of a key projects it for his use. Moreover, he has supplemental levers. On deflecting one of these levers to the right he runs through the book before him, each page in turn being projected at a speed which just allows a recognizing glance at each. If he deflects it further to the right, he steps through the book 10 pages at a time; still further at 100 pages at a time. Deflection to the left gives him the same control backwards.

”A special button transfers him immediately to the first page of the index. Any given book of his library can thus be called up and consulted with far greater facility than if it were taken from a shelf.

As he has several projection positions, he can leave one item in position while he calls up another. He can add marginal notes and comments, taking advantage of one possible type of dry photography, and it could even be arranged so that he can do this by a stylus scheme, such as is now employed in the telautograph seen in railroad waiting rooms, just as though he had the physical page before him.

"All this is conventional, except for the projection forward of present-day mechanisms and gadgetry. It affords an immediate step, however, to associative indexing, the basic idea of which is a provision whereby any item may be caused at will to select immediately and automatically another. This is the essential feature of the memex. The process of tying two items together is the important thing.

"When the user is building a trail, he names it, inserts the name in his code book, and taps it out on his keyboard. Before him are the two items to be joined, projected onto adjacent viewing positions. At the bottom of each there are a number of blank code spaces, and a pointer is set to indicate one of these on each item. The user taps a single key, and the items are permanently joined. In each code space appears the code word. Out of view, but also in the code space, is inserted a set of dots for photocell viewing; and on each item these dots by their positions designate the index number of the other item.

"Thereafter, at any time, when one of these items is in view, the other can be instantly recalled merely by tapping a button below the corresponding code space. Moreover, when numerous items have been thus joined together to form a trail, they can be reviewed in turn, rapidly or slowly, by deflecting a lever like that used for turning the pages of a book. It is exactly as though the physical items had been gathered together to form a new book. It is more than this, for any item can be joined into numerous trails.

"The owner of the memex, let us say, is interested in the origin and properties of the bow and arrow. Specifically he is studying why the short Turkish bow was apparently superior to the English long bow in the skirmishes of the Crusades. He has dozens of possibly pertinent books and articles in his memex. First he runs through an encyclopedia, finds an interesting but sketchy article, leaves it projected. Next, in a history, he finds another pertinent item, and ties the two together. Thus he goes, building a trail of many items. Occasionally he inserts a comment of his own, either linking it into the main trail or joining it by a side trail to a particular item. When it becomes evident that the elastic properties of available materials had a great deal to do with the bow, he branches off on a side trail which takes him through textbooks on elasticity and tables of physical constants. He inserts a page of longhand analysis of his own. Thus he builds a trail of his interest through the maze of materials available to him.

”And his trails do not fade. Several years later, his talk with a friend turns to the queer ways in which a people resist innovations, even of vital interest. He has an example, in the fact that the outraged Europeans still failed to adopt the Turkish bow. In fact he has a trail on it. A touch brings up the code book. Tapping a few keys projects the head of the trail. A lever runs through it at will, stopping at interesting items, going off on side excursions. It is an interesting trail, pertinent to the discussion. So he sets a reproducer in action, photographs the whole trail out, and passes it to his friend for insertion in his own memex, there to be linked into the more general trail.

”Wholly new forms of encyclopedias will appear, ready-made with a mesh of associative trails running through them, ready to be dropped into the memex and there amplified. The lawyer has at his touch the associated opinions and decisions of his whole experience, and of the experience of friends and authorities. The patent attorney has on call the millions of issued patents, with familiar trails to every point of his client's interest. The physician, puzzled by its patient's reactions, strikes the trail established in studying an earlier similar case, and runs rapidly through analogous case histories, with side references to the classics for the pertinent anatomy and histology. The chemist, struggling with the synthesis of an organic compound, has all the chemical literature before him in his laboratory, with trails following the analogies of compounds, and side trails to their physical and chemical behavior.

”The historian, with a vast chronological account of a people, parallels it with a skip trail which stops only at the salient items, and can follow at any time contemporary trails which lead him all over civilization at a particular epoch. There is a new profession of trail blazers, those who find delight in the task of establishing useful trails through the enormous mass of the common record. The inheritance from the master becomes, not only his additions to the world's record, but for his disciples the entire scaffolding by which they were erected.

”Thus science may implement the ways in which man produces, stores, and consults the record of the race. It might be striking to outline the instrumentalities of the future more spectacularly, rather than to stick closely to the methods and elements now known and undergoing rapid development, as has been done here. Technical difficulties of all sorts have been ignored, certainly, but also ignored are means as yet unknown which may come any day to accelerate technical progress as violently as did the advent of the thermionic tube. In order that the picture may not be too commonplace, by reason of sticking to present-day patterns, it may be well to mention one such possibility, not to prophesy but merely to suggest, for prophecy based on extension of the known has substance, while prophecy founded on the unknown is only a doubly involved guess.

"All our steps in creating or absorbing material of the record proceed through one of the senses - the tactile when we touch keys, the oral when we speak or listen, the visual when we read. Is it not possible that some day the path may be established more directly?

"We know that when the eye sees, all the consequent information is transmitted to the brain by means of electrical vibrations in the channel of the optic nerve. This is an exact analogy with the electrical vibrations which occur in the cable of a television set: they convey the picture from the photocells which see it to the radio transmitter from which it is broadcast. We know further that if we can approach that cable with the proper instruments, we do not need to touch it; we can pick up those vibrations by electrical induction and thus discover and reproduce the scene which is being transmitted, just as a telephone wire may be tapped for its message.

"The impulse which flow in the arm nerves of a typist convey to her fingers the translated information which reaches her eye or ear, in order that the fingers may be caused to strike the proper keys. Might not these currents be intercepted, either in the original form in which information is conveyed to the brain, or in the marvelously metamorphosed form in which they then proceed to the hand?

"By bone conduction we already introduce sounds into the nerve channels of the deaf in order that they may hear. Is it not possible that we may learn to introduce them without the present cumbersomeness of first transforming electrical vibrations to mechanical ones, which the human mechanism promptly transforms back to the electrical form? With a couple of electrodes on the skull the encephalograph now produces pen-and-ink traces which bear some relation to the electrical phenomena going on in the brain itself. True, the record is unintelligible, except as it points out certain gross mis-functioning of the cerebral mechanism; but who would now place bounds on where such a thing may lead?

"In the outside world, all forms of intelligence, whether of sound or sight, have been reduced to the form of varying currents in an electric circuit in order that they may be transmitted. Inside the human frame exactly the same sort of process occurs. Must we always transform to mechanical movements in order to proceed from one electrical phenomenon to another? It is a suggestive thought, but it hardly warrants prediction without losing touch with reality and immediateness.

"Presumably man's spirit should be elevated if he can better review his shady past and analyze more completely and objectively his present problems. He has built a civilization so complex that he needs to mechanize his record more fully if he is to push his experiment to its logical conclusion and not merely become bogged down part way there by overtaxing his limited memory. His excursion may be more enjoyable if he can reacquire the privilege of forgetting the manifold things

he does not need to have immediately at hand, with some assurance that he can find them again if they prove important.

”The applications of science have built man a well-supplied house, and are teaching him to live healthily therein. They have enabled him to throw masses of people against another with cruel weapons. They may yet allow him truly to encompass the great record and to grow in the wisdom of race experience. He may perish in conflict before he learns to wield that record for his true good. Yet, in the application of science to the needs and desires of man, it would seem to be a singularly unfortunate stage at which to terminate the process, or to lose hope as to the outcome.

Comments Related to Bush's Article

There are many significant items in the article, but the main ones upon which we shall comment here will be those relative to the use and implications of his Memex. The associative trails whose establishment and use within the files he describes at some length provide a beautiful example of a new capability in symbol structuring that derives from new artifact-process capability, and that provides new ways to develop and portray concept structures. Any file is a symbol structure whose purpose is to represent a variety of concepts and concept structures in a way that makes them maximally available and useful to the needs of the human's mental-structure development -- within the limits imposed by the capability of the artifacts and human for jointly executing processes of symbol-structure manipulation. The Memex allows a human user to do more conveniently (less energy, more quickly) what he could have done with relatively ordinary photographic equipment and filing systems, but he would have had to spend so much time in the lower-level processes of manipulation that his mental time constants of memory and patience would have rendered the system unusable in the detailed and intimate sense which Bush illustrates.

The Memex adds a factor of speed and convenience to ordinary filing-system (symbol-structuring) processes that would encourage new methods of work by the user, and it also adds speed and convenience for processes not generally used before. Making it easy to establish and follow the associative trails makes practical a new symbol-structuring process whose use can make a significant difference in the concept structuring and basic methods of work. It is also probable that clever usage of associative-trail manipulation can augment the human's process structuring and executing capabilities so that he could successfully make use of even more powerful symbol-structure manipulation processes utilizing the Memex capabilities. An example of this general sort of thing was given by Bush where he points out that the file index can be called to view at the push of a button, which implicitly provides greater capability to work within more sophisticated and complex indexing systems

Note, too, the implications extending from Bush's mention of one user duplicating a trail (a portion of his structure) and giving it to a friend who can put it into his Memex and integrate it into his own trail (structure). Also note the “wholly new forms of encyclopedia”, the profession of “trail blazers,” and the inheritance from a master including “the entire scaffolding” by which such additions to the world's record were erected. These illustrate the types of changes in the ways in which people can cooperate intellectually that can emerge from the augmentation of the individuals. This type of change represents a very significant part of the potential value in pursuing research directly on the means for making individuals intellectually more effective.

Some Possibilities with Cards and Relatively Simple Equipment

A number of useful new structuring processes can be made available to an individual through development and use of relatively simple equipment that is mostly electromechanical in nature and relatively cheap. We can begin developing examples of this by describing the hand operated, edge-notched card system that I developed and used over the past eight years.

An Existing Note and File System

The “unit records” here, unlike those in the Memex example, are generally scraps of typed or handwritten text on IBM-card sized edge-notchable cards. These represent little “kernels” of data, thought, fact, consideration concepts, ideas, worries, etc., that are relevant to a given problem area in my professional life. Each such specific problem area has its notecards kept in a separate deck, and for each such deck there is a master card with descriptors associated with individual holes about the periphery of the card. There is a field of holes reserved for notch coding the serial number of a reference from which the note on a card may have been taken, or the serial number corresponding to an individual from whom the information came directly (including a code for myself, for self-generated thoughts).

None of the principles of indexing or sorting used here is new: coordinate-indexing descriptors with direct coding on edge notched cards, with needle-sort retrieval. Mainly what is new is the use of the smaller units of information, in restricted-subject sets (notedecks) so that I gain considerable flexibility in the manipulations of my thought products at the level at which I actually work in my minute-by-minute struggle with analytical and formulative thought. Not only do my own thoughts produce results in this fashion, but when I digest the writings of another person, I find generally anyway that I have extracted from his structure and integrated into my own a specific selection of facts, considerations, ideas, etc. Often these different extracted items fit into different places in my structure, or become encased in special substructures as I modify or expand his concepts. Extracting such items or kernels and putting each on its own notecard helps this process considerably -- the role or position of each such item in the growth of the note structure is independent, and yet if

desired all can quickly be isolated and extracted by simple needle sorting on the reference-number notching field.

These notecards represent much more than just an information file. They provide a workspace for me, in which I can browse, make additions or corrections, or build new sets of thought kernels with a good deal of freedom. I can leave notes with suggestions or questions for myself that will drop out at an appropriate later time. I can do document-reference searches with good efficiency, too, by needle sorting for notes within relevant descriptor categories. Any notecard with relevant notes on it points to the original source (by the source serial number, which I always write, together with the page, at the top of the card). When I am in the process of developing an integrated write-up covering some or all of the notedeck's material, I can quickly needle out a set of cards relevant to the topic under consideration at the moment -- with all other cards in one pile to the side -- and I need do a very minimum of hand searching or stacking in special little category piles. If I utilize specific information from another person, I can register my acknowledgment in my draft write-up merely by writing in the source serial number that is at the top of the notecard -- it is a straight forward clerical job for a secretary later to arrange footnote entries and numbering.

Comments on the System

First, let me relate what has been described to the special terms brought out in previous sections. The writing contained on each notecard is a small-sized symbol structure, representing or portraying to me a small structure of concepts. The notches on the edges of the cards are symbols that serve to tie these card-sized symbol sub structures into a large symbol structure (the notedeck). One aspect of the structure is the physical grouping of the cards at a given time -- which happens to be the only aspect of the over-all structuring that my human capabilities can make direct use of -- and in this respect I can execute processes which produce restructuring (that is, physical re grouping) that helps me considerably to perceive and assimilate the concepts of worth to me. This restructuring is effected by composite processes involving me, a master code card, a sorting needle, and a work surface. I can add to the symbol structure by executing other composite processes which involve me, writing instruments (pen, pencil, or type writer), a master code card, and a card notcher.

If my mental processes were more powerful, I could dispense with the cards, and hold all of the card-sized concept structures in my memory, where also would be held the categorization linkages that evolved as I worked (with my feet up on the artifacts and my eyes closed). As it is, and as it probably always will be no matter how we develop or train our mental capabilities, I want to work in problem areas where the number and interrelationship complexity of the individual factors involved are too much for me to hold and manipulate within my mind. So, my mind develops conscious sets of concepts, or recognizes and selects them from what it perceives in the work of

others, and it directs the organization of an external symbol structure in which can be held and portrayed to the mind those concepts I cannot (reliably) remember or whose manipulations I cannot visualize. The price I pay for this augmentation shows up in the time and energy involved in manipulating artifacts to manipulate symbols to give me this artificial memory and visualization of concepts and their manipulation.

Associative-Linking Possibilities

But let us go further with discussing specific examples of means for augmenting our intellects. In using the edge-notched-card system described, I found several types of structuring which that system could not provide, but which would both be very useful and probably obtainable with reasonably practical artifact means. One need arose quite commonly as trains of thought would develop on a growing series of note cards. There was no convenient way to link these cards together so that the train of thought could later be recalled by extracting the ordered series of note cards. An associative-trail scheme similar to that outlined by Bush for his Memex could conceivably be implemented with these cards to meet this need and add a valuable new symbol-structuring process to the system. Straightforward engineering development could provide a mechanism that would be able to select a specific card from a relatively large deck by a parallel edge-notch sort on a unique serial number notched into each card, and the search mechanism could be set up automatically by a hole sensing mechanism from internal punches on another card that was placed in the sensing slot. An auxiliary notching mechanism could automatically give succeeding serial-number encoding to new notecards as they are made up.

Suppose that one wants to link Card B to Card A, to make a trail from A to B. He puts Card B into a slot so that the edge-notched coding of the card's serial number can automatically be sensed, and slips Card A under a hole-punching head which duplicates the serial-number code of Card B in the coding of the holes punched in a specific zone on Card A. Later, when he may have discovered Card A, and wishes to follow this particular associative trail to the next card, he aligns that zone on Card A under a hole-sensing head which reads the serial number for Card B therein and automatically sets up the sorting mechanism. A very quick and simple human process thus initiates the automatic extraction of the next item on the associative trail. It's not unreasonable to assume that establishing a link would take about three seconds, and tracing a link to the next card about three to five seconds.

There would still be descriptor-code notching and selection to provide for general grouping classifications -- and we can see that the system could really provide a means for working within the structure of the contained information.

An Experiment Illustrating Usage and Further System Possibilities

I once tried to use my cards, with their separate little “concept packets,” in the process of developing a file memo outlining the status and plans of a research project. I first developed a set of cards upon each of which I described a separate consideration, possibility, or specification about the memo -- in the disorderly sequence in which they occurred to me as my thoughts about the basic features of the memo evolved. Right off the bat I noticed that there were two distinct groups -- some ideas were about what the memo ought to accomplish, what time period it should cover, when it should be finished, what level and style of presentation should be used, etc., and some ideas were about the subject of the memo. As more thoughts developed, I found that the latter group also divided into ideas representing possible content and those representing possible organization.

