

NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:
The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

NOTES FOR A MODEL OF HUMAN PERFORMANCE IN ZOG

Allen Newell
August 5, 1977

Department of Computer Science
Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213

This work was supported by the Office of Naval Research under contract no. N00014-76-0874. It was also partially supported by the Defense Advanced Research Projects Agency under contract no. F44620-73-C-0074 and monitored by the Air Force Office of Scientific Research.

NOTES FOR A MODEL OF HUMAN PERFORMANCE IN ZOG¹

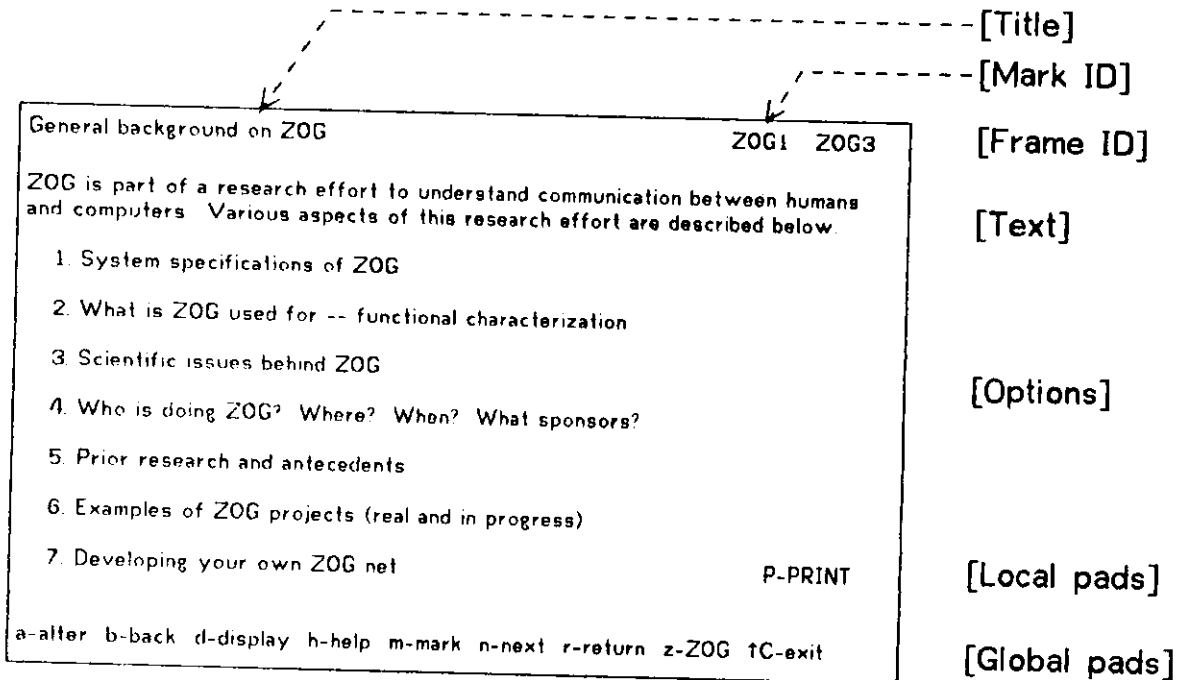
ZOG is an interface for man-computer communication currently under exploration at CMU (Robertson, Newell & Ramakrishna, 1977). This paper contains some initial considerations for how to analyze human performance using ZOG. We will give a brief introduction to ZOG as seen at the level of the user, in order to make this paper minimally self contained. The paper referenced above, which introduces ZOG and describes the system architecture, should be consulted for more background.

ZOG may be described as a rapid response, large network, menu selection technique. The user sits in front of video terminal on which is displayed a frame of information, as shown in Figure 1 (the names of the various parts of the display are also indicated on the figure). Thus communication from the computer to the user is via a display of formatted alphameric data. Communication from the man to the computer is via the selection of options (some of which are called pads). Selection is accomplished either by touching the screen at the place where the option is displayed or typing the selection-character that appears at the front of each option or pad. Selection generally produces another frame of similar format with additional options. Selection may also evoke an action, which may lead to additional data being displayed or other activities taking place. Though it will not be of much concern here, ZOG is actually a communications device to other programs on the computer; thus the actions that it takes are to communicate commands and data to other programs (of arbitrary character), which in turn output back through ZOG to the user. From the user's viewpoint ZOG might as well be simply a program capable of carrying out complex activities. The display of Figure 1 distinguishes three kinds of selections, which are logically identical, but serve to structure the information on a display for the user. The options lead to other frames of similar format and are strictly local to the frame. The global pads along the bottom are available on every frame and constitute a set of general search and orientation capabilities. The pads vertically along the right side are actions that hold over a subnet, eg, standard commands to a given program, if the subnet is being used as a guide program to executing a program.

What distinguishes ZOG from a standard menu-selection scheme, of which there are many, is that the response time for the next display is essentially instantaneous (eg, around a tenth to a quarter of a second) and that the total set of frames through which one can course is very large (eg, tens of thousands). These two features go together since rapid response in a small net is of only limited utility. With ZOG, the user is to be able to stay within the net essentially indefinitely, gaining knowledge and giving commands. ZOG's being a communication agent to other arbitrary programs is also a distinctive feature, though it relates to the scope of its applicability rather than to the nature of the man-computer communication to which it gives rise.

Rapid response coupled with the large network produces a qualitatively different

¹ I would like to thank George Robertson and Kamesh Ramakrishna for discussions on the topic of this paper.



Global pads:

a-alter	Enter alter mode to modify (edit) current frame
b-back	Go back to previous frame
d-display	Display again (in case display is lost or garbled)
h-help	Go to Help subnet to explain how ZOG works
m-mark	Mark this frame for return (will show in Mark ID)
n-next	Go to frame next to this one in option list "above"
r-return	Return to last marked frame (shown in Mark ID)
z-ZOG	Go to ZOG1, the root frame of the total ZOGNET
↑C-exit	Exit from the ZOG system to monitor mode

A selection (option or pad) is made by touching the screen or typing the character that appears in front of the selection text. When a selection is made: a selection action may occur, a new frame may be displayed, and its frame action may occur.

Figure 1. Typical ZOG Frame

man-computer communication philosophy from standard menu selection. The PROMIS laboratory (Problem Oriented Medical Information System), of the University of Vermont Medical School, was the group that first developed this scheme and demonstrated its effectiveness in a large application system (Hurst and Walker, 1972). Our interest is in determining whether the communication philosophy is generally applicable, in discovering its necessary parameters and in characterizing its performance. Thus, in many respects ZOG copies the essential features of the PROMIS communication interface, though it reflects some additional design constraints and requirements.

