COGNITIVE SCIENCE, 1, 193–215 (1977)

Problem Solving in Semantically Rich Domains: An Example from Engineering Thermodynamics*

R. BHASKAR AND HERBERT A. SIMON

Carnegie-Mellon University

Recent research on human problem solving has largely focused on laboratory tasks that do not demand from the subject much prior, task-related information. This study seeks to extend the theory of human problem solving to semantically richer domains that are characteristic of professional problem solving. We discuss the behavior of a single subject solving problems in chemical engineering thermodynamics. We use as a protocol-encoding device a computer program called SAPA which also doubles as a theory of the subject's problem-solving behavior. The subject made extensive use of means–ends analysis, similar to that observed in semantically less rich domains, supplemented by recognition mechanisms for accessing information in semantic memory.

In this paper we report on a study of the processes used to solve problems in engineering thermodynamics, of the kind normally encountered by chemical engineering majors in their sophomore year. The study is intended as a step toward extending the theory of human problem solving from the domain of laboratory tasks, requiring little specific semantic knowledge, to the kinds of semantically rich domains that are characteristic of professional problem solving.

We have used the general method of inquiry that has proved effective in recent years in problem-solving research: detailed analysis of thinking-aloud protocols of human problem-solving activities, and in parallel, the construction of computer programs aimed at simulating the human processes. The computer program that we have used is somewhat unusual, however, in that it is a hybrid between a program (a theory of behavior in the usual sense) for simulating the human subject and a program for coding his thinking-aloud protocol automatically. How

*This research was supported in part by Public Health Service Grant MH-07722 from the National Institute of Mental Health, in part by the Advanced Research Projects Agency of the Office of the Secretary of Defense (F44620-73-C-0074) which is monitored by the Air Force Office of Scientific Research, and in part by a grant from the Carnegie Corporation of New York. Reprints may be obtained from H. Simon, Department of Psychology, Carnegie-Mellon University, Pittsburgh, Pa. 15213.

these two functions are combined in a single program will be explained as we go along.

The study is organized into the following main sections: first, a fuller statement of the objectives and framework of the research; second, a characterization of the semantic domain of chemical engineering thermodynamics; third, a description of the computer program used for coding and simulating the subject's behavior; fourth, a description, in some detail, of the behavior of one subject; and finally, our comments and conclusions.

FRAMEWORK OF THE STUDY

For practical reasons, research on human problem solving has, until quite recently, been directed largely toward laboratory tasks that do not demand from the subject large amounts of prior semantic knowledge or task-related information. Nothing very specific in his previous learning or experience is of immediate relevance to the subject's performance on cryptarithmetic puzzles, so-called missionary and cannibal problems, or the Tower of Hanoi problem. Chess is a notable exception.

The state of the art in the analysis of human problem solving using the techniques of protocol analysis and simulation now permits the exploration of problem solving in domains that are at once well structured and semantically rich. Many college courses are good examples of such domains, and we have chosen one of them—chemical engineering thermodynamics—for our research.

It might be questioned whether there is anything special to be said about problem solving in particular problem domains. It is not implausible to suppose that the processes a person uses to solve problems in chemical engineering thermodynamics are the same ones he uses to play chess, or even to compose music. According to this view, problems are solved in any domain by using common problem-solving processes that then draw upon specific knowledge of that domain—and the specific knowledge is simply the information to be found in a good textbook on the subject. A theory of problem solving in chemical engineering thermodynamics would be constructed simply by juxtaposing the general, domain-independent theory of problem solving with a thermodynamics textbook.

There are large elements of truth in this position, but there are good reasons for suspecting that it is not the whole truth. First, if you hand a bright person who has not studied thermodynamics a textbook on the subject, he will not immediately be able to solve thermodynamics problems. He will not be able to do so after merely reading the textbook, even if he were to memorize its contents or to have constant access to it while working on the problems.

Second, if you were to store the contents of a thermodynamics textbook alongside a general problem-solving program in a digital computer, the computer would not thereby be enabled to solve thermodynamics problems. The information about the subject matter would have to be stored in such a way that it could be communicated to the problem-solving program, so that the latter could extract it as needed.