I separated the cards into three corresponding groups (which I shall call Specification, Organization, and Content), and began to organize each of them. I started with the Specification group (it being the “highest” in nature), and immediately found that there were several types of notes within that group. Just as there had been in the total group. Becoming immediately suspicious, I sorted through each of the other two main groups and found similar situations in each. In each group there was finally to emerge a definite set of statements (product statements) that represented that group's purpose -- e.g., the specifications currently accepted for the design of the memo -- and some of the cards contained candidate material for this. But there were also considerations about what these final statements might include or exclude or take into account, or conditions under which inclusion or modification might be relevant, or statements that were too bulky or brief or imprecise to be used as final statements.

It became apparent that the final issuance from my work, the memo itself, would represent but one facet of a complex symbol structure that would grow as the work progressed -- a structure comprising three main substructures, each of which had definite substructuring of its own that was apparent. I realized that I was being rather philosophically introspective with all of this analysis, but I was curious as to the potential value of future augmentation means in allowing me to deal explicitly with these types of structuring. So I went ahead, keeping the groups and sub-groups of cards separated, and trying to organize and develop them.

I found rather quickly that the job of extracting, re arranging, editing, and copying new statements into the cards which were to represent the current set of product statements in each grouping was rather tedious. This brought me to appreciate the value of some sort of copying device with which I could transfer specified strings of words from one card to another, thus composing new statements from fragments of existing ones. This type of device should not be too hard to develop and produce

for a price that a professional man could justify paying, and it would certainly facilitate some valuable symbol-structuring pro

I also found that there would have been great value in having available the associative-trail marking and following processes. Statements very often had implicit linkages to other statements in the same group, and it would have been very useful to keep track of these associations. For instance, when several consideration statements bore upon a given product statement, and when that product statement came to be modified through some other consideration, it was not always easy to remember why it had been established as it had. Being able to fish out the other considerations linked to that statement would have helped considerably.

Also, trial organizations of the statements in a group could be linked into trial associative trails, so that a number of such organizations could be constructed and considered without copying that many sets of specially ordered statements. Any of the previously considered organizations could be reconstructed at will.

In trying to do flexible structuring and restructuring within my experiment, I found that I just didn't have the means to keep track of all of the kernel statements (cards) and the various relationships between them that were important -- at least by means that were easy enough to leave time and thought capacity enough for me to keep in mind the essential nature of the memo-writing process. But it was a very provocative experience, considering the possibilities that I sensed for the flexible and powerful ways in which I could apply myself to so universal a design task if I but had the necessary means with which to manipulate symbol structures.

It would actually seem quite feasible to develop a unit record system around cards and mechanical sorting, with automatic trail establishment and trail-following facility, and with associated means for selective copying or data transfer, that would enable development of some very powerful methodology for everyday intellectual work. It is plain that even if the equipment (artifacts) appeared on the market tomorrow, a good deal of empirical research would be needed to develop a methodology that would capitalize upon the artifact process capabilities. New concepts need to be conceived and tested relative to the way the "thought kernels" could be knitted together into working structures, and relative to the conceptual presentations which become available and the symbol-manipulation processes which provide these presentations.

Such an approach would present useful and interesting research problems, and could very likely produce practical and significant results (language, artifacts, methodology) for improving the effectiveness of professional problem solvers. However, the technological trends of today foretell the obsolescence of such electromechanical information handling equipment. Very likely, by the

time good augmentation systems could be developed, and the first groups of users began to prove them out so that they could gain more widespread acceptance, electronic data processing equipment would have evolved much further and become much more prevalent throughout the critical-problem domains of our society where such ideas would first be adopted. The relative limitations of the mechanical equipment in providing processes which could be usefully integrated into the system would soon lead to its replacement by electronic computer equipment.

The next set of descriptive examples will involve the use of electronic computers, and their greatly increased flexibility and processing potential will be evident. Research based upon such electronic artifacts would be able to explore language and methodology innovations of a much wider range of sophistication than could research based upon limited and relatively inflexible electromechanical artifacts. In particular, the electronic-based experimental program could simulate the types of processes available from electromechanical artifacts, if it seemed possible (from the vantage of experience with the wide range of augmentation processes) that relatively powerful augmentation systems could be based upon their capabilities -- but the relative payoffs for providing even-more-sophisticated artifact capabilities could be assessed too so that considerations of how much to invest in capital equipment versus how much increase in human effectiveness to expect could be based upon some experimental data.

A Quick Summary of Relevant Computer Technology

This section may be of value both to readers who are already familiar with computers, and to those who are not. A little familiarity with computer technology, enough to help considerably in understanding the augmentation possibilities discussed in this report, can be gained by the uninitiated. For those already familiar with the technology, the following discussion can perhaps help them gain more understanding of our concepts of process and symbol structuring.

A computer is directly capable of performing any of a basic repertoire of very primitive symbol-manipulation processes (such as "move the symbol in location A to location 12417," or "compare the symbol in location A with that in location B, and if they are the same, set switch S to ON"). There may be from ten to over a hundred different primitive processes which a particular machine can execute, and all of the computer's more sophisticated processes are structured from these primitive processes. It takes a repertoire of surprisingly few such primitive processes to enable the construction of any symbol-manipulation process that can be explicitly described in any language.

Somewhat the same situation exists relative to symbol structures i.e., there are only a very few primitive symbols with which the machine can actually work, and any new and different symbol has to be defined to the machine as a particular structure (or organization) of its primitive symbols.

Actually, in every commercial digital computer, there are only two primitive symbols. Usually these are dealt with in standard-sized packets (called "words") of from eighteen to forty-eight primitive symbols, but arbitrary use can be made of individual primitives or of subgroups of the word.

To have the computer perform a non-trivial task or process, a structure of the primitive processes is organized (a computer program) and stored within the computer as a corresponding symbol structure. The computer successively examines the symbol substructure representing each primitive process in the program and executes that process -- which usually alters the total internal symbol structure of the machine in some way. It makes no difference to the computer whether the symbols involved in the re-structuring represent part of the computer program or part of the information upon which the program is operating. The ability to have the computer modify its own process structure (program) has been a very important factor in the development of its power.

Thus, some very sophisticated techniques for process and symbol structuring have evolved in the computer field, as evidenced by the very sophisticated processes (e.g., predicting election returns, calculating orbits, translating natural languages) that can be structured to manipulate very complex structures of symbols. Among the more interesting computer-process structures that have evolved are those that can automatically develop a structure of primitive computer processes to accomplish symbol manipulation tasks that are specified on a relatively high level of abstraction. Special languages have been evolved in several fields (e.g., ALGOL and FORTRAN for scientific calculations, COBOL for business processing) that enable explicit prescription of complex manipulation processes in a rapid and concise manner by a human, thinking about the processes in a rather natural manner, so that special computer programs or process structures (called Translators, Compilers, or sometimes in a slightly different sense, Interpreters) can construct the necessary structures of primitive processes and symbols that would enable the computer to execute the prescribed processes. This development has extended immensely our capability for making use of computers -- otherwise the specification of a complex process would often occupy a formidable number of man hours, and be subject to a great many errors which would be very costly to find and correct.

Computers have been used to simulate dynamic systems for which we humans had none but descriptive models, from which we otherwise could gain little feel for the way the system behaves. A very notable instance of this, for our consideration, has been in the area of the human thought processes. Newell, Shaw, and Simon initiated this approach) from which there has derived a number of features of interest to us. For one, they discovered that the symbol structures and the process structures required for such simulation became exceedingly complex, and the burden of organizing these was a terrific impediment to their simulation research. They devised a structuring technique for their symbols that is basically simple but from which stem results that are very

elegant. Their basic symbol structure is what they call a 'list,' a string of substructures that are linked serially in exactly the manner proposed by Bush for the associative trails in his Memex -- i.e., each substructure contains the necessary information for locating the next substructure on the list. Here, though, each substructure could also be a list of substructures, and each of these could also, etc. Their standard manner for organizing the data which the computer was to operate upon is thus what they term "list structuring."

They also developed special languages to describe different basic processes involved in list-structure manipulation. The most widely used of their languages, IPL-V (the fifth version of their Information Processing Languages), is described in a recent book edited by Newell. In these languages, both the data to be worked upon and the symbols which designate the processes to be executed upon that data are developed in list-structure form.

Other languages and techniques for the manipulation of list structures have been described by McCarthy, by Gelernter, Hansen, and Gerberich, by Yngve, by Perlis and Thornton, by Carr, and by Weizenbaum. The application of these techniques has been mainly of two types -- one of modeling complex processes and systems (e.g., the human thought processes), where the emphasis is upon the model and its behavior, and the other of trying to get computer behavior that is intelligent whether or not the processes and behavior resemble those of a human. The languages and techniques used in both types of application promise to be of considerable value to the development of radical new augmentation systems for human problem solvers, and we shall deal later with them in more detail.

Computers have various means for storing symbols so that they are accessible to it for manipulation. Assuming that the human might want to have a repertoire of sixty-four basic symbols (letters, numbers, special symbols), we can discuss various forms of storage in terms of their capacity for storing these kinds of symbols (each of which would be structured, in the computer and storage devices, as a group of six primitive computer symbols). Fast access to an arbitrary choice of a few neighboring symbols (of the human's repertoire) can be had to perhaps 100,000 such symbols within the period in which the computer can execute one of its primitive processes (from two to ten millionths of a second, depending upon the computer involved). This is the so-called high-speed, random-access working store, where space for the human's symbols might cost between sixty cents and \$1.50 per symbol.

Cheaper, larger-capacity backup storage is usually provided by devices to which access takes considerably longer (in the computer's time reference). A continuously rotating magnetic drum can hold perhaps a million of these symbols, for which access to a random storage position may average a thirtieth of a second (waiting for the drum to come around to bring that storage position under the magnetic reading head). This is short in the human's time scale, but a reasonably fast computer

could execute about 3,000 of its primitive processes during that time. Generally, information transfer between a drum and fast-access working storage takes place in blocks of data which are stored in successive positions around the drum. Such block-transfer is accomplished by a relatively small structure of primitive computer processes that cyclically executes the transfer of one word at a time until the designated block has been transferred. Drum storage costs about 5¢ per each of the basic symbols used by the human in our example.

Another type of backup storage uses a number of large, thin discs (about three feet in diameter), with magnetic coating on the surfaces. The discs are stacked with enough space between each so that a moveable read-record head can be positioned radially to line up over a specific circular track of symbol storage space. A commercially available disc storage system could hold over a hundred million of the human's basic symbols, to which random access would average about a tenth of a second, and where the cost per symbol-space would be about one seventh of a cent.

Magnetic tapes are commonly used for backup storage, too. For these, the random access time for storage blocks are of the order of a minute or two. Here, however, the actual storage units (the tape reels) can be taken off and shelf stored, so the total storage capacity may be very large -- however, the time to locate a reel and exchange reels on the tape transport adds to the above-quoted access time -- and this locating and reel changing are not generally automatic processes (i.e., a human has to do them). A transport unit, connected to the computer, might cost \$30,000, with tape reels at \$50 each holding about five million of the human's basic symbols. For one reel, storage space for each such symbol cost about two-thirds of a cent, but for twenty full reels in a "library" the cost comes down to about one-thirtieth of a cent per symbol space.

Other types of buffer storage for computer symbol structures are becoming available, and there is considerable economic demand spurring continuing research toward storage means that give high capacity at low cost, and with as short an access time as possible. Within the next ten years there would seem to be a very high probability of significant advances to this end.

For presenting computer-stored information to the human, techniques have been developed by which a cathode-ray-tube (of which the television picture tube is a familiar example) can be made to present symbols on their screens of quite good brightness, clarity, and with considerable freedom as to the form of the symbol. Under computer control an arbitrary collection of symbols may be arranged on the screen, with considerable freedom as to relative location, size, and brightness. Similarly, line drawings, curves, and graphs may be presented, with any of the other symbols intermixed. It is possible to describe to the computer, and thereafter use, new symbols of arbitrary shape and size. On displays of this sort, a light pen (a pen-shaped tool with a flexible wire to the

electronic console) can be pointed by the human at any symbol or line on the display, and the computer can automatically determine what the pen is pointing at.

A cathode-ray-tube display of this sort is currently limited in resolution to about 800 lines across the face of the tube (in either direction). The detail with which a symbol may be formed, and the preciseness with which the recurrent images of it may be located, are both affected by this figure so that no matter how large the screen of such a tube, the maximum number of symbols that can be put on with usable clearness remains the same.

The amount of usable information on such a screen, in the form of letters, numbers, and diagrams, would be limited to about what a normal human eye could make out at the normal reading distance of fourteen inches on a surface 3-1/2 inches square, or to what one could discern on an ordinary 8-1/2-by-11-inch sheet of paper at about three feet. This means that one could not have a single-tube display giving him an 8-1/2-by-11-inch frame to view that would have as much on it as he might be used to seeing, say on the page of a journal article.

The costs of such displays are now quite high -- ranging from \$20,000 to \$60,000, depending upon the symbol repertoire, symbol-structure display capacity, and the quality of the symbol forms. One should expect these prices to be lowered quite drastically as our technology improves and the market for these displays increases.

Much cheaper devices can "draw" arbitrary symbol shapes and diagrams on paper, at a speed for symbols that is perhaps a quarter of the rate that a typewriter can produce them. Also, special typewriters (at \$3,000 to \$4,000 apiece) can type out information on a sheet of paper, as well as allow the human to send information to the computer via the keyboard. But these two types of devices do not allow fast and flexible rearrangement of the symbols being displayed, which proves to be an important drawback in our current view of future possibilities for augmentation.

For communicating to the computer, considerable freedom exists in arranging push buttons, switches, and keysets for use by the human. The "interpretation" or response to be made by the computer to the actuation of any button, switch, or key (or to any combination thereof) can be established in any manner that is describable as a structure of primitive computer processes -- which means essentially any manner that is explicitly describable. The limitation on the flexibility and power of any explicit "shorthand" system with which the human may wish to utilize these input devices is the human's ability to learn and to use them.

There are also computer-input devices that can sense enough data from handwriting to allow a computer to recognize a limited number of handwritten symbols -- both as they are being written

and afterwards. Means for recognizing typescript are rather well developed and are already being designed into some large documentation and language translation systems. Also, a little progress has been made toward developing equipment that can recognize a limited spoken vocabulary. There is considerable economic pressure toward developing useful and cheap devices of this type, and we can expect relatively sophisticated capabilities to become available within the next ten years. Such equipment may play an important role in the individual-augmentation systems of the future (but our feeling is that a very powerful augmentation system can be developed without them).

An important type of development for our consideration of providing individual humans with close-coupled computer services is what is known as time sharing. Suppose a number of individual users connect to the same computer. The computer can be programmed to serve them under any of a wide variety of rules. One such could be similar to the way the telephone system gives you attention and service when you ask for it -- i.e., if too many other demands are not being made for service at that time, you get instant attention; otherwise, you wait until some service capacity is free to attend to you.

Our view of the interaction of human and computer in the future augmented system sees a large number of reactively simple processes (human scale of large and simple) being performed by the computer for the human -- processes which often will require only a few thousandths of a second of actual computer manipulation. Such a fast and agile helper as a computer can run around between a number of masters and seldom keep any of them waiting (at least, not long enough that they would notice it or be inconvenienced appreciably). Occasionally, of course, much larger periods of computer time will be needed by an individual, and then the other users might get their periodic milliseconds of service slipped in during these longer processes.

Other Related Thought and Work

When we began our search, we found a great deal of literature which put forth thought and work of general significance to our objective -- frankly, too much. Without having a conceptual framework, we could not efficiently filter out the significant kernels of fact and concept from the huge mass which we initially collected as a "natural first step" in our search. We feel rather unscholarly not to buttress our conceptual framework with plentiful reference to supporting work, but in truth it was too difficult to do. Developing the conceptual structure represented a sweeping synthesis job full of personal constructs from smatterings picked up in many places. Under these conditions, giving reference to a backup source would usually entail qualifying footnotes reflecting an unusual interpretation or exonerating the other author from the implications we derived from his work. We look forward to a stronger, more comprehensive, and more scholarly presentation evolving out of future work.