It is our intent to attend seriously to how to describe human performance in ZOG. The very arrangement of ZOG, with the user cycling through a repetitive cognitive operation as he moves from frame to frame, entices one to the belief that answers should be forthcoming to questions such as the following:

- (1) Time: How long does it take a user to accomplish some task using ZOG?
- (2) Errors: What sort of errors are committed in using ZOG? How serious are the consequences of these errors?
- (3) Learning: How long does it take a user to learn a given body of material or a collection of procedures using a ZOG net?
- (4) Motivation: How much is the user induced to stay with the system until the task is done or the material learned? To pick up additional related knowledge and skill because the net beckons him? To come back to use it for other purposes?

Each of these questions, as soon as asked, suggests a companion:

- (5) Comparison: How does behavior (for time, errors, learning, motivation) in ZOG compare to the behavior in attaining the same objectives by some other system?

These questions are "external" questions. They take ZOG as given and seek an overall evaluation of its effectiveness and place in relation to alternative techniques. Another set of questions are "internal". They deal with design alternatives of the ZOG system to increase its effectiveness and its efficiency. They are of little interest in any absolute sense. Some examples are:

- (1) How important is system response time to overall performance?
- (2) What is the trade-off between high-fanout shallow trees of options and low-fanout deep trees of options?

- (3) What is the trade-off between the amount of text to describe an option clearly and the time it takes to read and comprehend it?
- (4) What is the trade-off between large amounts of text on a single frame and splitting up the text on separate frames?
- (5) What is the optimal layout of the display?
- (6) What is the optimal set of common interaction functions?
- (7) How should knowledge be decomposed to be presented in a ZOG network?
- (8) What mixture of presentation, questions, exercises, etc., make for effective learning?

Each of these internal questions, in itself, is unanswerable without being referred back to some of the external questions about time, errors, learning or motivation.

Neither set of questions can be taken as properly posed. They must be recast in a form suitable for being answered experimentally and quantitatively. They may themselves not be the right central questions to ask; they might follow from some more fundamental characterizations of human performance in ZOG. However, they do indicate pretty clearly the sorts of issues that must be addressed ultimately by any theory of human performance for ZOG.

The questions are not properly posed in yet another way. ZOG may be used for a variety of different tasks: guidance, education, interviews, data retrieval, command language functions, programming, and more. Many of these lead to different styles of ZOGnets. The answers to the questions will be correspondingly varied, along with the overall usefulness of ZOG for the different tasks.

It will probably come as no surprise that the answers to the questions posed for ZOG are not available for any comparable system. Though there is a long tradition of human factors work, and though there has recently been established a journal devoted wholly to such issues (The International Journal of Man-machine Systems), still the literature provides help only for various limited aspects. Some recent work we have been involved in on studying interactive editing (Card, Moran & Newell, 1976) comes about as close as any previous effort to being directly applicable. Editing and using ZOG share some operational aspects as routine cognitive skills, but editing does not extend to knowledge acquisition, problem solving and search, such as occurs in a ZOGnet. The present paper draws strongly on the approach of the editing study that one can build quantitative models for cognitive skills, but its details will not be pertinent to the level of discussion here. In any event, the best the present essay can provide is preliminary guidance.

The purpose of this paper is to examine the structure of human action in ZOG to lay

the groundwork for answering these questions, by experimental investigation and the development of quantitative models. We discuss four components of an ultimate performance model. First is the four measures we identified as important to ZOG performance. Second is the determinants of the time per frame, ie, what happens when the user processes a frame. Third is the problem of orientation in the net. Fourth is the description of user strategies. These four components constitute a major part of the total requirements of a performance model, though not all.

Throughout it will be premature to distinguish carefully the different ways ZOG might be used for different tasks. However, to have a concrete task in mind, consider the one first being implemented: a guide to the facilities in the Computer Science Department at CMU. Much of this ZOGnet consists of a expanding tree of explanation and description that covers in increasing detail the people, projects, computers and software available -- what they are, who does what, how things work, where things happen, some history, and so on. The user explores this tree frame by frame to varying degrees of detail, as suits his motivation and objectives. When arriving at particular software systems, say, the ZOGnet will guide the user in executing the software, either showing him how to use the commands of the program or having him execute the program indirectly by making specific selections (eg, hit various pads). It will suggest various exercises to try with specific programs and may demonstrate these to the user by having him step through a partly preplanned sequence to witness how the program behaves. The user would return again and again to the same ZOGnet, exploring different corners of it as his interests dictated. Thus he would travel through some parts of the tree many times (on the way to other newer parts) and would become familiar with the general layout of the network in varying degrees. Overall he would probably visit many more frames that just held information than he would those that had actions associated with them, because all action frames would be surrounded by various additional informational paths to explain the actions and their consequences.

THE BASIC MEASURES

We characterized the external performance of ZOG in terms of four variables: time, errors, learning and motivation. A model must first define these in measurable terms before it can relate them to structural features of interest (eg, those mentioned in the second set of questions). Let us indicate briefly what needs to be studied to treat each of these.

Time

Time is the only variable of the four that raises no basic measurement issues. The total time for a user to accomplish something in ZOG is decomposable into the sum of the times taken for each frame. Given some uniformity in the design of ZOG frames, there will be characteristic times per frame, so approximating total times by the number of frames times the average time per frame will be useful. The time per frame depends basically on the internal make up of the frame (its content and arrangement), whereas the total number of frames depends on the subject matter (broadly interpreted) and the strategies of the user. This factorization is not perfect, eg, frames are dealt with quite differently when used as a source of knowledge or when used as an access path to other frames. Still, detailed study of the time per frame and its dependence on the internal structure of the frame would seem to be a necessary component for an analysis of the behavior in the large, ie, of the total time to accomplish a given objective.

Errors

In ZOG, as in many interactive systems, errors do not show up primarily as an independent measure of performance. Instead, they are detected by the user himself and converted into additional time -- time to undo the error's consequences and proceed correctly. This is only an approximation; ultimately some errors sneak by the user's detection and become errors in the final product, either errors of learning or errors of system use. Basically, however, the study of errors becomes folded in with the study of performance time.

The detail required in a theory of errors depends on the detail of the time predictions. Simple error amplification factors might do very well for some purposes. In so far as errors can be defined by frames encountered or selections made, they can be defined with accuracy and can be detected automatically. If errors have to be defined by human analysis and post hoc judgment, they become extremely difficult to work with.

Even a brief operation with ZOG reveals a critical source of errors to be the user's orientation in the net. Where am I? Where did I come from? Have I been here before? Do I know where to go next? Can I get back here, if I wander off in search of some information? These questions are not unique to ZOG. They show up in every extended activity from wandering through a city or a building, to reading a book. In most areas of life the

orientation is built into our experience with the situation and is not distinct. But with ZOG the space is a network, which is foreign to the user both in structure and in experience. Furthermore the frames have a homogeneity (being always text arranged in similar fashion) that makes orientation and recognition more difficult. Some attempt to deal with orientation seems a necessary component for treating the more subject-matter dependent aspects of ZOG.

Learning

The basic requirement for measuring learning is a measure of performance. Learning is indicated by changes in the performance measure over repeated occasions (suitably taken to avoid contamination by other phenomena such as fatigue). The performance measure can take essentially three forms: (1) a binary value of succeed/fail, which leads to performance being measured either by the time to completion or by the probability of success; (2) a measure of the number of errors committed, either on the way to completion or in the final product; (3) a measure of the quality of the final product along some scale, as in judgments of the aesthetic value. Quality measures usually involve human judgment. Besides raising problems of interjudge reliability, they are relatively expensive to obtain in massive quantities. It is tempting in ZOG to use what comes easiest, namely time to completion and succeed/fail measures, both of which are detectable automatically (providing that succeed/fail implies arriving at some frame or taking some selection).

The creation of performance tests is part of the design of ZOGnets, in any event. Good instruction implies providing various indications of whether the user has learned and feeding this back to him. For ZOG used as a guidance system for other programs, tests are mainly problem sets with detectable success or failure conditions. For ZOG used in a more CAI mode, these are multiple choice subject-matter tests.

Learning enters into ZOG in many ways and it is important to be clear in any instance what learning is being measured. If a ZOGnet is being used to convey a body of knowledge, then the learning of interest is acquisition of this knowledge. If a ZOGnet is being used to teach how to use a programming system, then the learning is exhibited by how effectively the programming system is used in isolation after the ZOG experience. If a ZOGnet is being used as a guide to accomplish a one-time activity, then learning occurs throughout the use of ZOGnet, but it is part and parcel of the performance task and no separate measure of it may be possible. In this latter respect all use of ZOG involves learning by the user, even though we talk about it as a performance.

Motivation

Motivation refers to what people want to do; cognition refers to how they attain what they want. Given that goals structures are recursive, so that what a person wants is determined by his method for attaining his immediate supergoal, it is unclear how there can

be any separation between motivational and cognitive analysis. Still, some questions of interest in ZOG seem to be "motivational" in nature. They relate to perseverance, exploitation of opportunities, enjoyment, attractiveness, etc. Motivational analysis seems called for when there is free, but mutually interfering, choices among activities with no texture of difficulty in implementing any choice.

The basic result of a motivational analysis, whether one speaks of reinforcement, incentive, drive or utility, is a description of objects in a choice situation in terms of some coin of the realm so that the actual choice can be easily computed. Unfortunately, there is yet no a priori analysis that can be applied to construct this measure. For complex objects multiple observation of choice behavior in varying situations with varying alternatives must be taken. The available conceptual frameworks focus on effective summarization once the choice data are in hand.

Thus we are driven to start with direct measures within the ZOG structure for detecting choices of user interest. For instance, compared to the main road how often are side branches taken? Introducing into such choice situations various standardized rewards and penalties permits the construction of a common basis of comparison. For instance, compared to various computer games which are available within the net, how is free time distributed? If such measures work out fruitfully, then one can work backward from them to make contact with more general approaches to motivation. But the start must be with quite idiosyncratic measures that have only intuitive face validity.

TIME PER FRAME

The ZOG user goes through a cyclic operation during which certain functions must necessarily be performed, given the structure of the ZOG system. Starting when the new frame is displayed, the user must first comprehend the information on the display, then he must select what to do among the options available, and finally he must wait for ZOG to display the next frame. These three functions -- comprehension (C), selection (S) and waiting (W) -- form a useful decomposition of the user's action, though what exactly constitutes these functions remains to be investigated.

For any frame f we would like to be able to write simply:

$$(E1) \quad T_f = T.C_f + T.S_f + T.W_f$$

where T_f is the total time to process frame f , $T.X_f$ is the time spent to accomplish function X for the frame, and the times for the components partition the total interval. This is the simplest additive model. Given this, a number of the internal questions could be directly addressed. However, some of its difficulties are worth noting.

Let us start with the assignment of times to the components of (E1). To do anything experimentally this must be possible, at least indirectly. The cycle starts when a new frame is displayed and ends when the next frame is displayed. Thus T_f , the total time for frame f , is easily measured. Similarly, the boundary between S and W is clearly marked by the act of touching the key or screen. However, the separation point between C and S is not clearly marked by an external action. The conceptual distinction remains, however, between comprehending the new knowledge in a frame and considering where one might go next. Experimental variations of the ZOG architecture might be used to assign a time to this boundary. The options might not appear until a pad OPTIONS was hit, at which point the text would disappear (to be redisplayed if desired by TEXT). (This would assume that all comprehension of the options is part of selecting where to go next and does not add to the knowledge acquired in a frame.) Such experimental variations distort the total operation, but could be calibrated against the standard scheme. An alternative strategy, called the additive factors method, of experimentally varying various task characteristics to separate the stages in cognitive process is also available.

The next difficulty is that the processing for a function will not stay neatly in one partition. There might be oscillation between comprehending and selecting. The user might wait for a while, and then realize he hasn't selected the option. Additional comprehension might occur while the user waits for the next frame. Some comprehension might have occurred on prior frames if there were redundant knowledge. And so on. Some of these complexities might be handled by classifying them as errors (as in waiting before selecting). A more complex partition than E1 could be used, with many increments possible for each function. This would multiply the difficulties of assigning times to each subinterval. Alternatively, E1 could simply define the aggregate time for each function, no matter where it occurred in the interval. This would require different techniques to obtain estimates of the components, since the boundaries would no longer have meaning.

A third difficulty is that a single processing action is not necessarily associated with a unique function -- it may serve multiple functions. This is certainly possible between comprehending and selecting, where the options may give much of the information of a frame. This difficulty is a direct denial of the additivity assumption in (E1), for it says that processing acts are not independent.

A final difficulty is the extent to which the user can process in parallel, so that even if additivity in the amount of processing holds, it does not translate into additivity in the time intervals taken. Given that there is some motor movement as well as some possibility for highly learned responses, the possibility of concurrent processing must be entertained.

The extent to which any of these problems are serious is an empirical question. At some level of detail they must raise difficulties, but simple approximations may be eminently useful. Let us consider each of the three functions (C, S, W) to pin down more clearly what is involved.

Comprehension

A frame provides knowledge that the user can acquire. Fortunately, the frame itself is limited enough and its language simple enough so that, without undue difficulty, we can take this knowledge as a set of elementary propositions, each of which is about some set of entities of interest. The language of the frame itself can be used as a first approximation, though a single sentence might need decomposing into several elementary propositions. Some more formalized notation might ultimately be useful. We can refer to these elementary propositions as knowledge elements.

The frame offers knowledge; the user need not comprehend it. He may already have the knowledge or some part of it. Having the knowledge does not automatically avoid all comprehension. If the user does not know that he knows what is in the frame, he must comprehend the frame enough to recognize that state of affairs. The user may also not comprehend all the knowledge in a frame through deliberate strategy or through inadvertence. To do so, of course, he must avoid attending to it, at least in detail.

The act of comprehending an element of knowledge might be taken as primitive without undue distortion. Given that its components are already known, ie, its predicate and argument entities, it requires a single "unit act" to comprehend. If the components are not all known, then the proposition cannot be learned, though it can initiate learning of the missing components. This is a strong continuity assumption, but its adoption as an approximation is reasonable both theoretically and experimentally.