However, we do want to acknowledge thoughts and work we have come across that bear most directly upon the possibilities of using a computer in real-time working association with a human to improve his working effectiveness. These fall into two categories. The first category, which would include this report, presents speculations and possibilities but does not include reporting of significant experimental results. Of these, Bush is the earliest and one of the most directly stimulating. Licklider provided the most general clear case for the modern computer, and coined the expression, "man-computer symbiosis" to refer to the close interaction relationship between the man and computer in mutually beneficial cooperation. Ulam has specifically recommended close man-computer interaction in a chapter entitled, "synergism," where he points out in considerable detail the types of mathematical work which could be aided. Good includes some conjecture about possibilities of intellectual aid to the human by close cooperation with a computer in a rather general way, and also presents a few interesting thoughts about a network model for structuring the conceptual kernels of information to facilitate a sort of self-organizing retrieval system. Ramo has given a number of talks dealing with the future possibilities of computers for "extending man's intellect," and wrote several articles. His projections seem slanted more toward larger bodies of humans interacting with computers, in less of an intimate personal sense than the above papers or than our initial goal. Fein in making a comprehensive projection of the growth and dynamic inter-relatedness of "computer-related sciences," includes specific mention of the enhancement of human intellect by cooperative activity of men, mechanisms, and automata. He coined the term "synnoetics" as applicable generally to the cooperative interaction of people, mechanisms, plant or animal organisms, and automata into a system whose mental power is greater than that of its components, and presented a good picture of the integrated way in which many currently separate disciplines should be developed and taught in the future to do justice to their mutual roles in the important meta-discipline defined as "synnoetics."

In the second category, there have been a few papers published recently describing actual work that bears directly upon our topic. Licklider and Clark, and Culler and Huff, in the 1962 Spring Joint Computer Conference, gave what are essentially progress reports of work going on now in exactly this sort of thing -- a human with a computer backed display getting minute-by-minute help in solving problems. Teager reports on the plans and current development of a large time-sharing system at MIT, which is planned to provide direct computer access for a number of outlying stations located in scientists' offices, giving each of these users a chance for real-time utilization of the computer.

There are several efforts that we have heard about, but for which there are either no publications or for which none have been discovered by us. Mr. Douglas Ross, of the Electronic Systems Laboratory at MIT has, we learned by direct conversation, been thinking and working on real-time

man-machine interaction problems for some years. We have recently learned that a graduate student at MIT, Glenn Randa, has developed the design of a remote display console under Ross for his graduate thesis project. We understand that another graduate student there, Ivan Sutherland, is currently using the display-computer facility on the TX-2 computer at Lincoln Lab to develop cooperative techniques for engineering-design problems. And at RAND, we have learned by personal discussion that Cliff Shaw, Tom Ellis, and Keith Uncapher have been involved in implementing a multi-station time-sharing system built around their JOHNNIAC computer. Termed the JOHNNIAC Open-Shop System (JOSS for short), it apparently is near completion, and will use remote type writer stations.

Undoubtedly, there are efforts of others falling into either or both categories that have been overlooked. Such oversight has not been intentional, and it is hoped that these researchers will make their pertinent work known to us.

HYPOTHETICAL DESCRIPTION OF COMPUTER-BASED AUGMENTATION SYSTEM

Let us consider some specific possibilities for redesigning the augmentation means for an intellectually oriented, problem-solving human. We choose to present those developments of language and methodology that can capitalize upon the symbol-manipulating and portraying capabilities of computer-based equipment. The picture of the possibilities to pursue will change and grow rapidly as research gets under way, but we need to provide what pictures we can -- to give substance to the generalities developed in Section II, to try to impart our feeling of rich promise, and to introduce a possible research program (Section IV).

Although our generalizations (about augmentation means, capability hierarchies, and mental-, concept-, symbol-, process-, and physical structuring) might retain their validity in the future -- for instance, our generalized prediction that new developments in concept, symbol, and process structuring will prove to be tremendously important -- the specific concepts, symbol structures, and processes that evolve will most likely differ from what we know and use now. In fact, even if we in some way could know now what would emerge after say, ten years of research, it is likely that any but a general description would be difficult to express in today's terminology.

Background

To try to give you (the reader) a specific sort of feel for our thesis in spite of this situation, we shall present the following picture of computer-based augmentation possibilities by describing what

might happen if you were being given a personal discussion-demonstration by a friendly fellow (named Joe) who is a trained and experienced user of such an augmentation system within an experimental research program which is several years beyond our present stage. We assume that you approach this demonstration-interview with a background similar to what the previous portion of this report provides -- that is, you will have heard or read a set of generalizations and a few rather primitive examples, but you will not yet have been given much of a feel for how a computer-based augmentation system can really help a person.

Joe understands this and explains that he will do his best to give you the valid conceptual feel that you want -- trying to tread the narrow line between being too detailed and losing your over-all view and being too general and not providing you with a solid feel for what goes on. He suggests that you sit and watch him for a while as he pursues some typical work, after which he will do some explaining. You are not particularly flattered by this, since you know that he is just going to be exercising new language and methodology developments on his new artifacts -- and after all, the artifacts don't look a bit different from what you expected -- so why should he keep you sitting there as if you were a complete stranger to this stuff? It will just be a matter of "having the computer do some of his symbol-manipulating processes for him so that he can use more powerful concepts and concept-manipulation techniques," as you have so often been told.

Joe has two display screens side by side, but one of them he doesn't seem to use as much as the other. And the screens are almost horizontal, more like the surface of a drafting table than the near-vertical picture displays you had somehow imagined. But you see the reason easily, for he is working on the display surface as intently as a draftsman works on his drawings, and it would be awkward to reach out to a vertical surface for this kind of work. Some of the time Joe is using both hands on the keys, obviously feeding information into the computer at a great rate.

Another slight surprise, though -- you see that each hand operates on a set of keys on its own side of the display frames, so that the hands are almost two feet apart. But it is plain that this arrangement allows him to remain positioned over the frames in a rather natural position, so that when he picks the light pen out of the air (which is its rest position, thanks to a system of jointed supporting arms and a controlled tension and rewind system for the attached cord) his hand is still on the way from the keyset to the display frame. When he is through with the pen at the display frame, he lets go of it, the cord rewinds, and the pen is again in position. There is thus a minimum of effort, movement, and time involved in turning to work on the frame. That is, he could easily shift back and forth from using keyset to using light pen, with either hand (one pen is positioned for each hand), without moving his head, turning, or leaning.

A good deal of Joe's time, though, seems to be spent with one hand on a keyset and the other using a light pen on the display surface. It is in this type of working mode that the images on the display frames changed most dynamically. You receive another real surprise as you realize how much activity there is on the face of these display tubes. You ask yourself why you weren't prepared for this, and you are forced to admit that the generalizations you had heard hadn't really sunk in -- "new methods for manipulating symbols" had been an oft-repeated term, but it just hadn't included for you the images of the free and rapid way in which Joe could make changes in the display, and of meaningful and flexible "shaping" of ideas and work status which could take place so rapidly.

Then you realized that you couldn't make any sense at all out of the specific things he was doing, nor of the major part of what you saw on the displays. You could recognize many words, but there were a good number that were obviously special abbreviations of some sort. During the times when a given image or portion of an image remained unchanged long enough for you to study it a bit, you rarely saw anything that looked like a sentence as you were used to seeing one. You were beginning to gather that there were other symbols mixed with the words that might be part of a sentence, and that the different parts of what made a full-thought statement (your feeling about what a sentence is) were not just laid out end to end as you expected. But Joe suddenly cleared the displays and turned to you with a grin that signaled the end of the passive observation period, and also that somehow told you that he knew very well that you now knew that you had needed such a period to shake out some of your limited images and to really realize that a "capability hierarchy" was a rich and vital thing.

"I guess you noticed that I was using unfamiliar notions, symbols, and processes to go about doing things that were even more unfamiliar to you?" You made a non-committal nod -- you saw no reason to admit to him that you hadn't even been able to tell which of the things he had been doing were to cooperate with which other things -- and he continued. "To give you a feel for what goes on, I'm going to start discussing and demonstrating some of the very basic operations and notions I've been using. You've read the stuff about process and process-capability hierarchies, I'm sure. I know from past experience in explaining radical augmentation systems to people that the new and powerful higher-level capabilities that they are interested in -- because basically those are what we are all anxious to improve -- can't really be explained to them without first giving them some understanding of the new and powerful capabilities upon which they are built. This holds true right on down the line to the type of low-level capability that is new and different to them all right, but that they just wouldn't ordinarily see as being 'powerful.' And yet our systems wouldn't be anywhere near as powerful without them, and a person's comprehension of the system would be rather shallow if he didn't have some understanding of these basic capabilities and of the hierarchical structure built up from them to provide the highest-level capabilities."

Single-Frame Composition

”For explanation purposes here, let's say that the lowest level at which the computer system comes into direct play in my capability hierarchy is in the task of what I'll call 'single-frame composition.' We'll stick to working with prose text in our examples -- most people can grasp easily enough what we are doing there without having to have special backgrounds in mathematics or science as they would to gain equal comprehension for some of the similar sorts of things we do with diagrams and mathematical equations. This low-level composition task is just what you normally do with a pen or pencil or typewriter on a piece of paper -- that is, assemble a bunch of symbols before your eyes in order to portray something which you have in mind.”

You listened and watched as Joe showed you some of the different ways in which the composition of straightforward text was made easier for him in this system. With either hand, Joe could “type” (the keysets didn't look at all like typewriter keyboards) individual letters and numbers, and if he had directed it to do so, the computer would put each successive symbol next to its predecessor just as a typewriter does -- only here there was completely automatic “carriage return” service. This didn't impress you very much, since an automatic carriage-return feature was sort of a trivial return on the investment behind all of this equipment -- but then you reflected that, as long as the computer was there anyway, to help do all the flashy things you had witnessed earlier, one might as well use it in all of the little helpful ways he could.

But there were other ways in which help was derived for this composition task. He showed you how he could call up the dictionary definition to any word he had typed in, with but a few quick flicks on the keyset. Synonyms or antonyms could just as easily be brought forth. This also seemed sort of trivially obvious, and Joe seemed to know that you would feel so. “It turns out that this simple capability makes it feasible to do some pretty rough tasks in the upper levels of the capability hierarchy -- where precise use of special terms really pays off, where the human just couldn't be that precise by depending upon his unaided memory for definitions and 'standards,' and where using dictionary and reference-book lookup in the normal fashion would be so distracting and time-consuming that the task execution would break down. We've tried taking this feature away in some of these processes up there, and believe me, the result was a mess.

You could get some dim feeling for what he meant, having watched him working for a while, but you were nevertheless much relieved to find the next thing he showed you to be more directly impressive. He showed you how he could single out a group of words (called the “object symbol string,” or simply “object string”) and define an abbreviation term, composed of any string of symbols he might choose, that became associated with the object string in computer storage. At any later time (until he chose to discard that particular abbreviation from his working vocabulary) the

typing of the abbreviation term would call forth automatically the “printing” on the display of the entire object string. Joe showed you another way in which this abbreviation feature might work. He “arranged” for the computer to print the abbreviation on the display, just the way he typed it in. At a subsequent reading, if he had forgotten what the abbreviation stood for, he could call for substitution of the full object string to refresh his memory.

Then he showed you how this sort of facility had been extended, in a refined way, to provide a rather powerful sort of shorthand. He could hit a great many combinations of keys on his keyset -- i.e., any one stroke of his hand could depress a number of keys, which gave him over a thousand unique single-stroke signals to the computer with either hand. Some of these signals were used as abbreviations for entire words. It seems that, for instance, the 150 most commonly used words in a natural language made up about half of any normal text in that language. Joe said that it was thus quite feasible to learn and use the single-stroke abbreviations for about half of the words he used, but beyond that each added percent began to require him to have too many abbreviations under his command. But he said that there were a lot of word endings, letter pairs (diagrams), and letter triplets (trigrams) that were so common as to make it pay to abbreviate them to a single stroke. A whole word so abbreviated saved typing all the letters as well as the spaces at either side of the word, and a word-ending abbreviated by a single stroke saved typing the letters and the end-of-word space. He claimed that he could comfortably rattle off about 180 words a minute -- faster than he could comfortably talk. You believed him after he transcribed your talking for a minute or so, and it gave you an eerie feeling to see the near instantaneous appearance of your words and sentences in neat printed form.

Joe said that there were other miscellaneous simple features, and some quite sophisticated features to help the composition process. He made some brief references to statistical predictions that the computer could make regarding what you were going to type next, and that if you got reasonably skillful you could “steer through the extrapolated prediction field” as you entered your information and often save energy and time. You gathered that he thought you would saturate about there on this particular subject, because he went on to the next.

Single-Frame Manipulation

”Even if I couldn't actually specify new symbols here any faster than with a typewriter, the extreme flexibility that this computer system provides for making changes in what is presented on the display screen would make me very much more effective in creating finished text than I could ever be on a typewriter.” With this statement, Joe proceeded to show you what he meant. The frame full of your transcribed speech was still showing, and it represented the clumsy phrasing and illogical progression of thought so typical of extemporaneous speech. Joe took the light gun in his right

hand, and with a deft flick of it, coordinated with a stroke of his left hand on its keyset, caused the silent and instantaneous deletion of a superfluous word. The word disappeared from the frame, and the rest of the text simultaneously readjusted to present the neat, no-gap, full-line appearance it had had.

With but slightly more motion of his light pen, he could similarly delete any string of words or letters. He demonstrated this by cutting out what I thought to be some relevant prose, and then he showed how the system allowed for second thoughts about such human-directed processes -- those words were automatically saved for a brief period in case he wanted to call them back. Leaving his light pen pointed at the space where a deleted symbol string used to be, Joe could reinstate it instantaneously with one stroke of his left hand.

Adding one more light-pen pointing to what it took to delete an arbitrary string of symbols, Joe could direct the computer to move that string from where it was to insert it at a new point which his light pen designated. Again it would disappear instantaneously from where it had been, but now the modified display would show the old text to have been spread apart just enough at the indicated point to hold this string. The text would all still look as neat as if freshly retyped. With similar types of keyset and light-pen operations, Joe could change paragraph break points, transpose two arbitrary symbol strings (words, sentences, paragraphs, etc., or fragments thereof), readjust margins of arbitrary sections of text -- essentially being able to affect immediately any of the changes that a proofreader might want to designate with his special marks, only here the proofreader is always looking at clean text as if it had been instantaneously retyped after each designation had been made.

Joe also demonstrated how he could request that each instance of the use of a given term be changed to a newly designated term, and this would again be instantaneously accomplished. Also, he could arbitrarily set the margins between which any section of text must appear, and its line lengths and number of lines would automatically be adjusted. He showed how this was useful in displaying parallel or counter arguments -- although he said that actual use of this feature was a bit more sophisticated -- by squeezing each into half width and putting them side by side (with a vertical line suddenly separating them). One of the sections of text was about a third longer than the other -- but two quick strokes with Joe's left hand caused the computer to adjust the display automatically. The middle separator line was moved toward the shorter piece of text, and the line lengths of the two sections were adjusted so that they occupied the same length along the display frame. Yes, you were beginning to get a feel for what the expression "flexible new methods for manipulating symbol structures" might really imply, at least on this basic-capability level.

Structuring an Argument

”If we want to go on to a higher-level capability to give you a feeling for how our rebuilt capability hierarchy works, it will speed us along to look at how we might organize these more primitive capabilities which I have demonstrated into some new and better ways to set up what we can call an 'argument.' This refers loosely to any set of statements (we'll call them 'product statements') that represents the product of a period of work toward a given objective. Confused? Well, take the simple case where an argument leads to a single product statement. For instance, you come to a particular point in your work where you have to decide what to do for the next step. You go through some reasoning process -- usually involving statements -- and come up with a statement specifying that next step. That final statement is the product statement, and it represents the product of the argument or reasoning process which led to it.

”You usually think of an argument as a serial sequence of steps of reason, beginning with known facts, assumptions, etc., and progressing toward a conclusion. Well, we do have to think through these steps serially, and we usually do list the steps serially when we write them out because that is pretty much the way our papers and books have to present them -- they are pretty limiting in the symbol structuring they enable us to use. Have you even seen a 'scrambled-text' programmed instruction book? That is an interesting example of a deviation from straight serial presentation of steps.