The user can comprehend a knowledge element whenever it is presented and he is prepared. An element on the given frame could be obtained before arriving at the frame, either prior to the entire ZOG run or from earlier frames, or after the given frame from later frames. The possibilities for sharing between frames can be determined from the knowledge

elements for the frames. Likewise the extent to which some frames are a necessary preparation for a given frame (for the necessary prior learning of components of a proposition) is also determinable.

The knowledge elements have so far been taken to be the actual propositions expressed on the frame. But the user can infer or induct knowledge as well as simply absorb what is presented. The extent to which users go beyond the knowledge given would seem to be a fairly tricky affair. But again, simple approximations may be possible. Users readily form single-element generalizations. When presented with some knowledge in concrete form, what is comprehended (and learned) is a more general rule. This is often narrowed or otherwise shaped by subsequent experience. On the other hand, inferences involving chains of deductions from sets of knowledge elements may be much rarer, especially in knowledge-conveying, as opposed to problem-posing, situations.

Comprehension and learning must be distinguished. Comprehension implies the ability to use knowledge in some way; learning implies the maintenance of some comprehension (once acquired) over some period of time. Good evidence exists that learning will take place if the knowledge is used. (In the current metaphor of psychology, the deeper the level of processing of some material, the greater the probability that it will be learned.) However, this is confounded with the amount of use, so that knowledge elements on a frame undoubtedly can be comprehended and used locally without being learned so as to be available at some later time. Thus, we may need to distinguish both an amount of time for the unit act of comprehension and the amount of time for long term learning of the knowledge element.

We have identified three distinct processes that occur in the comprehension phase for a frame: recognition (that what is expressed is already known), comprehension, and learning. These each apply to the single knowledge element. The total comprehension activity for a frame consists of a sum of such processes over the set of knowledge elements in the frame, plus any control processes that are required. These latter might make the situation much more complicated than the additive model implicit in the knowledge element scheme. What might these control processes be?

One possibility is a general orienting response to the new display. Such orienting responses are readily observed in startling situations. It could add a set-up time of a few seconds to each frame, which might constitute a limiting factor to rapid use of the net. Another possibility is a decision process about which elements to consider. With large bodies of material considerable conditional behavior can occur that just ignores part of the material. However, the single frame contains only a small body of knowledge and the user may adopt a simple exhaustive strategy. Finally, comprehension may not occur in a single pass, especially with several levels of processing available for a knowledge element. However, there may be no way to separate out the passes experimentally.

What can be said about times for the various processes we have talked about? Take comprehension first. There is some directly relevant data in the work of Kintsch (1974), who measured the time to read material incorporating different numbers of elementary

propositions. He finds rates of .5 to 4 secs/proposition, with the variation depending on the difficulty of the material (.5 is for adults reading children's stories) and on the length of the total material, where it takes appreciably longer per proposition (independent of type of material) as the total length increases. The data are not extensive, but it indicates what we can expect in experimental measurements on ZOG. From a practical point of view it might be argued that one could simply use reading rates (eg, words/sec), since number of words and number of propositions must be highly correlated. Ultimately, as a practical matter, this might be right, though Kintsch is careful to produce times per proposition independent of how many words it takes to express the proposition.

Comprehension was the middle of our three processes. There does not appear to be any similar data for recognition, ie, the time per proposition to recognize that you know a given proposition. There is much data on practiced recognition of single sentences (eg, "a canary is a bird") but it is not clear these times are relevant. Off the cuff one might assume recognition at about half the time of comprehension, though much depends on the conditions under which the process is examined (eg, reading or scanning) on the a priori expectations. For learning on the other hand there are a number of estimates on the time per chunk for verbal material. A typical one is 8 secs (Gilmartin, 1975), but it ranges as low as 2 and as high as 10. However, these are not for propositions but for words and other verbal material.

Selection

The standard ZOG frame has three classes of selections. First are the options, which form a vertical array below the text. These are local to the frame in that the user does not know what they are until he sees them, though he may be able to predict them from the text plus general expectation. Second are the standard pads along the bottom line, which are available on every frame. The experienced ZOG user becomes totally familiar with these. Third are the pads in the righthand vertical column, which are local to a group of frames. The user learns about these on some frames and applies them repeatedly in accomplishing a task. They signify invariant highly meaningful functions. The user will be in various states of familiarity about them. Each of these three types of selections may require separate treatment.

Consider a special case of selection, where the user, having finished comprehending, knows what he wishes to select (and knows that it will be there). His task is only to locate the option on the display and move his finger to touch the display. In this situation something is known about how long it takes to make the selection, ensconced in Fitts Law (Welford, 1968):

$$(E2) \quad T = K1 + K2 * \log_2(D/S + .5) \text{ seconds}$$

T is the selection time; D is the distance to be moved; S is the size of the target; K1 and K2 are constants that depend on the subject and the task arrangement; and log₂ is the base 2 logarithm. Fitts Law essentially asserts that guidance rather than movement is the governing

consideration. Thus, relative accuracy (D/S) appears in the formula and the overall function is logarithmic rather than linear, essentially bits of uncertainty. There is some evidence that K2 is never faster than about .1 sec per bit. K1 will be in the range of 1 to 2 seconds. Learning can reduce the total time appreciably (25% in some experiments).

We can analyze the selection phase by starting with this particular decision model and inquiring how the ZOG situations differ.

The formulation might be useful directly in certain conditions. It probably applies moderately well to the selection of standard pads. The pads are all about the same size (about 2 cm) and the finger will be located at distances up to about 20 cm away, but more like 6-10 cm typically. These lead to the logarithmic term being between 2 and 3, with occasional values towards 4. Total times might run therefore 1.5 to 3 sec. The same structure might also describe a set of vertical pads, used repeatedly within a subnet.

Other selection situations differ from this along two directions. The first approaches typing. The user proceeds through a sequence of familiar frames. The decision about the sequence is made at the beginning, and he knows not only the options to be taken (semantically) but also their locations on the display. We can think of this as knowing that 1-3-2-1-4 is to be selected, except that the knowledge is motor knowledge not verbal knowledge, like the knowledge of how to move the fingers to type B-A-C-K. In such a situation there is substantial anticipation and the user's fingers move continuously from one pad to the other in a typical skilled gesture.

How closely sequence selection will approach typing is unclear. Many features limit it: the tactile character of the touch surface (flat, unyielding); the use of fewer fingers; the difference between option sequence and words (character sequences), both in frequency of occurrence and in redundant substructure, the difficulty of knowing exact frames, and probably others. Yet one can expect some aspects of skilled sequence behavior to show up. Indeed, one justification for the speed of ZOG is that standard menu selection techniques seem cumbersome for the experienced user as a device for specifying command sequences he already knows well.