”Conceptually speaking, however, an argument is not a serial affair. It is sequential, I grant you, because some statements have to follow others, but this doesn't imply that its nature is necessarily serial. We usually string Statement B after Statement A, with Statements C, D, E, F, and so on following in that order -- this is a serial structuring of our symbols. Perhaps each statement logically followed from all those which preceded it on the serial list, and if so, then the conceptual structuring would also be serial in nature, and it would be nicely matched for us by the symbol structuring.

”But a more typical case might find A to be an independent statement, B dependent upon A, C and D independent, E depending upon D and B, E dependent upon C, and F dependent upon A, D, and E. See, sequential but not serial? A conceptual network but not a conceptual chain. The old paper and pencil methods of manipulating symbols just weren't very adaptable to making and using symbol structures to match the ways we make and use conceptual structures. With the new symbol-manipulating methods here, we have terrific flexibility for matching the two, and boy, it really pays off in the way you can tie into your work.

This makes you recall dimly the generalizations you had heard previously about process structuring limiting symbol structuring, symbol structuring limiting concept structuring, and concept structuring limiting mental structuring. You nod cautiously, in hopes that he will proceed in some

way that will tie this kind of talk to something from which you can get the “feel” of what it is all about. As it turns out, that is just what he intends to do.

”Let's actually work some examples. You help me.” And you become involved in a truly fascinating game. Joe tells you that you are to develop an argument leading to statements summarizing the augmentation means so far revealed to you for doing the kind of straight-text work usually done with a pencil and eraser on a single sheet of paper. You unconsciously look for a scratch pad before you realize that he is telling you that you are going to do this the “augmented way” by using him and his system -- with artful coaching from him. Under a bit of urging from him, you begin self-consciously to mumble some inane statements about what you have seen, what they imply, what your doubts and reservations are, etc. He mercilessly ignores your obvious discomfort and gives you no cue to stop, until he drops his hands to his lap after he has filled five frames with these statements (the surplus filled frames disappeared to somewhere -- you assume Joe knows where they went and how to get them back).

”You notice how you wandered down different short paths, and crisscrossed yourself a few times?” You nod -- depressed, no defense. But he isn't needling you. “Very natural development, just the way we humans always seem to start out on a task for which we aren't all primed with knowledge, method, experience, and confidence -- which is to include essentially every problem of any consequence to us. So let's see how we can accommodate the human's way of developing his comprehension and his final problem solution.

”Perhaps I should have stopped sooner -- I am supposed to be coaching you instead of teasing you - - but I had a reason. You haven't been making use of the simple symbol-manipulation means that I showed you -- other than the shorthand for getting the stuff on the screens. You started out pretty much the way you might with your typewriter or pencil. I'll show you how you could have been doing otherwise, but I want you to notice first how hard it is for a person to realize how really unquestioning he is about the way he does things. Somehow we implicitly view most all of our methods as just sort of 'the way things are done, that's all.' You knew that some exotic techniques were going to be applied, and you'll have to admit that you were passively waiting for them to be handed to you.”

With a non-committal nod, you suggest getting on with it. Joe begins, “You're probably waiting for something impressive. What I'm trying to prime you for, though, is the realization that the impressive new tricks all are based upon lots of changes in the little things you do. This computerized system is used over and over and over again to help me do little things -- where my methods and ways of handling little things are changed until, lo, they've added up and suddenly I can do impressive new things.”

You don't know. He's a nice enough guy, but he sure gets preachy. But the good side of your character shows through, and you realize that everything so far has been about little things -- this is probably an important point. You'll stick with him. Okay, so what could you have been doing to use the simple tricks he had shown you in a useful way? Joe picks up the light pen, poises his other hand over the keyset, and looks at you. You didn't need the hint, but thanks anyway, and let's start rearranging and cleaning up the work space instead of just dumping more raw material on it.

With closer coaching now from Joe, you start through the list of statements you've made and begin to edit, re-word, compile, and delete. It's fun -- "put that sentence back up here between these two" - - and blink, it's done. "Group these four statements, indented two spaces, under the heading "shorthand," and blinko, it's done. "Insert what I say next there, after that sentence." You dictate a sentence to extend a thought that is developing, and Joe effortlessly converts it into an inserted new sentence. Your ideas begin to take shape, and you can continually re-work the existing set of statements to keep representing the state of your "concept structure."

You are quite elated by this freedom to juggle the record of your thoughts, and by the way this freedom allows you to work them into shape. You reflected that this flexible cut-and-try process really did appear to match the way you seemed to develop your thoughts. Golly, you could be writing math expressions, ad copy, or a poem, with the same type of benefit. You were ready to tell Joe that now you saw what he had been trying to tell you about matching symbol structuring to concept structuring -- when he moved on to show you a succession of other techniques that made you realize you hadn't yet gotten the full significance of his pitch.

So far the structure that you have built with your symbols looks just like what you might build with pencil-and-paper techniques -- only here the building is so much easier when you can trim, extend, insert, and rearrange so freely and rapidly. But the same computer here that gives us these freedoms with so trivial an application of its power, can just as easily give us other simple capabilities which we can apply to the development and use of different types of structure from what we used to use. But let me unfold these little computer tricks as we come to them.

"When you look at a given statement in the middle of your argument structure, there are a number of things you want to know. Let's simplify the situation by saying that you might ask three questions, 'What's this?', 'How come?', and 'So what?' Let's take these questions one at a time and see how some changes in structuring might help a person answer them better.

You look at a statement and you want to understand its meaning. You are used to seeing a statement portrayed in just the manner you might hear it -- as a serial succession of words. But, just as with the statements within an argument, the conceptual relationship among the words of a sentence is not

generally serial, and we can benefit in matching better to the conceptual structure if we can conveniently work with certain non-serial symbol-structuring forms within sentences.

”Most of the structuring forms I’ll show you stem from the simple capability of being able to establish arbitrary linkages between different substructures, and of directing the computer subsequently to display a set of linked substructures with any relative positioning we might designate among the different substructures. You can designate as many different kinds of links as you wish, so that you can specify different display or manipulative treatment for the different types.”

Joe picked out one of your sentences, and pushed the rest of the text a few lines up and down from it to isolate it. He then showed you how he could make a few strokes on the keyset to designate the type of link he wanted established, and pick the two symbol structures that were to be linked by means of the light pen. He said that most links possessed a direction, i.e., they were like an arrow pointing from one substructure to another, so that in setting up a link he must specify the two substructures in a given order.

He went to work for a moment, rapidly setting up links within your sentence. Then he showed you how you could get some help in looking at a statement and understanding it. “Here is one standard portrayal, for which I have established a computer process to do the structuring automatically on the basis of the inter-word links.” A few strokes on the keyset and suddenly the sentence fell to pieces - - different parts of it being positioned here and there, with some lines connecting them. “Remember diagramming sentences when you were studying grammar? Some good methods, plus a bit of practice, and you’d be surprised how much a diagrammatic breakdown can help you to scan a complex statement and untangle it quickly.

”We have developed quite a few more little schemes to help at the statement level. I don’t want to tangle you up with too much detail, though. You can see, probably, that quick dictionary-lookup helps.” He aimed at a term with the light pen and hit a few strokes on the keyset, and the old text jumped farther out of the way and the definition appeared above the diagram, with the defined term brighter than the rest of the diagram. And he showed you also how you could link secondary phrases (or sentences) to parts of the statement for more detailed description. These secondary substructures wouldn’t appear when you normally viewed the statement, but could be brought in by simple request if you wanted closer study.

”It proves to be terrifically useful to be able to work easily with statements that represent more sophisticated and complex concepts. Sort of like being able to use structural members that are

lighter and stronger -- it gives you new freedom in building structures. But let's move on -- we'll come back to this area later, if we have time.

"When you look at a statement and ask, 'How come?', you are used to scanning back over a serial array of previously made statements in search of an understanding of the basis upon which this statement was made. But some of these previous statements are much more significant than others to this search for understanding. Let us use what we call 'antecedent links' to point to these, and I'll give you a basic idea of how we structure an argument so that we can quickly track down the essential basis upon which a given statement rests."

You helped him pick out the primary antecedents of the statement you had been studying, and he established links to them. These statements were scattered back through the serial list of statements that you had assembled, and Joe showed you how you could either brighten or underline them to make them stand out to your eye -- just by requesting the computer to do this for all direct antecedents of the designated statement. He told you, though, that you soon get so you aren't very much interested in seeing the serial listing of all of the statements, and he made another request of the computer (via the keyset) that eliminated all the prior statements, except the direct antecedents, from the screen. The subject statement went to the bottom of the frame, and the antecedent statements were neatly listed above it.

Joe then had you designate an order of "importance to comprehension" among these statements, and he rearranged them accordingly as fast as you could choose them. (This choosing was remarkably helped by having only the remainder statements to study for each new choice -- another little contribution to effectiveness, you thought.) He mentioned that you could designate orderings under several different criteria, and later have the display show whichever ordering you wished. This, he implied, could be used very effectively when you were building or studying an argument structure in which from time to time you wanted to strengthen your comprehension relative to different aspects of the situation.

"Each primary antecedent can similarly be linked to its primary antecedents, and so on, until you arrive at the statements representing the premises, the accepted facts, and the objectives upon which this argument had been established. When we had established the antecedent links for all the statements in the argument, the question 'So what?' that you might ask when looking at a given statement would be answered by looking for the statements for which the given statement was an antecedent. We already have links to these consequents -- just turn around the arrows on the antecedent links and we have consequent links. So we can easily call forth an uncluttered display of consequent statements to help us see why we needed this given statement in the argument.

”To help us get better comprehension of the structure of an argument, we can also call forth a schematic or graphical display. Once the antecedent-consequent links have been established, the computer can automatically construct such a display for us.” So, Joe spent a few minutes (with your help) establishing a reasonable set of links among the statements you had originally listed. Then another keyed-in request to the computer, and almost instantaneously there appeared a network of lines and dots that looked something like a tree -- except that sometimes branches would fuse together. “Each node or dot represents one of the statements of your argument, and the lines are antecedent-consequent links. The antecedents of one statement always lie above that statement -- or rather, their nodes lie above its node. When you get used to using a network representation like this, it really becomes a great help in getting the feel for the way all the different ideas and reasoning fit together -- that is, for the conceptual structuring.”

Joe demonstrated some ways in which you could make use of the diagram to study the argument structure. Point to any node, give a couple of strokes on the keyset, and the corresponding statement would appear on the other screen -- and that node would become brighter. Call the antecedents forth on the second screen, and select one of special interest -- deleting the others. Follow back down the antecedent trail a little further, using one screen to look at the detail at any time, and the other to show you the larger view, with automatic node-brightening indication of where these detailed items fit in the larger view.

”For a little embellishment here, and to show off another little capability in my repertoire, let me label the nodes so that you can develop more association between the nodes and the statements in the argument. I can do this several ways. For one thing, I can tell the computer to number the statements in the order in which you originally had them listed, and have the labeling done automatically.” This took him a total of five strokes on the keyset, and suddenly each node was made into a circle with a number in it. The statements that were on the second screen now each had its respective serial number sitting next to it in the left margin. “This helps you remember what the different nodes on the network display contain. We have also evolved some handy techniques for constructing abbreviation labels that help your memory quite a bit.

”Also, we can display extra fine-structure and labeling detail within the network in the specific local area we happen to be concentrating upon. This finer detail is washed out as we move to another spot with our close attention, and the coarser remaining structure is compressed, so that there is room for our new spot to be blown up. It is a lot like using zones of variable magnification as you scan the structure -- higher magnification where you are inspecting detail, lower magnification in the surrounding field so that your feel for the whole structure and where you are in it can stay with you.”

General Symbol Structuring

"If you are tangling with a problem of any size -- whether it involves you for half an hour or two years -- the entire collection of statements, sketches, computations, literature sources, and source extracts that is associated with your work would in our minds constitute a single symbol structure. There may be many levels of substructuring between the level of individual symbols and that represented by the entire collection. You and I have been working with some of the lower-ordered substructures -- the individual statements and the multi-statement arguments -- and have skimmed through some of the ways to build and manipulate them. The results of small arguments are usually integrated in a higher level network of argument or concept development, and these into still higher-level networks, and so on. But at any such level, the manner in which the interrelationship between the kernels of argument can be tagged, portrayed, studied and manipulated is much the same as those which we have just been through.

"Substructures that might represent mathematical or formal-logic arguments may be linked right in with substructures composed of the more informal statements. Substructures that represent graphs, curves, engineering drawings, and other graphical forms can likewise be integrated. One can also append special substructures, of any size, to particular other substructures. A frequent use of this is to append descriptive material -- something like footnotes, only much more flexible. Or, special messages can be hung on that offer ideas such as simplifying an argument or circumventing a blocked path -- to be uncovered and considered at some later date. These different appended substructures can remain invisible to the worker until such time as he wants to flush them into view. He can ask for the cue symbols that indicate their presence (identifying where they are linked and what their respective types are) to be shown on the network display any time he wishes, and then call up whichever of them he wishes. If he is interested in only one type of appended substructure, he can request that only the cues associated with that type be displayed.

"You should also realize that a substructure doesn't have to be a hunk of data sitting neatly distinct within the normal form of the larger structure. One can choose from a symbol structure (or substructure, generally) any arbitrary collection of its substructures, designate any arbitrary structuring among these and any new substructures he wants to add, and thus define a new substructure which the computer can untangle from the larger structure and present to him at any time. The associative trails that Bush suggested represent a primitive example of this. A good deal of this type of activity is involved during the early, shifting development of some phase of work, as you saw when you were collecting tentative argument chains. But here again, we find ever more delightful ways to make use of the straightforward-seeming capabilities in developing new higher-level capabilities -- which, of course, seem sort of straight forward by then, too.

"I found, when I learned to work with the structures and manipulation processes such as we have outlined, that I got rather impatient if I had to go back to dealing with the serial-statement structuring in books and journals, or other ordinary means of communicating with other workers. It is rather like having to project three-dimensional images onto two-dimensional frames and to work with them there instead of in their natural form. Actually, it is much closer to the truth to say that it is like trying to project n-dimensional forms (the concept structures, which we have seen can be related with many many nonintersecting links) onto a one-dimensional form (the serial string of symbols), where the human memory and visualization has to hold and picture the links and relationships. I guess that's a natural feeling, though. One gets impatient any time he is forced into a restricted or primitive mode of operation -- except perhaps for recreational purposes.

"I'm sure that you've had the experience of working over a journal article to get comprehension and perhaps some special-purpose conclusions that you can integrate into your own work. Well, when you ever get handy at roaming over the type of symbol structure which we have been showing here, and you turn for this purpose to another person's work that is structured in this way, you will find a terrific difference there in the ease of gaining comprehension as to what he has done and why he has done it, and of isolating what you want to use and making sure of the conditions under which you can use it. This is true even if you find his structure left in the condition in which he has been working on it -- that is, with no special provisions for helping an outsider find his way around. But we have learned quite a few simple tricks for leaving appended road signs, supplementary information, questions, and auxiliary links on our working structures -- in such a manner that they never get in our way as we work -- so that the visitor to our structure can gain his comprehension and isolate what he wants in marvelously short order. Some of these techniques are quite closely related to those used in automated-instruction programming -- perhaps you know about 'teaching machines?'

"What we found ourselves doing, when having to do any extensive digesting of journal articles, was to type large batches of the text verbatim into computer store. It is so nice to be able to tear it apart, establish our own definitions and substitute, restructure, append notes, and so forth, in pursuit of comprehension, that it was generally well worth the trouble. The keyset shorthand made this reasonably practical. But the project now has an optical character reader that will convert our external references into machine code for us. The references are available for study in original serial form on our screens, but any structuring and tagging done by a previous reader, or ourselves, can also be utilized.

"A number of us here are using the augmented systems for our project research, and we find that after a few passes through a reference, we very rarely go back to it in its original form. It sits in the archives like an orange rind, with most of the real juice squeezed out. The contributions from these

references form sturdy members of our structure, and are duly tagged as to source so that acknowledgment is always implicitly noted. The analysis and digestion that any of us makes on such a reference is fully available to the others. It is rather amazing how much superfluous verbiage is contained in those papers merely to try to make up for the pitifully sparse possibilities available for symbol structuring in printed text.”