The other direction is toward greater decision making, rather than less. In one major use of ZOG the typical situation will be that the options are unfamiliar, relating to the content of the text presented on the frame. The user needs to comprehend an option before he can select it. There are several possibilities for the user's internal situation, each of which suggests a first order approach to the decision process.

The user may have developed some strong expectations and preferences of the options he would like to take. His task is to comprehend the options that exist and identify the one that can be interpreted as the one he wants. When he finally spots it, he is in a situation not unlike the one to which the basic model applies. Additional processing for comprehension would be expected, but only for the options that occurred before the one selected. Furthermore, the one he wants is not necessarily distributed at random in the list.

The user can be expected to read the options from the top down, having finished reading the text. The appropriate design strategy puts toward the top those options which are relatively frequently selected. (This design strategy, of course, reinforces the strategy of linear, top-down investigation.) Thus, a more plausible model has the user spending some expected time (c) comprehending each option with a probability (p) that he will select the current option. Under these conditions the expected time (T) to determine the desired option in a list of K options is:

$$\begin{aligned}
 \text{(E3)} \quad T &= cp + 2cp(1-p) + 3cp(1-p)^2 + \dots + Kcp(1-p)^{K-1} \\
 &= cp[1 + 2(1-p) + 3(1-p)^2 + \dots + K(1-p)^{K-1}] \\
 &= (c/p)[1 + (1 + Kp)(1-p)^K] \\
 &= (c/p)[1 + (1 + m)\exp(-m)] \text{ (approx)}
 \end{aligned}$$

Where $m = Kp =$ expected search length

An interesting feature of this expression is its virtual independence of K , the length of the list, if there is any positive bias in probability toward the front of the list (ie, $p > 1/K$ implies $m > 1$ and the negative exponential kills the term involving K).

A second situation has the user uncertain about what to do. Then he will use the options to suggest a course. In this case, all of the options will be comprehended and some sort of linear dependence on the size of the option list would be expected. Little guidance exists for exactly how this process will go. Will the user be able to select the one he wants after considering them all, or will he require another pass, this time in a decision rather than a comprehending mode? Would this second pass, even if it existed, be seen in the data, or simply be absorbed in a larger constant time?

A final point can be imagined along the dimension of uncertainty where the user behaves as if the decision were a complex and important one. Some evidence exists that deciding in such situations is a series of successive passes, each rejecting some alternatives and narrowing down the choice set (Tversky, 1972). However, ZOG has been designed to eliminate the occurrence of such situations, by making easy the selection of further frames to gather more information. Several standard pads (back, mark, return, next) are devoted to this, as is the design principle of always providing additional paths in the network to explain further. Only experience with various networks will show the extent to which deliberation has been eliminated.

ZOG is not monolithic, despite the common frame structure and display arrangement. If ZOG is being used to run System X (being effectively a command language for X) then System X itself imposes a structure on the decisions and the explorations for knowledge. This structure may be the dominant determiner of time and errors. This can probably be

determined just by seeing whether models based on general ZOG structure fail when used to predict behavior using System X via ZOG. However, we may have to examine the use of System X via more standard command languages to disentangle the features of the system from the features of ZOG.

As selection becomes more complex, the problem of delineating a boundary between comprehension and selection becomes more difficult. In the selection models described above, comprehension concerns what selecting an option would provide; thus in principle it is distinguishable from the acquisition of knowledge for later use. But frames sometimes use the option texts to convey knowledge, where the option itself is the implicit choice to get more information about the knowledge element in the option text (eg, a list of useful principles for how to use ZOG nets). Far from being rare, this may be an effective way to design frames. Conceivably, in nets that are heavily knowledge oriented, it will make no sense to separate comprehension from selection. This could happen either because they are in fact one continuous process or because the comprehension time simply overwhelms any separable selection time.

Waiting

The duration of the third phase of the ZOG cycle, waiting, is governed by ZOG and the computer system within which it lives. Design specs for ZOG place the wait time at less than .1 sec, which is instantaneous from the viewpoint of human action (though not quite from the viewpoint of perception). But with the present system design this is only maintainable on pure knowledge frames. Any attempt at complex action requires execution of a user's program and produces variable delays up to several seconds. The attempt to use ZOG on more conventional systems (such as the CMU time-shared PDP10s with alphameric terminals) produces delays of 5 seconds just to display the frame at 1200 baud (a typical fast rate), and more like 20 seconds at 300 baud (a typical slow rate). In addition to the expected value of the delay, the variability may also be relevant.

The important question is what effects waiting has on human performance. The simplest view is that it is just lost time with no other effects. It seems unlikely to have any positive functions, since waiting can occur under the user's control at any point he wishes. Selection has a clear termination, so that nothing remains to be done during the waiting phase. If the user needs a short preparatory phase to build up expectations for the new frame, then a short delay might be indistinguishable from zero delay. The user would produce the equivalent delay in orienting response to the next frame even if it came up instantly.

Two negative effects (besides lost time itself) seem possible. One is the memory barrier. Any state being carried from one frame to another, whether in terms of expectations or data elements (when using ZOG to perform a task) must be held over the waiting period. There is no reason to assume this memory load is excessive. However, any load at all will tend to assure that the waiting period is used only for waiting. And if the

period is too long, the probability of losing state will certainly become appreciable. What is too long is not known (try greater than 10 seconds?), but could probably be measured without too much difficulty in a system with controllable delays.

The second negative effect is on motivation. The delays may make the system seem a drag to use, killing either the tendency to use it or, more subtly, the tendency to use it freely to explore for other knowledge. In the design of ZOG (and of the PROMIS system) this concern has been central in insisting on rapid response.

Poor (and variable) system response is endemic to current computer systems. However, no studies are known to us on the quantitative effects of this on human performance. Some design adaptation has occurred. An interesting one is type-ahead, in which the user is allowed to type a sequence of inputs to the machine, based on his prediction of appropriate response. (This wins if the delay is caused by the time between computing quanta allocated to the user, so an entire sequence of inputs can be processed if available at the beginning of the user's quanta.) The extent to which select-ahead can work in ZOG needs to be explored. It may turn out to be strongly affected by selection modality, type-ahead working quite well, while touch-ahead doesn't.

ORIENTATION

The general problem is easily enough described. How does the user orient himself in working in a ZOGnet, how does it affect performance (time and errors), and how does it depend on the structure of the net and of the frames? The straightforward approach starts with the question: What is an orientation? What sort of cognitive structure does a user have when he has an orientation? Given this, it is sensible to ask how such a thing is built up and how it is used in working with a net. Given this second step, the evaluative questions can be asked of whether an orientation is useful or misleading, adequate or inadequate. Psychology is not currently in a good position to deliver the answers to such questions, but the current possibilities can be indicated briefly.