Process Structuring

There was a slight pause while Joe apparently was reflecting upon something. He started to speak, thought differently of it, and turned to flash something on a screen. You looked quickly, anticipating that now you would comprehend. Well, more of the display looked meaningful to you than when you had first watched him going about his work, but you realized that you were still a bit uneducated. I've developed a sequence for presenting the different basic features of our augmentation system that seems to work pretty well, and I just wanted to be sure I was still following it reasonably closely.”

He noticed you wrinkle your face as you looked at the display. “It's time to shift the topic a bit, and some of the things on the screen that are probably puzzling you can make a starting point for a new discussion phase. See, when I outlined a delivery for giving a feel for these techniques to the uninitiated, I could have sketched out the subject matter in a skeletal argument structure. From what we've been through so far, you might expect it to be like that. What I did, though, was to treat the matter as a process that I was going to execute the process of giving you a lecture demonstration. It is a rather trivial exercise of the techniques we have for developing and manipulating processes, but anyway that's the form I chose for making the notes.

”A process is something that is designed, built, and used -- as is any tool. In the general sense in which we consider processes to be a part of our augmentation system, it is absolutely necessary that there be effective capability for designing and building processes as well as for using them. For one thing, the laying out of objectives and a method of approach for a problem represent a form of process design and building, to our way of looking at it. And an independent problem solver certainly has to have this capability. Indeed, we find that designing and coordinating one's sequence of steps, in high levels or in low levels of such process structuring, is an extremely important part of the total activity.

”One of our research guys in the early phases of our augmentation development was considered (then) to be a bug on this topic. He maintained that about ten percent of the little steps we took all day accounted for ninety percent of the progress toward the goals we claimed to pursue -- that is, that ninety percent of our actions and thoughts were coupled to our net progress in only a very

feeble way. Well, we can't analyze the old ways of doing things very accurately to check his estimated figures, but we certainly have come to be in general sympathy with his stand. We have developed quite a few concepts and methods for using the computer system to help us plan and supervise sophisticated courses of action, to monitor and evaluate what we do, and to use this information as direct feedback for modifying our planning techniques in the future.

”There are, of course, the explicit computer processes which we use, and which our philosophy requires the augmented man to be able to design and build for himself. A number of people, outside our research group here, maintain stoutly that a practical augmentation system should not require the human to have to do any computer programming -- they feel that this is too specialized a capability to burden people with. Well, what that means in our eyes, if translated to a home workshop, would be like saying that you can't require the operating human to know how to adjust his tools, or set up jigs, or change drill sizes, and the like. You can see there that these skills are easy to learn in the context of what the human has to learn anyway about using the tools, and that they provide for much greater flexibility in finding convenient ways to use the tools to help shape materials.

”It won't take too much time to give you a feel for the help methods we have for working on computer-process structures -- or programs -- because there is quite a bit of similarity in concept to what you have seen in the symbol-structuring techniques. No matter what language you use -- whether machine language, list language, or ALGOL, for instance -- you build up the required process structure by organizing statements in that language. Each statement specifies a given process to your computer. Well, you have already seen how you can get help in developing precise and powerful statements, or in gaining quick comprehension of statements, by charting or diagramming them and using special links between the different parts. “Look here.” And he went after what he said was a typical process structure, to give you an example of what he was talking about. In several brief, successive frame displays, before he got to the one he wanted, you got glimpses of network schematics that reminded you of those used in symbol structuring. But, what he finally had on the display frame was quite different from the argument statements you had seen.

”In explaining symbol-structuring to you, I used the likely questions, 'What's this?' 'How come?' and 'So what?' to point out the usefulness of some of our structuring methods. Here, in process structuring, corresponding questions about a statement might be: 'What does it say to do?' 'What effect will that have?' and 'Why do we want that done?' Let's take a quick look at some of the ways you can get help in answering them.

”The language used to compose these process-description statement for the computer is considerably more compact and precise than is a natural language, such as English, and there is

correspondingly less advantage to be gained by appending special links and tags for giving us humans a better grasp of their meaning. However, as you see in this left-hand section of the statement portrayal, geometrical grouping, linking, and positioning of the statement components are used in the blown-up statement display. But this portrayal doesn't stem from special appended information, it can be laid out like this automatically by the computer, just from the cues it gets from the necessary symbol components of the statement. The different significant relationships are more perceptible to a human in this way of laying it out, and an experienced human thus gets quite a bit of help in answering the first question: 'What does it say to do?'

For the second question, relative to what effect the specified action will have, some of these symbols to the right give you a quick story about the very detailed and immediate effect on the state of the symbol structure which this process structure is manipulating. Other symbols here provide keys which a light-pen selection can activate to bring to you displays of that symbol structure, usually a choice of several relevant views at different levels of the structure. Then I can use the keyset to ask for the preceding statement, if I'm a little puzzled about the detailed manipulation -- or, I can request a specific higher-level view of the process structure by light-pen selection on one of these remaining symbols here.

So saying, Joe selected one of these symbols with his pen, and a new and different display popped into view. "This is the next level up in the process structure. It consists of lists of compactly abbreviated statements, and some condensed notes about their effects. If we want, we can blow up one at a time as we study over the list. In this context, one can get some answer to the larger picture of what effect will a given statement have, and also some answer to the question about why we want a given effect produced. But this is a sort of a holdover from old programming habits, and most of us nowadays are making considerably more use of the schematic techniques that evolved out of the program flow-charting techniques and out of our symbol-structuring techniques.

"I know that you have less previous familiarity with the nature of programs than you do with the nature of arguments, so I'll just give you a few quick views of what these process-structure schematic portrayals look like, and not try to explain them in any detail. He flashed a few on the screen, and indicated how some of the different features could give the human a quick appreciation of how different component processes were cooperating to produce a more sophisticated process. You could appreciate some of the tricks of linking in explanatory and descriptive substructure and the general means of using all the different symbol-structuring tricks for representing to the human the considerations, critical features, and interdependencies involved in the process structure.

"Most of this portrayal technique actually represents special structuring of what we previously defined in a loose way as arguments. The human who wants to approach an established process

structure in order to modify it, needs to gain comprehension of the relevant features both of the functioning and of the design of the structure. You saw how this could be facilitated by our symbol-structuring techniques. And if he is building a new process structure or changing an existing one, he needs to structure the argument or reasoning behind the design. We have developed a number of special symbol-structuring techniques that allow us to match especially well to the concepts involved in designing processes.

”But there is a very significant feature involved in this particular type of process structuring that I should tell you about. It is based upon the fact that the process-description language for the computer is formal and precise. Because of this fact, we can establish explicit rules for treating statements in this language, and for treating symbol structures composed of these statements, such that computer processes based upon these rules can be said to extract meaning from these statements and to do operations based upon this meaning. The result is that the computer is able to find answers to a much wider range of questions about a specified process structure than it could if only the structural characteristics were discernible to it.

”In our studying and designing process structures, we have found many ways to capitalize upon this more sophisticated question-answering capability now possessed by the computer. We are learning, for instance, how to get the computer to decide whether or not some types of design specifications are met, and if not, where the limitation exists. Or, perhaps we approach an already designed process structure which we think we can modify, or from which we can extract some useful sub-process that we contemplate incorporating into another process we are designing. We are getting terrific help in this type of instance, since we can now ask the computer direct questions about types of capability and limitation in this structure. The computer can even lead us directly to the particular design features from which these capabilities or limitations stem, and it is simple then to examine the descriptive and explanatory arguments linked thereto in order to see why these features were designed into the structure.

”But I don't want to spend a disproportionate amount of time on the computer processes. The augmented man is engaged more often in structuring what we call composite processes than he is in structuring computer processes. For instance, planning a research project, or a day's work, are examples of structuring composite processes. A composite process, remember, is organized from both human processes and computer processes -- which includes, of course, the possible inclusion of lower-order composite processes. The structuring here differs from that of a computer process mainly in the sophistication of the sub-processes which can be specified for the human to do. Some of these specifications have to be given in a language which matches the human's rich working framework of concepts -- and we have been demonstrating here with English for that purpose -- but quite a few human-executed processes can be specified in the high-level computer-processing

language even though we don't know how to describe them in that language. This means that there are a few composite-process structures about which the computer can answer very useful questions for us.

”But to be more specific -- we find that setting up objectives, designing a method of approach, and then implementing that method are of course our fundamental operating sequence -- done over and over again in the many levels of our activity. We mentioned above what the characteristic structural difference was between computer processes and composite processes. But perhaps more important to us is the difference in the way we work with composite-process structures. Here is a crude but succinct way to put this. With the human contributing to a process, we find more and more as the process becomes complex that the value of the human's contribution depends upon how much freedom he is given to be disorderly in his course of action. For instance, we provide him as much help as possible in making a plan of action. Then we give him as much help as we can in carrying it out. But we also have to allow him to change his mind at almost any point, and to want to modify his plans. So, we provide augmentation help to him for keeping track of his plans where he is in them, what has been happening in carrying them out to date -- and for evaluating possibilities that might occur to him for changing the plans. In fact, we are even learning how the computer can be made to watch for some kinds of plan-change possibilities, and to point them out to the human when they arise.

”Here's a simple example of this sort of help for the human. Last winter, we designed a computer process that can automatically monitor the occurrence of specified types of computer usage over a specified period of time, and which, from the resulting data, can deduce a surprising amount of information regarding how the human made use of that time. This was quite helpful to us for evaluating our ways of doing things. Then we added more features to the program, in which the computer occasionally interrupts the human's activity and displays some questions to be answered. From these answers, together with its normal monitoring data, the program can provide evaluative data regarding the relative success of his different work methods. Our augmentation researchers became intrigued by this angle and bore down a little on it. They came up with a package process which gives the human many different types of feedback about his progress and way of doing things. Now, as part of my regular practice, I spend about five minutes out of each hour exercising with this package. This almost always reveals things to me that change at least the slant of my approach during the next hour, and often stimulates a relatively significant change in my short-range plans

”You appreciate, of course, that I accomplish many more meaningful steps in an hour now than I used to, or than would be your norm now. This once-an-hour review for me now might compare with a once-a-day review for you, as far as the distance travelled between reviews is concerned.

”Our way of structuring the statement of our objectives, the arguments which lead to the design of our plans, and the working statements of our plans, has been influenced by this review process. We found special types of tags and descriptive codes which we could append to these respective planning structures as we developed them which later facilitated our man-computer cooperative review of them. Also, our methods of developing these structures have evolved to facilitate their later modification. For instance, every basic consideration upon which a given planning statement is based is linked to that statement as a matter of standard argument structuring. But we have taken to linking special tagging codes into these argument structures involving our planning, to identify for the computer some of the different types of dependency relationships in the antecedent linkages. Later, if we consider changing the plan, these special tags often enable us to make use of some special computer processes that automatically isolate the considerations relevant to a particular type of change we have in mind.

”Maybe an example will help here. There is a plan I am currently using for the way I go about entering miscellaneous scraps of information into my total symbol structure. It is designed so that there will be a good chance for these scraps later to be usefully integrated. It turns out that this plan is closely coupled in its design argument to the general plan for reviewing process structures -- and symbol-structures, too, for that matter. Recently, I got an idea as to how I might add a little feature to that process that specially suited my particular way of wanting to deal with miscellaneous thoughts that I get. By various means, I very quickly learned that this would be easy to do if I could but reverse the order in which I execute the sub-process Steps A and B, when I enter a piece of information. I had to find out if I could safely reverse their order without getting into trouble someplace in my system.

”This I could do relatively rapidly, by your standards, by snooping down the antecedent trails, looking for statements relevant to this timing question. There is, in fact, a semi-automatic processes available to me for speeding just such searches. The computer keeps track of where I have looked, where I've marked things as yes, or no, or possible, and does the bookkeeping and calculating necessary to guide me through an optimum search strategy. But the special tagging we do when we make a process structure lets this search be fully automatic when certain kinds of relationships are involved -- and relative timing happens to be one of these relationships.

”So I phrased a question which essentially asked for considerations relevant to the order in which these two steps were executed, and turned the computer loose. It took about three seconds for the results to be forthcoming -- you haven't yet seen me request a task that took a noticeable period of machine time, have you? But anyway, the computer discovered a relevance trail that ended up showing that reversing the order of Steps A and B during the information-scrap entry process would

cripple a certain feature in the planning-review process, where miscellaneous thoughts and possibilities are gleaned from this store to be considered relative to the planning.

”But let's try to back away from details for a bit, now, and see if we can get a feeling for the significance of the things we've been talking about. Comparison with other working domains would be helpful, perhaps. If you were an inventor of useful mechanisms, you would like to have a wide range of materials-processing and shaping techniques available to you. This would give you more freedom and more interesting possibilities in the way you worked and designed. But many of these techniques are very specialized; they require special equipment, special skills to execute the processing and shaping, and special knowledge about applicability and possibilities for the techniques.

”Suppose you were told that you could subscribe to a community- owned installation of special equipment -- containing all sorts of wonderful instruments tools and machines for measuring and processing with such as chemical, optical, mechanical, electronic, pneumatic, vacuum, metallurgy, and human factors. But this wasn't all that was included in the subscription. There would be a specialist assigned to you, instantly available for consultation and help whenever you requested it. He wouldn't have high-level theoretical training His specialty would be familiarity with the special manuals compiled from what the theoreticians, equipment builders, and technicians know, and being able to pinpoint relevant data and apply complex rules and specifications.

”A lot of questions you might ask he couldn't answer directly, but in such a case he could often lead you quickly to some relevant pages in his books. You discovered that usually a succession of well-chosen questions of the sort he could answer, interspersed with your occasional study of succinct and relevant material he'd dig up for you, could very rapidly develop answers to conceptually sophisticated questions. His help in your minute-by-minute designing work could be extremely valuable -- availing you of quick and realistic consideration of a great many new design possibilities.

”Similarly, when it came to carrying out a planned set of operations, it turned out that he couldn't carry out all of the processes for you -- he could manage complex rules and procedures beautifully, but he would break down when it came to steps that required what you might call a larger view of the situation. But this wasn't so bad. The set of routine processes which he could manage all alone still provided you with a great deal of help -- in fact, you got to developing ways to build things so as to capitalize upon his efficiency at these tasks. Then the processes which were too much for him would be done by the two of you together. He filled in all the routine stuff and you took care of the steps that were beyond his capability Often the steps you had to take care of were buried in the middle of a complex routine whose over-all nature didn't have to be understood by either of you for

proper execution. Your helper would keep track of the complex procedure and execute all the steps he could. When he came to a step that was too big for him, he would hand you enough information to allow you to take that step, whereupon he would take over again until he met another such step.

”As an inventor and builder of devices that solve needs, you could become a great deal more versatile and productive, applying your imagination, intuition, judgment, and intelligence very effectively over a much wider range of possibilities. You could tackle much more complex and sophisticated projects, you could come up with very much better results -- neater, cheaper, more reliable, more versatile, higher-quality performance -- and you could work faster. Your effectiveness in this domain of activity would be considerably increased.

”So let's turn back to the working domain which we are considering here. It is an intellectual one, where the processing and shaping done is of conceptual material rather than physical material. But between these two types of working domains we nonetheless find closely analogous conditions relative to the variety and sophistication of the processes and techniques applicable to what nonroutine workers do. Consider the intellectual domain of a creative problem solver, and listen to me rattle off the names of some specialized disciplines that come to mind. These esoteric disciplines could very possibly contribute specialized processes and techniques to a general worker in the intellectual domain: Formal logic -- mathematics of many varieties, including statistics -- decision theory -- game theory -- time and motion analysis -- operations research -- classification theory -- documentation theory -- cost accounting, for time, energy, or money -- dynamic programming -- computer programming. These are only a few of the total, I'm sure.