There is a relatively new subfield of psychology and geography directly concerned with understanding spatial images and orientations. It seems exactly what is needed. It started with Kevin Lynch's *The Image of the City* (1960), which was undertaken from a concern with urban design rather than psychology, and attempted to characterize how people viewed the city in which they lived -- what image they had of it and how they oriented themselves within it. The principle techniques are getting people to externalize their mental maps of geographic area (cities, countries) by drawing maps and describing terrains and tours (Downs and Stea, 1973). But the resultant data (hand drawn maps, etc.) leave us a long way from a proposal of what an orientation is, especially in ZOG, which is, in its own way, quite abstract.

The most promising area is called proximity analysis (Shepard, Romney and Nerlove, 1972). This is a psychometric technique for taking similarity (proximity) measures among a collection of objects and constructing a multi-dimensional metric space which reproduces the given data. This is now a well developed area (with a range of related techniques) and has been used successfully in ways that fit the basic requirements of our problem. For example, taking as proximities the confusions of teletype operators of morse code, it can show that these confusions can be described in a two dimensional space, one dimension being length-like and the other symmetry-like; further, that this space differs between novice and expert operators. Indeed, most current hopes of becoming quantitative in the first-described area of geographic images and orientation rest on utilizing this technique. A conceivable approach to orientation in ZOGnets is to obtain judgments of proximity (how far is information X from Y, the current frame) and see how the resulting space can be characterized. The technique requires substantial data and rather stable organizations, but it might work. One peculiarity is that the objective space (as represented by the net itself), may not have a finite dimensional character, but branch outward in exponential fashion. This might make proximity analysis, which forces low order dimensionality on the data, genuinely deceptive. Contrariwise, if humans insist on organizing any world in which they can get "oriented" into some low dimensional space, so as to make it familiar, then this might be a superb way of discovering it.

The notion of spatial image seems inextricably linked to the notion of orientation. In both areas above the answer to "What is an orientation?" would be "A space with such and

such properties". This simply pushes off the question of how such a space is used in maneuvering through a net, but given an actual spatial representation to work with, this latter problem is surely approachable. Psychology has been deeply concerned with imagery generally in the last decade (Pavio, 1971; Sheehan, 1972). A noteworthy feature of this work is that it has not yet been able to produce good concrete proposals for the nature of the image. One strand has avoided the issue, working experimentally to show that aspects that presumably correlate with imagery (eg, concrete as opposed to abstract nouns) have strong and stable effects. A second strand has attempted to demonstrate internal representations (images) that are space-like, eg, have properties of continuity and proportionality to distances in an imagined space. But nothing has proceeded further than simply affirming that the image is indeed generally like external space. A third strand takes the view that images are symbolic structures (of the kind familiar in artificial intelligence). This approach is concrete enough, but in effect enjoins abandoning spatial aspects as central and concentrating on the semantics involved directly. Its answer to the question "What is an orientation?" is "A symbolic structure". But it does not yet add much in the way of details.

Let us attempt a little direct analysis of the ZOG situation to see what the notion of orientation should do for us in explaining performance and what seems like natural definitions within the restricted framework of ZOG.

We are concerned with the following phenomena:

- (1) A novice user hesitates for a relatively long time at various frames, declaring himself to be uncertain about what to do.
- (2) An expert user skips nimbly from frame to frame, even in new subnets.
- (3) A user declares himself lost and says he doesn't know what to do.
- (4) A user keeps returning to some particular frame as a homing place, even though this means going through redundant sequences of frames repeatedly.
- (5) Before a user will do anything (in the sense of executing operations) in a new subnet, he explores the frames several times.
- (6) A user always tries all options systematically from top to bottom, even though he has a specific task to perform and the options seem clearly marked.

We have used spatial language to indicate how an observer or the user himself might naturally describe this behavior. The behavior has definite consequences for performance, slowing it down, speeding it up, producing errors, etc. As described, it operates as a bias over the behavior or perhaps as a partial component of a strategy, affecting many frames rather than just a single one.

Consider the notion of feeling "at home" in a place and its contrary of feeling

"strange". Strange things can happen in strange places, things that one is not prepared to cope with. Paralysis of action and heightened vigilance occurs in strange places, grounded in part to devoting capacity to being prepared, in part to the hypothesis that action produces (stirs up?) strange responses and in part to the hypothesis that movement (to a new place) will yield an even stranger (more hostile?) place. Contrariwise, at home there is a sense of control. Whatever needs to be done will be seen and accomplished -- even if it is not yet foreseen in detail. There is relaxation to attend to the current task fully. To be at home implies that the territory is familiar, with expectations of what it contains in the way of possibilities. Movement though the environment seems safe, leading to places which themselves are familiar.

A possible way to treat this phenomena is to take it as a strategic response to specifiable concerns, each of which add some processing to be done in dealing with the frame. (1) If there is a concern that something strange may happen, then observing behaviors must occur with a certain frequency. (2) If a movement into a stranger place may result from a selection, then the possibility must be tested for each contemplated choice (in addition to other considerations); if the concern with strange results is strong enough then it may become a primary requirement on the choice of option. (3) If one is lost, then each frame must be searched for clues that would help to locate it. And so on. The slow down in behavior comes from layers of these concerns. The speed-up from a sense of familiarity comes from eliminating them from processing of a frame.

These concerns may hold independent of the design of the ZOG system. Deliberate design has dealt with some aspects of this problem. ZOG is essentially a passive system under complete control of the user. Therefore, there need be no concern that strange things should suddenly happen. Likewise, a design principle of "No sudden death" has been adhered to rather strictly: it is always possible to get back immediately to where one started (the back pad); no irreversible action will be taken without informing the user and getting confirmation; etc. Therefore some layers of concern should extinguish quickly upon general familiarity with ZOG. Tests are needed to see whether in fact they linger; if so, special studies may be required to understand why they cannot be eliminated (or eliminated more rapidly).

Other concerns cannot be eliminated in such obvious ways. Much ZOG use (ie, for knowledge acquisition) implies that frames are being entered for the first time; hence, in some objective sense they are in fact strange. Whether familiarity/strangeness concerns are evoked is a separate issue. It may still be possible to model the behavior of less experienced users by means of a scheme of incremental concerns.

A feature of this approach is that it bypasses a direct representation of what the cognitive structure of "an orientation" is, replacing it by a collection of response systems for dealing with the issues that derive from lack of orientation and a set of environmental clues about when these systems will be evoked.

GOALS, STRATEGIES AND ERRORS

The total number of frames visited by a user to deal with a given domain of knowledge or action results from the interaction of three aspects:

- (1) The encoding of that domain into the ZOGnet.
- (2) The acquisition or operation strategy adopted by the user.
- (3) The sorts of errors made by the user that have implications for touring through additional frames.

It might seem that (1), the actual network, would simply be an exogenous variable as far as a theory of human behavior for ZOG is concerned. A theory should deal with behavior through an arbitrary network. However, a user's behavior arises from his goals plus his analysis of how to attain these goals. This latter depends on the structure of the network as he perceives it. In discussing orientation, we already noted some potential examples where principles of universal back-up may have large effects on the way users treat the network.