”This implies the range of potentially applicable processes. Realize that there is also a correspondingly large list of specialized materials potentially usable in the fabrications of the intellectual worker. I speak, of course, about the conceptual material in the many different fields of human interest. The things that I have been demonstrating to you this afternoon were designed to increase significantly the range of both processes and materials over which a human can practically operate within this intellectual domain. You might say that we do this by providing him with a very fast, agile vehicle, equipped with all sorts of high-performance sensory equipment and navigational aids, and carrying very flexible, powerful, semi-automatic devices for operating upon the materials of this domain. Not only that, but to provide an accurate analogy, we have to give him a computer to help him organize and monitor his activity and assess his results. We get direct help on many levels of activity in our system, you see.

”But back to the topic of tools, and the analogy of the inventor who was given the equipment and the helper. Our augmented intellectual worker gets essentially this same kind of service, only more so -- a compounding of this kind of service. Structuring our processes with care and precision

enables the computer to answer limited questions, to guide you to relevant descriptions and specifications within its structure, to execute complex but limited-grasp processes on its own, and to take care of complex rule and procedure-following bookkeeping in guiding the execution of sophisticated composite processes. This actually makes it practical to use many specialized processes and techniques from very esoteric fields -- to assess their applicability and limitations quickly, to incorporate them intelligently into the design and analysis of possible courses of action, and to execute them efficiently.

"Our specialized processes represent a beautiful collection of special tools. These tools are designed by specialists, and they come equipped with operating instructions, trouble-shooting hints, and complete design data. Furthermore, we are provided with other tools that help us determine the applicability of these tools by automatically operating upon the instruction manual for us. Further, if something goes wrong with one of these tools, if we want to design a new tool of our own and make use of one of its modular components, or if we want to rearrange some of its adjustable features, we get considerable help in learning what we have to know about its design, and in making adjustments or coupling a part of it to another tool. Our shop contains an efficient tool-making section, where we can design and build our own tools from scratch, or by incorporating parts or all of any other tools we have.

"Let me tell you of an interesting feature stemming from my using such improved Process-structuring techniques. An effective job of breaking down a complex problem into humanly manageable steps -- and this is essentially what we seek in our process structuring -- will provide the human with something to do at every turn. This may be to ponder or go searching, true enough - - we aren't saying that the steps are necessarily straightforward. But the point I want to make is that no longer am I ever at a loss as to what to do next. I get stuck at times, to be sure, but when I do I have clean and direct ways to satisfy myself that I should just beat away at that roadblock for the time being.

"And then, for beating away at the roadblock, my bookkeeping regarding what I've tried, what possibilities I've collected, and what my assumptions and objectives are, is good enough to help tremendously in keeping me from getting into loops and quandaries, in carefully exhausting possibilities, and in really analyzing my assumptions and objectives. What's more, I'm not generating reams of cyclic arguments, lists, calculations, or the like -- either I'm checking the validity of what I've already structured, or I am correcting or expanding the structure. In other words, it seems that the growth of my comprehension is sure and steady up to the point at which I succeed or give up. If I give up, I leave a structure which is very well organized to accommodate a subsequent revisit with new data, possibilities, assumptions, objectives, or tools. Also, I set up a

sentinel process that will operate in the future to help alert me to concepts which may clear the block.

"This feature, of always having satisfying actions to perform, and having a good feeling that they are what I should be doing at that time, gives a surprisingly contented, eager, and absorbing flavor to my work. I guess it's an adult instance of the sort of change observed in students when they were given teaching machines that provided continuous participation and reinforcement.

"Anyway, with the quick flexibility available to me for structuring arguments, and semi-automatic application of special tagging and linking rules, I find it really quite easy to construct, use, or modify sophisticated process structuring. And I can turn right around and apply this toward improving my ability for structuring arguments and processes. The initial, straightforward capabilities for manipulating symbol structures, that were more or less obviously availed me by the computer have given to me a power to participate in more sophisticated processes that capitalize more fully upon the computer's capability -- processes which are very significant to my net effectiveness, and yet which weren't particularly apparent to us as either possible or useful in the days before we started harnessing computers to the human's workaday activities in this direct way.'

Team Cooperation

"Let me mention another bonus feature that wasn't easily fore seen. We have experimented with having several people work together from working stations that can provide inter-communication via their computer or computers. That is, each person is equipped as I am here, with free access to the common working structures. There proves to be a really phenomenal boost in group effectiveness over any previous form of cooperation we have experienced. They can all work on the same symbol structure, wherever they might wish. If any two want to work simultaneously on the same material, they simply duplicate and each starts reshaping his version -- and later it is easy to merge their contributions. The whole team can join forces at a moments notice to 'pull together' on some stubborn little problem, or to make a group decision. Most points of contention are resolved quite naturally, over a period of time, as the developing structure of argument bears out one, or the other, or neither stand.

"No one can dominate the show, since seldom do you have to 'listen' to the person concurrent to the developments he is pursuing -- and yet at any time another person can tune in on what he has done and is doing. One can either take immediate personal issue with another about some feature, anywhere in the structure where he might find something done by the other to which he wants to take issue, or he can append his objection and the associated argument there where the disagreement lies, and tag this with a special cue that signals a point of contention that must ultimately be resolved. Any idea of the moment by any member can easily be linked to where it can do some

good. It gets to be like a real whing-ding free-for-all -- tremendously stimulating and satisfying, and things really get done. You find yourself 'playing over your head' almost all of the time.

"We have been experimenting with multi-disciplinary teams and are becoming especially excited over the results. For instance, there is a great reduction of the barrier that their different terminologies used to represent, where one specialist couldn't really apply his experience, intuition, or conceptual feel very well unless the situation could be stated and framed in his accustomed manner, and yet the others couldn't work with his terminology. Here, they meet at their concept and terminology interface and work out little shifts in meaning and use which each can find digestible in his system, and which permit quite precise definitions in each system of the terms and concepts in the others. In studying the other's structuring then, either of them can have his own definitions automatically substituted for the other's special terms. Reduce this language barrier, and provide the feature of their being able to work in parallel independence on the joint structure, and what seems to result is amplification of their different capabilities.

"Remember the term, synergesis, that has been associated in the literature with general structuring theory? Well, here is something of an example. Three people working together in this augmented mode seem to be more than three times as effective in solving a complex problem as is one augmented person working alone -- and perhaps ten times as effective as three similar men working together without this computer-based augmentation. It is a new and exhilarating experience to be working in this independent-parallel fashion with some good men. We feel that the effect of these augmentation developments upon group methods and group capability is actually going to be more pronounced than the effect upon individuals methods and capabilities, and we are very eager to increase our research effort in that direction."

Miscellaneous Advanced Concepts

"I have dragged you through a lot of different concepts and methods so far. I haven't been complete because we won't have the time. But I have selected the sample features to present to you with an eye toward giving you a maximum chance to identify these as being something significant to your own type of work. I avoided discussing techniques applicable to esoteric problem-solving processes -- although some of them display especially stimulating possibilities to those with appropriate backgrounds. The ability to structure arguments organized in English-language statements, and to make use of the linking and tagging capabilities at all levels of the structure, can be seen to lead to many interesting and promising new capabilities for organizing your thoughts and actions. I think you could picture learning these tricks and using them in your own work.

"What I hoped to avoid by presenting the system in this way, was losing your identification with these possibilities by letting you get the mistaken impression that an individual couldn't harness

these techniques usefully unless he first learned a lot of very sophisticated new language, logic and math. It is true that the more of the sophisticated tricks you learn, the more computer power you can harness and the more powerful you become -- but very significant and personally thrilling practical problem-solving capabilities have been developed by quite a few subjects who were given only fifteen hours of training at one of these stations. The training, incidentally, was all provided by the computer without the presence of a human instructor. And the people were of such diverse fields as sociology, biology, engineering management, applied mathematics, and law. These were all relatively high-level people, and they were completely and unreservedly unanimous in their faith that their increased capability would easily justify the capital and operating outlay that we predicted for work stations of this sort in five years, if the computer industry really were to take this type of potential market seriously.

”What these people became capable of was somewhat less than the range of capabilities that we have discussed so far -- but they would find it very natural to develop further techniques on their own, and new teaching programs could be provided them so that they could continue learning the improved techniques turned out by a research group such as ours here.

”But let me give you a brief view of some of the more advanced concepts and techniques that have evolved here, compatible with, but beyond, what I have so far shown you. And evolved is a good word to use here, because our appreciation for the potential worth of possibilities to be developed had to evolve too, and only came with the experience and perspective gained in our earlier work.

”For instance, we initially felt that defining categories and relationships, and making a plan for action, were things to be done as quickly as possible so that we could get on with the work. But, as our means developed for dealing with definitions and plans more precisely, easily, and flexibly, we began to realize that they in reality might be the most significant part of that work. With our immensely increased capability for complex bookkeeping relative to our interlaced hierarchies of objectives, plans, and arguments, we found that defining a new category, searching for members or instances of it, or applying its selection criteria were becoming ever conscious and specific tasks.

”For instance, we began to find it more and more useful to distinguish different categories or types of process, different types of arguments, different types of relationships, and different types of descriptions. For a specific example, Ranganathan¹ once cited five specific relationships that could obtain between two terms, where one modifies the other. He called these phase relations, and named how one term could relate to the other as either biasing it, being a tool used to study it, being an aspect of it, being in comparison with it, or influencing it. Vickery gave more examples, saying one could also have an effect on the other, be a cause of it, be a use for it, be a substitute for it, a source

for it, an implication of it, be an explanation of it, or be a representation of it. There are even more categories mentioned in the literature.

”It was easy to form tags and links, and we experimented with the gains to be made by consciously specifying and indicating categories. It turned out to be a very invigorating innovation, and we began to take more pains with our structuring. It took longer to set up links and nodes in our structures, to be sure, but we found on the one hand that the structures became much cleaner and required fewer members, and on the other hand that we could get considerably more sophisticated help from the computer in doing significant chores for us.

”We began to work up processes that would help us establish categories, give them good definitions, check their relationship with other established categories, decide whether something fit a given category or not, search for all possible members of it within a given substructure, and so forth. The very fact of using this careful classification within our structures allowed us to get more powerful help from the computer in these classification processes. I should mention that the relationships among the terms in a sentence -- the syntax if you wish -- had been given further specification tags than those I showed you earlier, to remove ambiguities that hindered the computer from going back to a statement and resolving the syntactical structure. Also, ambiguities in the meaning of the terms began to limit us, and we developed methods for removing a good deal of this semantic ambiguity. This slowed us down, as I've mentioned, but not as much as you'd think.

”Let me demonstrate one of the advanced processes which has evolved. It is heavily dependent upon the very care in building structures that it so nicely facilitates, and also upon several other developments. One of these other developments stems from the concepts and techniques of the semantic differential, as first introduced by Osgood, Suci, and Tannenbaum² back in 1957, and from some subsequent work by Mayer and Bagley³ on what they called semantic models. These offered useful possibilities for establishing quite precisely what meaning a concept has to an individual, relative to his general conceptual framework, and for representing this meaning in a specific way that was amenable to computer manipulation.

”The other development upon which this process to be exhibited is based, was stimulated by our realizing that flexible cooperation with the computer was calling for lots of little interactions. Our working repertoire of small-task requests for computer service was getting quite large, and it was proving to be extremely valuable to use them and to be able to remember automatically their procedures and designation codes. One of our research psychologists had worked on human-memory phenomena before he came with us, and had interested himself in mnemonic aids of all sorts. He has developed some useful techniques for us to use in connection with this, and other

problems. Now let me demonstrate this example of an advanced process for helping work with categories.

”Suppose that I want to establish a new category. Let's say that I have developed its description in what you and I have been calling an argument structure. I want to give it a name -- a short and meaningful one -- and I want a good definition. In fact, I want a definition that the computer can later work with. Look, I'll dig up a description that is awaiting such a definition, and you can watch what happens.” So saying, Joe drummed on his keysets for a moment, with one interruption when the computer flashed something on the screen that was apparently a question about what he was asking the computer to find for him. He finally had a network display on one screen and a set of “exploded” statements on the upper half of the other.

”I'm initiating the naming and defining process now, and designating to it the argument structure represented by this network as what I want named and defined. Watch what happens.” A few more strokes on the keyset, and he picked up his light pen in anticipation and waited a few moments. A statement appeared in the lower half of the second frame. He studied it a moment, then looked at the statements above, picked out a node on the network with the pen, and hit the keyset a few strokes. Another statement flashed on almost immediately, with two familiar adjectives placed below and a graduated line between them. Joe studied this, referred to the statements above, flipped through several levels of network portrayals, through a few statements representing a couple of low-level nodes, reflected a moment, and then pointed his light pen at a point on the graduated line, part way between the adjectives, and pressed its button.

”Actually, right now I'm demonstrating a cooperative process execution technique. This process is applying some very sophisticated criteria and using some very sophisticated analytical techniques, and it is set up so that it is actually the computer that is now in the executive seat. I called for the process, but its execution essentially involves the computer's asking me questions, and feeding me successive questions according to how I've answered the previous ones. It also is doing a lot of work on the symbol structure that represents my description. It, with some small help from me, is proceeding through a quite complex analysis of the meaning that this incipient concept has to me, and of certain types of mental associations that I may have with it. I don't have to remember the special rules and forms of analysis involved -- nevertheless, a very sophisticated little capability is mine to use at will, taxing neither me nor the computer.”

After a little over a minute of these question-answer interactions, the process apparently terminated, with four lines of special terms remaining on the screen. “This first line gives me two suggested names for this category or concept. The first term is a newly coined formal name, while the remaining three terms represent a compound expression, involving established concepts, that can be

used also as a designation of the new category. The second line furnishes me with an association chain to use for a mnemonic aid in remembering the new name -- linking the name to several characteristics of the concept. The name itself was selected under mnemonic criteria, as well as to have a structure that goes with its syntactic and semantic categories. The third line lists the names of some previously defined categories or concepts that are the closest to this in meaning -- these before the break were found to overlap, and the rest are just close.

"The fourth line you recognize as a statement form, perhaps. This is the definition, as developed by the computer. It's in a special language and I won't try to explain. I'll just mention that I can now study it, take it apart, check its references, so to speak, and perhaps even see if the computer and I might work out any changes or improvements. But this process has been worked on pretty hard, and we're getting definitions that are hard to improve.

"This special language, in which I said the definition was stated, is a recent development. We had found that the types of structuring we were developing had a lot of extra tags and links that were traceable to the complexity of the rules and combinatorial possibilities of the English language with which the statements were constructed. We finally got a clear enough picture of the requirements we place upon a language in our use here that we could consider designing our own special language. It turned out to be a straightforward and rather simple language compared with English, but much more precise and powerful. It proves rather inflexible and awkward to use for speaking, but it provides plenty of flexibility and power for expressing things in the visual-symbol forms that we use. Its precision leaves no syntactic ambiguity in a well-formed statement, and makes it much easier to reduce semantic ambiguity to the point where the computer can deal with our statements much as it can with mathematical or formal-logic expressions.

"It is worth mentioning, too, that we are experimenting with standard ways of structuring arguments at levels higher than the statements -- sort of a super grammar or syntax, with rules for assembling argument modules of different function into what becomes a well-formed higher-level argument module. There are some mixed feelings around here about this possibility, but I myself have become very much excited by it.

"Also we have been introducing formal methods for manipulating what you might call reasonable statements -- as opposed to absolute true/false statements which the more familiar formal logic can manipulate. This finds approval and faith in all of us here, but it is going a bit slowly.

"Let's run over some of the results we've seen to date, stemming from this new language and the new semantic awareness thus given the computer. If it can get hold of and manipulate important aspects of the meaning that is contained in our structures, it can develop answers to some questions

for which there existed only conceptually implicit data. With practice and good strategy, asking questions like this proves to be a tremendously effective way to gain comprehension about a structure. We even have special processes and symbol-structuring methods to help organize the questioning and the answers. Some of the answers are a bit costly, however -- in computer time and charges -- and we have to watch the way we ask questions. Some of our researchers are studying the language and structuring techniques relative to this problem, and they think they see ways to change them to make question answering generally more efficient. But this sort of thing will likely always have its cost problems, as far as we can see now.”

He went on to say that the computer now represents such an intelligent helper -- although much less so than any human helper they would hire -- that they refer to it as the Clerk. They can make a tentative new statement in the development of a structure, and have the clerk look over the structure to detect inconsistency or redundancy. The Clerk can also point out some of the weaknesses in the statement, as well as some of the effects of the statement upon the rest of the structure. They find that they need to give less and less human concern for the details of structure building -- in-fact, the roles have reversed a little. Where the human used to set up tags and links so the computer could find its way around the structure as it ran errands for him, they now have the computer studiously installing similar things that are for the benefit of the human when he is studying the structure.

He also mentioned a recently developed computer process that could go back over a record of the human actions involved in establishing a given argument structure and do a creditable job of picking out the steps which contributed the most to the final picture -- and also some of those that contributed least. This process, and some of the past data collected by its use, were becoming an important addition to the planning review sessions, as well as to the continuing development of improved methods. And apparently, it had a surprisingly positive psychological effect upon members of a cooperating team, where an objective means of relative scoring was thus available.

- - - Let yourself be disengaged now from your role in the above discussion-demonstration. You have been through an experience that was designed to give you a feel for the sort of future developments that (to us) are predictable from our conceptual framework. What is presented in Section II is an attempt at giving a “straight” presentation of the various conceptual segments of this framework, and Section III hopefully supplemented the formal presentation to provide you with a more complete picture of how we are oriented and what sorts of possibilities impel us.

Assuming that we have communicated our conceptual framework in some reasonable form, we proceed below to discuss the question of what to do about it. Our approach to this question is with the view that energetic pursuit of this research could be of considerable significance to society, and that research should stem from a big enough picture of the over-all possibilities so that the

contribution of any program, large or small, could have maximum long-range significance. Our recommendations are fairly general, and are cast in rather global terms, but we assert that they can be readily recast into the specific terms required of research planning to be done for a given project, within a given set of sub-goals and research-activity constraints. In fact, we are now engaged in the process of so recasting these general recommendations into specific plans (for the experimental research to be pursued here at Stanford Research Institute).

RESEARCH RECOMMENDATIONS

OBJECTIVES FOR A RESEARCH PROGRAM

The report has put forth the hypothesis that the intellectual effectiveness of a human being is dependent upon factors which are subject to direct redesign in pursuit of an increase in that effectiveness. A conceptual framework is offered to help in giving consideration to this hypothesis, and an extensive and personalized projection into possible future developments is presented to help develop a feeling for the possibilities and promise implicit in the hypothesis and conceptual structure.

If this hypothesis and its glowing extrapolations were borne out in future developments, the consequences would be most exciting and assumedly beneficial to a problem-laden world. What is called for now is a test of this hypothesis and a calibration on the gains if any that might be realized by giving total-system design attention to human intellectual effectiveness. If the test and calibration proved to be favorable, then we can set to work developing better and better augmentation systems for our problem solvers.

In this light, we recommend a research program approach aimed at (Goal 1) testing the hypothesis, (Goal 2) developing the tools and techniques for designing better augmentation systems, and (Goal 3) producing real-world augmentation systems that bring maximum gains over the coming years to the solvers of tough, critical problems. These goals and the resulting design for their pursuit are idealized, to be sure, but the results nonetheless have valuable aspects.

BASIC RESEARCH CONDITIONS

This should be an empirical approach on a total-system basis -- i.e., doing coordinated study and innovation, among all the factors admitted to the problem, in conjunction with experiments that provide realistic action and interplay among these variables. The question of limiting these factors is considered later in the section. The recommended environment for this empirical, total-system approach, is a laboratory providing a computer-backed display and communication system of the general sort described in Section III-B. There should be no stinting on the capabilities provided -- it is very important to learn what value any given artifact feature may offer the total system, and the only way to learn the value is to experiment with the feature. At this point no time will be taken to develop elaborate improvements in the art of time sharing, to provide real-time service to many users. This kind of development should be done as separate, backup work. The experimental lab should take the steps that are immediately available to provide all the service to the human that he needs in the experimental environment.

Where economy demands that a computer not be idle during the time the augmented subject is not using it (which would be a rather large net fraction of the time, probably), and where sharing the computer with other real-time users for which demand delays are a problem, then the only sharing that should be considered is that with off-line computations for which there are no real-time service demands to be met. The computer can turn away from off-line users whenever the on-line worker needs attention of any sort.

WHOM TO AUGMENT FIRST

The experimental work of deriving, testing, and integrating innovations into a growing system of augmentation means must have a specific type of human task to try to develop more effectiveness for, to give unifying focus to the research. We recommend the particular task of computer programming for this purpose -- with many reasons behind the selection that should come out in the following discussion. Some of the more direct reasons are these:

The programmer works on many problems, including large and realistic ones, which can be solved without interaction with other humans. This eases the experimental problem.

Typical and realistic problems for the programmer to solve can be posed for experimental purposes that do not involve large amounts of working and reference information. This also eases the experimental problem.

Much of the programmers working data are computer programs (he also has, we assume, his own reasoning and planning notes), which have unambiguous syntactic and semantic form so that getting

the computer to do useful tasks for him on his working data will be much facilitated -- which helps very much to get early experience on the value a human can derive from this kind of computer help.

A programmer's effectiveness, relative to other programmers, can probably be measured more easily than would be the case for most other complex-problem solvers. For example, few other complex solutions or designs beside a program can so easily be given the rigorous test of 'Does it actually work?'

The programmer's normal work involves interactions with a computer (although heretofore not generally on-line), and this will help researchers use the computer as a tool for learning about the programmer's habits and needs.

There are some very challenging types of intellectual effort involved in programming. Attempting to increase human effectiveness therein will provide an excellent means for testing our hypothesis.

Successful achievements in evolving new augmentation means which significantly improve a programmer's capability will not only serve to prove the hypothesis, but will lead directly to possible practical application of augmentation systems to a real-world problem domain that can use help.

Computer programmers are a natural group to be the first in the real world to incorporate the type of augmentation means we are considering. They already know how to work in formal methodologies with computers, and most of them are associated with activities that have to have computers anyway, so that the new techniques, concepts, methods, and equipment will not seem so radical to them and will be relatively easy for them to learn and acquire.

Successful achievements can be utilized within the augmentation-research program itself, to improve the effectiveness of the computer programming activity involved in studying and developing augmentation systems. The capability of designing, implementing, and modifying computer programs will be very important to the rate of research progress.

Workers in an augmentation-research laboratory are the most natural people in the world to be the very first users of the augmentation means they develop, and we think that they represent an extremely important group of people to make more effective at their work.

BASIC REGENERATIVE FEATURE

The feature brought forth in Reason 9 above is something that offers tremendous value to the research objectives -- i.e., the feeding back of positive research results to improve the means by

which the researchers themselves can pursue their work. The plan we are describing here is designed to capitalize upon this feature as much as possible, as will be evident to the reader as he progresses through this section. This positive-feedback (or regenerative) possibility derives from the facts that: (1) our researchers are developing means to increase the effectiveness of humans dealing with complex intellectual problems, and (2) our researchers are dealing with complex intellectual problems. In other words, they are developing better tools for a class to which they themselves belong. If their initial work needs the unifying focus of concentrating upon a specific tool, let that tool be one important to them and whose improvement will really help their own work.

TOOLS DEVELOPED AND TOOLS USED

This close similarity between tools being developed and the tools being used to do the developing, calls for some care in our terminology if we want to avoid confusion in our reasoning about their relationship. 'Augmentation means' will be used to name the tools being developed by the augmentation research. 'Subject information' will be used to refer to description and reasoning concerned with the subject of these tools (as opposed to the method of research), and 'subject matter' will refer to both subject information and physical devices being incorporated as artifacts in the augmentation means being developed. 'Tools and techniques' will be used to name the tools being used to do that research, and are likely here to include special additions to language, artifact, and methodology that particularly improve the special capabilities exercised in doing the research.

An integrated set of tools and techniques will represent an art of doing augmentation research.

Although no such art exists ready-made for our use, there are many applicable or adaptable tools and techniques to be borrowed from other disciplines. Psychology, computer programming and physical technology, display technology, artificial intelligence, industrial engineering (e.g., motion and time study), management science, systems analysis, and information retrieval are some of the more likely sources. These disciplines also offer initial subject matter for the research. Because this kind of diagramming can help more later on, we represent in Fig. 3 the situation of the beginning research drawing upon existing disciplines for subject matter and tools and techniques.

The program begins with general dependence upon other, existing disciplines for its subject matter (solid arrow) and its tools and techniques (dashed arrow). Goal 1 has been stated as that of verifying the basic hypothesis that concerted augmentation research can increase the intellectual effectiveness of human problem solvers.

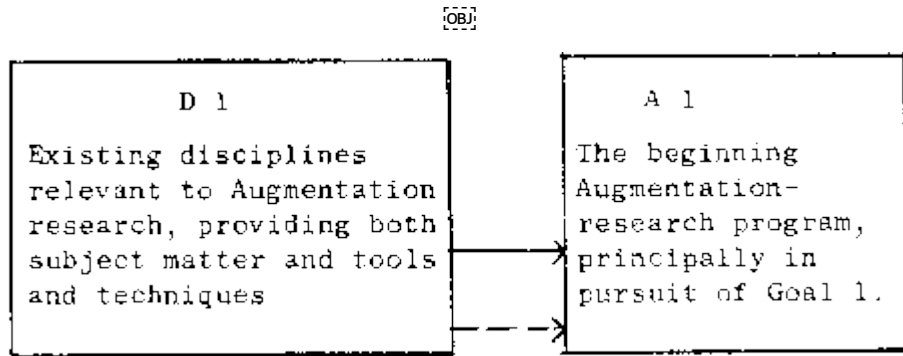


Fig. 3 Initial Augmentation-Research Program

RESEARCH PLAN FOR ACTIVITY A 1

The dominant goal of Activity A 1 (Goal 1, as in Fig 3) is to test our hypothesis. Its general pursuit of augmenting a programmer is designed to serve this goal, but also to be setting the stage for later direct pursuit of Goals 2 and 3 (i.e., developing tools and techniques for augmentation research and producing real-world augmentation systems).

Before we discuss the possible subject matter through which this research might work, let us treat the matter of its tools and techniques. Not too long ago we would have recommended (and did), in the spirit of taking the long-range and global approach, that right from the beginning of a serious program of this sort there should be established a careful and scientific methodology. Controlled experiments, with special re search subjects trained and tested in the use of experimental new augmentation means, careful monitoring, record-keeping, and evaluative procedures, etc. This was to be accompanied by a thorough search through disciplines and careful incorporation of useful findings.

Still in the spirit of the long-range and global sort of planning, but with a different outlook (based, among other things, upon an increased appreciation for the possibilities of capitalizing upon regeneration), we would now recommend that the approach be quite different. We basically recommend A 1 research adhering to whatever formal methodology is required for (a) knowing when an improvement in effectiveness has been achieved, and (b) knowing how to assign relative value to the changes derived from two competing innovations.

Beyond this, and assuming dedication to the goal, reasonable maturity, and plenty of energy, intelligence, and imagination, we would recommend turning loose a group of four to six people (or a number of such groups) to develop means that augment their own programming capability We would recommend that their work begin by developing the capability for composing and modifying

simple symbol structures, in the manner pictured in Section III-B-2, and work up through a hierarchy of intermediate capabilities toward the single high-level capability that would encompass computer programming. This would allow their embryonic and free wheeling “art of doing augmentation research” to grow and work out its kinks through a succession of increasingly complex system problems -- and also, redesigning a hierarchy from the bottom up somehow seems the best approach

As for the type of programming to tell them to become good at -- tell them, “the kind that you find you have to do in your research.” In other words, their job assignment is to develop means that will make them more effective at doing their job. Figure 4 depicts this schematically, with the addition to what was shown in Fig. 3 of a connection that feeds the subject-matter output of their research (augmentation means for their type of programming problems) right back into their activity as improved tools and techniques to use in their research.




Fig. 4 Regeneration

If they are making head way, it won't take any carefully worded criterion of effectiveness nor any great sophistication in measurement technique to tell that they are more effective with the augmentation means than without -- being quicker to “design and build” a running program to meet given processing specifications or being quicker to pick up a complex existing program, gain comprehension as necessary, and find its flaws or rebuild it. On the other hand, if no gains are really obvious after a year or so, then it is time to begin incorporating more science in their approach. By then there will be a good deal of basic orientation as to the nature of the problem to which “science” is to be applied.

What we are recommending in a way is that the augmented capability hierarchy built by this group represent more a quick and rough scaffolding than a carefully engineered structure. There is orientation to be derived from climbing up quickly for a look that will be of great value. For instance, key concepts held initially, that would have been laboriously riveted into the well-

engineered structure, could well be rendered obsolete by the “view” obtained from higher in the hierarchy. And besides, it seems best to get the quick and rough improvements built and working first, so that the research will benefit not only from the orientation obtained, but from the help that these improvements will provide when used as tools and techniques to tackle the tougher or slower possibilities. As progress begins to be made toward Goal 1, the diagram of Fig. 3 will become modified by feeding the subject-matter output (augmentation means for computer programmers) back into the input as new tools and techniques to be used by the researchers.

We would suggest establishing a sub-activity within A 1, whose purpose and responsibility is to keep an eye on the total activity, assess and evaluate its progress and try to provide orientation as to where things stand and where attention might be beneficial.

A few words about the subject matter through which Activity A 1 may progress. The researchers will think of simple innovations and try them in short order -- and perhaps be stimulated in the process by realizing how handy some new feature would be that would help them whip up trial processes in a hurry. They will know of basic capabilities they want to work toward for structuring their arguments their planning, their factual data, etc., so that they can more easily get computer help in developing them in analyzing and pursuing comprehension within them and in modifying or extending them. They will try different types of structuring and see how easy it is to design computer processes to manipulate them or composite processes to do total useful work with them.

They can work up programs that can search through other programs for answers to questions about them -- questions whose answers serve the processes of debugging, extending, or modifying. Perhaps there will be ways they adopt in the initial structuring of a program -- e.g., appending stylized descriptive cues here and there -- that have no function in the execution of that program, but which allow more sophisticated fact retrieval therein by the computer. Perhaps such cue tagging would allow development of programs which could automatically make fairly sophisticated modifications to a tagged program. Maybe there would evolve semi-automatic “super-compilers,” with which the programmer and the computer leap-frog over the obstacles to formulating exact specifications for a computer (or perhaps composite) process and getting it into whatever programming language they use.

A Second Phase In The Research Program

The research of A 1 could probably spiral upwards indefinitely, but once the hypothesis (see Section IV-A) has been reasonably verified and the first of our stated objectives satisfied, it would be best to re-organize the program. To describe our recommendation here, let us say that two research activities, A 2 and A 3, are set up in place of A 1. Whether A 1 is split, or turned into A 2

and a new group formed for A 3, does not really matter here -- we are speaking of separate activities, corresponding to the responsible pursuit of separate goals, that will benefit from close cooperation.

To Activity A 2 assign the job of developing augmentation means to be used specifically as tools and techniques by the researchers of both A 2 and A 3. This establishes a continuing pursuit for Objective 2 of Section IY-A. A 2 will now set up a sub-activity that studies the problems of all the workers in A 2 and A 3 and isolates a succession of capabilities for which the research of A 2 will develop means to augment. Activity A 2 should be equipped with the best artifacts available to an experimental laboratory.

To Activity A 3 assign the job of developing augmentation systems that can be practically adopted into real-world problem situations. This provides a direct and continuing pursuit of Goal 3 of Section IV-A. It is to be assumed that the first real-world system that A 3 will design will be for computer programmers. For this it might well be able to clean up the "laboratory model" developed in A 1, modify it to fit the practical limitations represented by real-world economics, working environments, etc., and offer it as a prototype for practical adoption. Or Activity A 3 might do a redesign, benefitting from the experience with the first model.

Activity A 3 will need a sub-activity to study its potential users and guide the succession of developments that it pursues. Activity A 2 in its continued pursuit of increased effectiveness among workers in idealized environment, will be the source for basic subject matter in the developments of A 3, as well as for its tools and techniques. From the continuously expanding knowledge and developments of A 2, A 3 can organize successive practical systems suitable for ever more general utilization.