ZOG presents a peculiar parallelism. From the user's point of view, ZOG is a task environment. He brings to it his own goals (eg, to learn something or to accomplish some processing); his problem is to use ZOG to accomplish these goals, reading material and taking options as seems appropriate in the light of his running analysis. The user builds up internally a hierarchical structure of goals and subgoals for doing this, just as he does in any other task. He may plan ahead, try different alternative routes, etc.

On the other hand, ZOG itself is analogous to a goal structure. Frames are goals; they explicitly state problems to be solved. Their options link to subgoals, ie, to other goal-like frames which state subproblems to be solved. Some are AND-like, where the options decompose the problem into a total set of tasks, all of which have to be done (eg, all of the various aspects of a subject matter that have to be learned). Some are OR-like, where the options present alternative ways to solve the problem (eg, the different commands in a programming system). Thus, in general the ZOG tree is analogous to the familiar AND-OR goal trees of problem solving programs. Some frames are operators, rather than goals, giving the final knowledge that is the solution to the problem. Likewise, some selections (usually pads rather than options) are operators that accomplish final actions.

From this latter "goal-net" point of view, the user is to be seen not as an independent problem solver, but as an interpreter who simply traverses the fixed pre-generated goal tree of ZOG, in response to the data (or lack of it) which is available to him. Carried to the extreme, there is no independent way that ZOG can be used. A user can conform to ZOG or he can fail to conform. He can bring to it his own goals, electing to move within those parts of the net that match these goals. But he cannot use ZOG as a carpenter uses a tool chest, or as a programmer uses a regular programming language. He can of course use ZOG as a whole to accomplish tasks (eg, to program), but only by doing it ZOG's way.

This goal-net view of ZOG is, of course, a particular design philosophy. Other philosophies exist. One would be to decompose knowledge domain in some logical manner, providing the total material in terminal frames, with nets of index frames to provide selective access. A paradigmatic example would be a taxonomic reference of, say, birds. Another would be to let users modify the net to fit their own styles (a technique we are in fact exploring). We will examine only the goal-net view, as befits a preliminary investigation. The goal-net view seems attractive, both in the amount of guidance it might give to net construction and the leverage it might give to understanding human performance.

Encoding of the Domain

Let us spell out the design philosophy in more detail. What we describe is neither operational nor without difficulties, requiring expansion and debugging against actual cases, but it will allow discussion of user strategies. Assume the user is an experienced ZOG user, though not necessarily familiar with the particular net to be designed. (Introducing ZOG poses additional problems.) Then, to construct a ZOGnet according to the goal-net philosophy:

- (GN1) Determine a set of possible top goals that the user can solve with the net. Identify each of these with a specific frame.
- (GN2) Create an explanatory index of frames that lets the user get from the top frame of the net to these goals. The user must find out what goals the net solves.
- (GN3) For each goal create one or more methods for attaining the goal. A method is a conditional sequence of actions (activities that produce some results without further problem solving) and subgoals (which recursively involve further methods and goals). Each method can be encoded in a network. Actions are options that lead to information frames, or pads that command an action-program. Subgoals are frames. The control structure (sequencing of actions and logic of goal selection) is encoded into sequences of frames, with appropriate explanatory intermediate frames if necessary.
- (GN4) If there are alternative methods for a goal, create a choice-net that leads from the goal-frame to the methods, explaining the factors on which the choice rests.
- (GN5) Specify the alternative states of relevant knowledge that the user can have. Provide alternative "short-circuit" paths in the net so that the user is not forced to obtain knowledge from the net that he already has.
- (GN6) Specify the ways the user might be confused about what to do, or about the meaning or intent of various terms and phrases. Provide alternative "long-circuit" paths at the point of use that provide explanations.

User Strategy

From this description we can see that the rational strategy for the user is as follows:

- (U1) The user approaches the ZOGnet in some particular state of knowledge and with some particular goal (we ignore the possibility of several goals).
- (U2) Entering at the top frame, the user explores the top-goal index net until he is able to identify which (if any) of the goals of the net match his goal.
- (U3) The user simply follows the directions of the net. At each goal frame he must select a method and for each method he must carry out its provisions. There are various nets at each point which tell which choice to make as a function of the details of the user's goal and the data that the user has.
- (U4) If the method-operations involve executing a program, then the user observes its results, perhaps recording them if the program itself does not have such facilities. This adds knowledge for making further choices in the net or for the final solution.
- (U5) If the method-operations constitute presentation of knowledge, the user attempts to comprehend the knowledge and retain it (as indicated). This adds knowledge for making further choices in the net or for the final solution.

According to this model, there is a single uniform strategy for user behavior. The actual course of frames (hence the total number of frames, which is what we would like to predict) is governed entirely by three things: the net, the goal of the user, and the knowledge the user has. Given a theory for how the individual frame is processed, as discussed earlier, detailed simulations should be possible. Their accuracy will be limited by how good this detailed theory is. Note that the user's goal and knowledge must be given. Basically, there is no way out of this. This information produces qualitatively different choices throughout the net, causing indefinitely large effects in the total behavior. (One can, of course, attempt to derive various averages or bounds on behavior given various minimal assumptions about the user's goals and knowledge; the ability to do this is beyond discussion here.)

What scope is left for the user's own goal directed behavior? There does not seem much freedom to use the net itself in ways other than its designed modes. This does not say that the user's total behavior is determined by the net. For instance consider a net for using a program, eg, a text editor, such as LINED on the PDP10. The system consisting of the LINED-NET plus LINED is still a text editor, and the user manipulates this total system to accomplish some editing goals that he has. LINED-NET + LINED does not determine what editing the user wants done. What is determined is the way to use LINED-NET to evoke the editor, given the user's goals. This would also seem to hold for the acquisition of knowledge from a ZOGnet, a slightly less obvious proposition.

One freedom still left open is the use of various short-circuits and long-circuits (GN4 and GN5 above). On the assumption that the user is completely certain what he knows and doesn't know, the choice in any instance should be clear cut. However, clarity may not be possible, so that the user must make a judgment. If so, the decision situation is quite precise: the user wishes to predict his actual state of knowledge. If he guesses wrong he will have wasted time either exposing himself to knowledge a second time or going down a wrong path. An additional alternative is taking too long to search his own mind to see whether or not he has the knowledge -- he might just as well have taken the long-circuit again. His ability to solve this decision problem will affect the total number of frames and hence the total time.

Error

Given this view, the important questions, relative to determining the total set of frames, address what can go wrong. The picture above, and any prediction based on it, assumes error free performance. Additional psychological features of the user enter in through errors of various kinds. These errors are not purely a feature of the user -- decrements to be assigned to his psychology, so to speak. They may be induced by the net, either from a failure to carry out the goal-net design philosophy or from the philosophy itself being inherently mismatched to human capabilities.

We have already had reason to discuss two specific sources of errors, those stemming from orientation problems (in the prior section) and those stemming from uncertainty over the need for explanation (just above). We need to characterize more generally the sources of errors and their consequences. Consider first the varieties of user error, taking the net as given. There seem to be six fundamental types of errors, having to do with: selection, acquisition of final knowledge, retrieval of control knowledge, memory load, uncertainty, and the interpreter role.

The user can make a misjudgment in some selection (option or pad). The source of this may vary from knowing what to select but accidentally touching the wrong option or pad to being mistaken about the meaning of the selection to misjudging his own state of knowledge (as discussed above). In all cases, the effect is to go down a wrong path until it is so recognized, then recover and proceed down the correct path. In accordance with similar tasks one can expect the bulk of these error-recognize-recover-resume sequences to be simple undo's, leaving a short loop-like path that took time but with no other effect. A few will have some side effects, so that the recovery path will lead to a different frame than the error-initiation frame or so that the user has acquired some new useful (or harmful) knowledge. A very few will convert into serious errors, in which the user gets confused or lost or, in general, loses his orientation; these may take a very long time for recovery.

The user can fail to acquire some final knowledge provided by a frame, ie, knowledge that constitutes part of the solution to the user's goal but that is not used within the ZOGnet. Such an error may translate directly into an error in the final goal. The goal-net design supposedly tests for correct acquisition. However, there are a pandora's box of difficulties

in doing so. and it may not even be possible in principle. If ZOG is used to run a program, there is no way that ZOG can know what constitutes the correct program according to the user's intentions. Even if ZOG is used only to acquire "pure" knowledge, many issues exist about when and how to test. If the user himself detects the missing knowledge, then the result of the error is an added loop to acquire the knowledge from the point of error detection.

The user can fail to acquire knowledge used for further choices, ie, control knowledge. Unlike final knowledge, control knowledge has some built in error detection capability, since the user will find himself without the knowledge at some point. Actually, what counts is not acquisition but retrieval at the critical point. In any event, the user must engage in another extra exploration to reacquire the knowledge. Without better knowledge of the design one cannot say whether all such loops will be nearby (as GN6 suggests).

There may be explicit short term memory requirements when using other systems through ZOG. ZOG may only be of limited help, either because of memory limitations in ZOG itself (as opposed to the program being executed) or because giving the knowledge to ZOG (which did not generate it) is too much bother. The user can fail to remember such information, if he becomes overloaded. In either case, the normal consequences is re-execution of the program involving retransversal of the net.

The user may be uncertain of what to do and rather than get additional knowledge, may simply focus on the given frame. Some of the effects of disorientation would show up in this way. The result is an expansion of time at a frame, rather than extra frames being traversed.

Finally the user may not use ZOG as the goal-view intends. He may reject or not understand the role of problem-solving interpreter assumed by the goal-net view and taken as the basis of performance analysis. Conceivably he may find another way of using ZOG that works. Some of this is also covered in the notion of disorientation. But there may be a number of small ways in which the user simply deviates from adopting the interpreter role. It is difficult to describe the consequences without having a sample of such behavior as a guide.

The purpose of going through all these types of errors is to observe that most of them (though not all) tend to have their consequences describable as additional explorations of the net, taken under some definite initial conditions. Thus they translate into additional numbers of frames (ie, additional time), as we observed earlier.

The user is not the only one who makes errors. So do the net builders. That is, they deviate from the design called for under the goal-net specifications. We can ignore gross errors that make the net incapable of solving the stated goals and the like. But there will still be errors that make using the net more difficult. These may be classified as errors that (1) fail to provide appropriate explanations; (2) fail to provide appropriate short-circuits; and (3) provide misleading text either as explanations or as options. Without going through the

enumeration, these either result in additional paths to be traced or offer additional opportunities for the user to make errors.

We have distinguished failures of the actual net to adhere to the goal-net design philosophy from failures of the user to be a perfect goal-net interpreter. There are also inherent limits of the goal-net design. However, they do not show up directly as a third source. Rather they are revealed by an inability to decrease errors of the first two kinds to a negligible effect or by failures of a "perfectly functioning" total system (net plus user) to compete with alternative schemes (other ZOGnet design philosophies or other communication technologies).

CONCLUSION

ZOG seems to be extremely well placed to permit an approach to analyzing its performance as a total man-machine system. Our interest in doing such an analysis goes well beyond our desire to answer the specific questions enumerated at the beginning. As we observed, the quantitative analysis of total man-machine systems is hardly yet a well developed art. We might hope to contribute to that. The fractionation that ZOG permits, coupled with its rapid response, may permit some basic chronometric analyses that would also contribute to fundamental cognitive psychology. How humans acquire knowledge, how they make decisions in the service of acquisition, even the basic nature of orientation, all seem approachable.

The components we have considered -- the time per frame, the orientation of the user, and the user's strategies -- appear on analysis to be sufficiently central and sufficiently well structured that an immediate approach can be made to them, both theoretically and experimentally.

REFERENCES

- Card, S. K., T. P. Moran and A. Newell, The Manuscript Editing Task: A Routine Cognitive Skill, SSL-76-1, Xerox Palo Alto Research Center, December 1976.
- Downs, R. M. and David Stea (eds.), Image and Environment: Cognitive Mapping and Spatial Behavior, Chicago Press, 1973.
- Gilmartin, Kevin, An Information Processing Model of Short-Term Memory, Ph.D dissertation, CMU-Psychology Department, 1975.
- Hurst, J. and K. Walker (eds.), The Problem-Oriented System, MEDCOM Press, New York, 1972.
- Kintsch, Walter, The Representation of Meaning, Erlbaum Associates, 1974.
- Lynch, Kevin, The Image of the City, MIT Press, 1960.
- Pavio, Allan, Imagery and Verbal Processes, Holt, Rinehart and Winston, 1971.
- Robertson, George, A. Newell and K. Ramakrishna, ZOG: A man-machine communication philosophy, CMU-Computer Science Department Technical Report, August 1977.
- Sheehan, P. W. (ed.), The Function and Nature of Imagery, Academic Press, 1972.
- Shepard, R. N., A. K. Romney and S. B. Nerlove (eds.), Multidimensional Scaling, vol. 1 and 2, Seminar Press, 1972
- Tversky, Amos, Elimination by aspects: A theory of choice, Psychological Review, vol. 79, pp. 281-299, 1972.
- Welford, A. T., Fundamentals of Skill, Methuen Press, 1968.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) NOTES FOR A MODEL OF HUMAN PERFORMANCE IN ZOG		5. TYPE OF REPORT & PERIOD COVERED Interim
7. AUTHOR(s) Allen Newell		6. PERFORMING ORG. REPORT NUMBER
9. PERFORMING ORGANIZATION NAME AND ADDRESS Carnegie-Mellon University Computer Science Dept. Pittsburgh, PA 15213		8. CONTRACT OR GRANT NUMBER(s) F44620-73-C-0074 N00014-76-C-0874
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research Arlington, VA 22217		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
14. MONITORING AGENCY NAME & ADDRESS (If different from Controlling Office)		12. REPORT DATE August 5, 1977
		13. NUMBER OF PAGES 29
		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) None		