We have assumed that what was developed in A 1 was primarily language and methodology, with the artifacts not being subject to appreciable modification during the research. By this second phase, enough has been learned about the trends and possibilities for this type of on-line man-computer cooperation that some well-based guidance can be derived for the types of modifications and extensions to artifact capability that would be most valuable. Activity A 2 could continue to derive long-range guidance for equipment development, perhaps developing laboratory innovations in computers, display systems, storage systems, or communication systems, but at least experimenting with the incorporation of the new artifact innovations of others.

An example of the type of guidance derived from this research might be extracted from the concepts discussed in Section-C-5 (Structure Types). We point out there that within the computer there might be built and manipulated symbol structures that represent better images of the concept structures of

interest to the human than would any symbol structure with which the human could work directly. To the human, the computer represents a special instrument which can display to him a comprehensible image of any characteristic of this structure that may be of interest. From our conceptual viewpoint, this would be a source of tremendous power for the human to harness, but it depends upon the computer being able to 'read' all of the stored information (which would be in a form essentially incomprehensible to a human). Now, if this conjecture is borne out there would be considerably less value in micro-image information-storage systems than is now generally presumed. In other words, we now conjecture that future reference information will be much more valuable if stored in computer-sensible form. The validity of this and other conjectures stemming from our conceptual framework could represent critical questions to manufacturers of information systems.

It is obvious that this report stems from generalized "large-view" thinking. To carry this to something of a final view, relative to the research recommendations, we present Fig. 5, which should be largely self-explanatory by this time. Activity A 2 is lifting itself by the bootstraps up the scale of intellectual capability, and its products are siphoned to the world via A 3. Getting acceptance and application of the new techniques to the most critical problems of our society might in fact be the most critical problem of all by then, and Activity A 4 would be one which should be given special help from A 3.

There is another general and long-range picture to present. This is in regard to a goal for a practically usable system that A 3 would want to develop as soon as possible. You might call this the first general Computer Augmentation System -- CAUG-I (pronounced "cog-one").

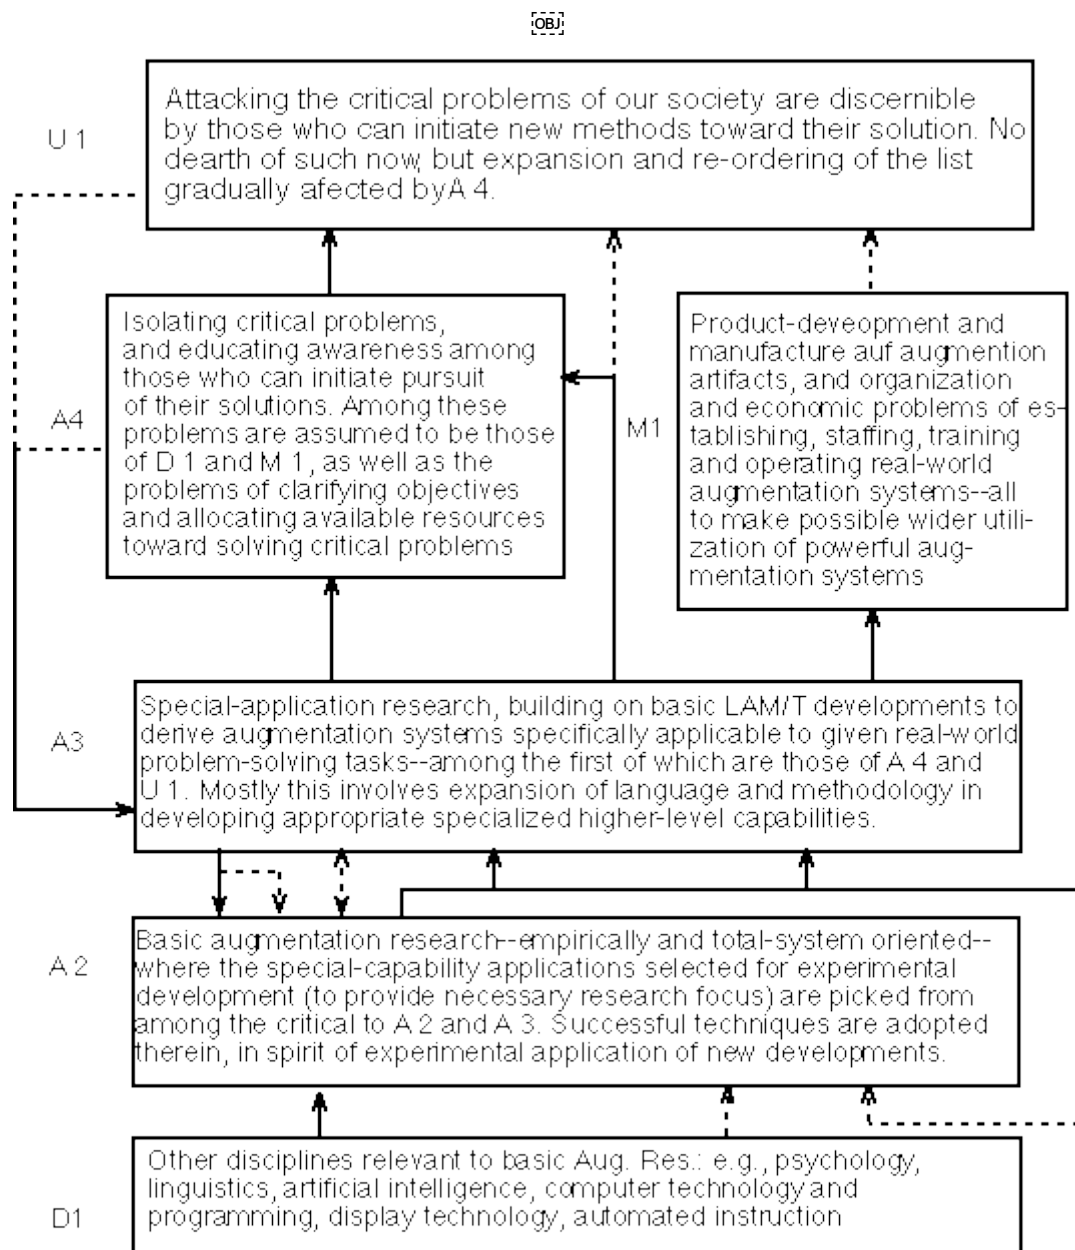


Fig. 5: A Total Program

Suggested relationship among the major activities in achieving the stated objective (essentially, of significantly boosting human power in A 4 and U 1). Solid lines represent subject information or artifacts used or generated within an activity, and dashed lines represent special tools and techniques for doing the activity in the box to which they connect. Subject product of an activity (output solid) can be used as working material (input solid) or as tools and techniques (input dashed). Tools and techniques as used or needed in an activity (output dashed) can be used either to work on (input solid) or as tools and techniques to work with (input dashed).

It would be derived from what was assessed to be the basic set of capabilities needed by both a general-problem-solving human and an augmentation researcher. Give CAUG-I to a real-world problem solver in almost any discipline, and he has the basic capabilities for structuring his arguments and plans, organizing special files, etc., that almost anyone could expect to need. In addition to these direct-application on capabilities, however, are provided those capabilities necessary for analyzing problem tasks, developing and evaluating new process capabilities, etc., as would be required for him to extend the CAUG-I system to match to the special features of his problem area and the way he likes to work.

In other words, CAUG-I represents a basic problem-solving tool kit, plus an auxiliary tool-makers tool kit with which to extend the basic tool kit to match the particular job and particular worker. In subsequent phases, Activity A 3 could be turning out successive generations (CAUG-II, CAUG-III, etc.) each incorporating features that match an ever-more-powerful capability hierarchy in an ever-more-efficient manner to the basic capabilities of the human.

SUMMARY

This report has treated one over-all view of the augmentation of human intellect. In the report the following things have been done:

An hypothesis has been presented.

A conceptual framework has been constructed.

A “picture” of augmented man has been described.

A research approach has been outlined. These aspects will be re viewed here briefly:

An hypothesis has been stated that the intellectual effectiveness of a human can be significantly improved by an engineering-like approach toward redesigning changeable components of a system.

A conceptual framework has been constructed that helps provide a way of looking at the implications and possibilities surrounding and stemming from this hypothesis. Briefly, this framework provides the realization that our intellects are already augmented by means which appear to have the following characteristics:

The principal elements are the language artifacts, and methodology that a human has learned to use.

The elements are dynamically interdependent within an operating system.

The structure of the system seems to be hierarchical, and to be best considered as a hierarchy of process capabilities whose primitive components are the basic human capabilities and the functional capabilities of the artifacts -- which are organized successively into ever-more-sophisticated capabilities.

The capabilities of prime interest are those associated with manipulating symbols and concepts in support of organizing and executing processes from which are ultimately derived human comprehension and problem solutions.

The automaton of the symbol manipulation associated with the minute-by-minute mental processes seems to offer a logical next step in the evolution of our intellectual capability.

A picture of the implications and promise of this framework has been described, based upon direct human communication with a computer. Here the many ways in which the computer could be of service, at successive levels of augmented capability, have been brought out. This picture is fanciful, but we believe it to be conservative and representative of the sort of rich and significant gains that are there to be pursued.

An approach has been outlined for testing the hypothesis of Item (1) and for pursuing the "rich and significant gains" which we feel are promised. This approach is designed to treat the redesign of a capability hierarchy by reworking from the bottom up, and yet to make the research on augmentation means progress as fast as possible by deriving practically usable augmentation systems for real-world problem solvers at a maximum rate. This goal is fostered by the recommendation of incorporating positive feedback into the research development -- i.e., concentrating a good share of the basic-research attention upon augmenting those capabilities in a human that are needed in the augmentation-research workers. The real-world applications would be pursued by designing a succession of systems for specialists, whose progression corresponds to the increasing generality of the capabilities for which coordinated augmentation means have been evolved. Consideration is given in this rather global approach to providing potential users in different domains of intellectual activity with the basic general-purpose augmentation system from which they themselves can construct the special features of a system to match their job, and their ways of working -- or it could be used on the other hand by researchers who want to pursue the development of special augmentation systems for special fields.

CONCLUSIONS

Three principal conclusions may be drawn concerning the significance and implications of the ideas that have been presented.

First any possibility for improving the effective utilization of the intellectual power of society's problem solvers warrants the most serious consideration. This is because man's problem-solving capability represents possibly the most important resource possessed by a society. The other contenders for first importance are all critically dependent for their development and use upon this resource. Any possibility for evolving an art or science that can couple directly and significantly to the continued development of that resource should warrant doubly serious consideration.

Second, the ideas presented are to be considered in both of the above senses: the direct-development sense and the 'art of development' sense. To be sure, the possibilities have long-term implications, but their pursuit and initial rewards await us now. By our view, we do not have to wait until we learn how the human mental processes work, we do not have to wait until we learn how to make computers more intelligent or bigger or faster, we can begin developing powerful and economically feasible augmentation systems on the basis of what we now know and have. Pursuit of further basic knowledge and improved machines will continue into the unlimited future, and will want to be integrated into the "art" and its improved augmentation systems -- but getting started now will provide not only orientation and stimulation for these pursuits, but will give us improved problem-solving effectiveness with which to carry out the pursuits.

Third, it becomes increasingly clear that there should be action now -- the sooner the better -- action in a number of research communities and on an aggressive scale. We offer a conceptual framework and a plan for action, and we recommend that these be considered carefully as a basis for action. If they be considered but found unacceptable, then at least serious and continued effort should be made toward developing a more acceptable conceptual framework within which to view the over-all approach, toward developing a more acceptable plan of action, or both.

This is an open plea to researchers and to those who ultimately motivate, finance, or direct them, to turn serious attention toward the possibility of evolving a dynamic discipline that can treat the problem of improving intellectual effectiveness in a total sense. This discipline should aim at producing a continuous cycle of improvements -- increased understanding of the problem, improved means for developing new augmentation systems, and improved augmentation systems that can serve the world's problem solvers in general and this discipline's workers in particular. After all, we

spend great sums for disciplines aimed at understanding and harnessing nuclear power. Why not consider developing a discipline aimed at understanding and harnessing “neural power?” In the long run, the power of the human intellect is really much the more important of the two.

REFERENCES

Kennedy, J. L and Putt, G. H., “Administration of Research in a Research Corporation,” RAND Corporation Report P-847 (20 April 1956).

Ashby, Ross, *Design For a Brain* (John Wiley & Sons, New York City, N. Y., 1960).

Ashby, Ross, “Design for an Intelligence-Amplifier, *Automata Studies*, edited by C. E. Shannon and J. McCarthy, pp. 215-234 (Princeton University Press, 1956).

Korzybski, A , *Science and Sanity*, 1st Ed. (International non Aristotelian Library Publishing Co., Lancaster, Pennsylvania, 1933).

Whorf, B. L., *Language, Thought, and Reality* (MIT & John Wiley & Sons, Inc., New York City, N.Y., 1956).

Bush, Vannevar, “As We May Think,” *The Atlantic Monthly* (July 1945).

Newell, A. (editor), *Information Processing Language-V Manual* (Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1961).

McCarthy, J., “LISP 1.5 Programmer's Manual,” Computation Center and Research Laboratory of Electronics, MIT (14 July 1961).

Gelernter, H., Hansen, J. R., and Gerberich, C. L., “A Fortran Compiled List-Processing Language,” *Journal of the Assoc. for Computing Machinery* (April 1960).

Yngve, V. H., “Introduction to COMIT Programming,” Technical Report, Research Laboratories of Electronics and Computation Center, MIT (5 November 1961).

Yngve, V. H., “COMIT Programmer's Reference Manual,” Technical Report, Research Laboratories of Electronics and Computation Center, MIT (5 November 1961).

Perlis, A. J. and Thornton, C., "Symbol Manipulation by Threaded Lists," Communications of the ACM, Vol. 3, No. 4 (April 1960).

Carr, J. W., III, "Recursive Subscripting Compilers and List-Type Memories," Communications of the ACM, Vol. 2, pp. 4-6 (February 1959).

Weizenbaum, J., "Knotted List Structures," Communications of the ACM, Vol. 5, No. 3, pp. 161-165 (March 1962).

Licklider, J. C. R., "Man-Computer Symbiosis," IRE Transactions on Human Factors in Electronics (March 1960).

Ulam, S. M., A Collection of Mathematical Problems, p. 135 (Interscience Publishers, Inc., New York, N.Y., 1960).

Good, I. J., "How Much Science Can You Have at Your Fingertips?" IBM Journal of Research and Development, Vol. 2, No. 4 (October 1958)

Ramo, Simon, "A New Technique of Education," IRE Trans. on Education (June 1958).

Ramo, Simon, "The Scientific Extension of the Human Intellect," Computers and Automation (February 1961).

Fein, Louis, "The Computer-Related Science (Synnoetics) at a University in the Year 1975," unpublished paper (December 1960).

Licklider, J. C. R. and Clark, W. E., "On-Line Man-Computer Communication," Proceedings Spring Joint Computer Conference, Vol. 21, pp. 113-128 (National Press, Palo Alto, California, May 1962).

Culler, G. J. and Huff, R. W., "Solution of Non-Linear Integral Equations Using On-Line Computer Control," Ramo-Wooldridge, Canoga Park, California, paper for presentation at SJCC, San Francisco (May 1962).

Teager, H. M., "Real-Time, Time-Shared Computer Project," report, MIT, Contract Nonr-1841(69) DSR 8644 (1 July 1961).

Teager, H. M., "Systems Considerations in Real-Time Computer Usage," paper presented at ONR Symposium on Automated Teaching (12 October 1961).

Randa, Glenn C., “ “

Design of a Remote Display Console,” Report ESL, R-132, MIT, Cambridge, Massachusetts (available through ASTIA) (February 1962).

Vickery, B. C., Classification and Indexing in Science, p. 42 (Academic Press, Inc., New York, 1959).

Osgood, C. E., Suci, G. J., and Tannenbaum, P. H., The Measurement of Meaning (University of Illinois Press, Urbana, Illinois, 1957).

Current Research and Development in Scientific Documentation No. 6 NSF-60-25, p. 104 (National Science Foundation, May 1960).

© *Doug Engelbart, 1962-2008.*

www.bootstrap.org