The question, then, is not so much *what* information must be available in solving problems in this domain, as *how* that information must be organized and stored in memory so as to be available to general problem-solving processes. To what extent is the information stored as data; to what extent as program? How must the information be indexed in memory so that particular items will be evoked on the occasions when they are relevant to the problem-solving effort? It is these kinds of questions about organization and representation of information in semantic memory that must be answered by a problem-solving theory for semantically rich domains.

The reasons for choosing thermodynamics as the domain for study were largely tactical. We had been engaged for some time in constructing programs capable of generating thermodynamics problems automatically, to be used in individually paced instruction in a course at Carnegie-Mellon University. Writing those programs gave us considerable insight into the structure of the subject; and indeed, it gradually dawned on us that the knowledge that had to be built into the programs to enable them to generate challenging problems for students was much the same as the knowledge that the students needed to solve the problems. Since thermodynamics appeared to be not atypical of a wide range of technical subjects that draw upon a substantial amount of well-structured information, it seemed an appropriate domain for our initial explorations.

The human data for this study were provided by six problem-solving protocols of a single subject who was a teaching assistant in a self-paced course in chemical engineering thermodynamics, and hence was reasonably proficient at solving these kinds of problems. We presented him with a sequence of six problems, which he solved in a single session of about two hours. Two of the problems were generated by our earlier programs; the other four were designed specifically to explore various facets of his handling of the task. The subject was asked to think aloud while he worked the problems with paper and pencil, and his verbalizations were recorded on tape. To help us analyze the protocols, and to increase the objectivity and consistency of our encoding of them, we constructed an interactive computer program that doubled as a semiautomated coding system and as a protocol simulator. We will say more about the program after we have described the task and the way in which the study was carried out.

THE STRUCTURE OF THERMODYNAMICS PROBLEMS

The reader is presumed to be familiar with the general character of thermodynamics, but he is not assumed to possess detailed knowledge of the laws of thermodynamics or knowledge of how to solve thermodynamics problems. In general, thermodynamics is the study of the transformation of energy from one form into another, especially, but not exclusively, the transformation of heat into work. Chemical engineering thermodynamics is particularly concerned with energy transformations in gases and liquids as they flow through devices or sequences of devices like compressors, pipes, turbines, pumps, and nozzles. Here is an example of a problem in chemical engineering thermodynamics, one of those we used in our study:

> Nitrogen flows along a constant area duct. It enters at 40°F and 200 psi. It leaves at atmospheric pressure and at a temperature of −210°F. Assuming that the flow rate is 100 lb/min, determine how much heat will be transferred to the surroundings.

In this problem, there is a single *process,* heat exchange, which takes place in a *device,* a duct of constant area, and produces a change in a *substance,* nitrogen. The substance enters the process in a particular input *state,* in this case, a particular temperature and pressure; it leaves in a different state, also characterized in this case by its temperature and pressure. Thus a substance has properties, and when certain of these have been specified (two, for a single homogeneous substance), its state is determined, the value of all its other state-determined properties can be calculated. Given the pressure and temperature of a quantity of gaseous nitrogen, one can compute the volume it occupies, the energy it contains, and certain other thermodynamic properties such as its enthalpy and its entropy. In general, the values of *any* two of the thermodynamic variables, or *state variables* as they are usually called, determine the values of the remainder.

The two most general and fundamental laws of thermodynamic analysis are the laws of conservation of mass and of energy. By the former, the quantity of a substance that leaves a process is equal to the quantity that entered, less the net transfer outside the system. By the latter, the total energy, of all forms, at the input to a process is equal to the total energy at the output, less the net transfer of energy to the system's surroundings. In conditions of steady-state flow, the mass conservation equation is replaced by the so-called "continuity equation," which performs the same function.

Since different units are used to measure energy in its different forms, solving thermodynamics problems requires a knowledge of the conversion ratios (for example, the ratio for conversion of foot pounds, which are units of work energy, into BTUs, which are units of heat energy). The laws of conservation of mass and energy are expressed as algebraic equations in consistent units.

Two additional kinds of relations besides the conservation laws enter into thermodynamics problems. The first of these are definitional equations (e.g., the definition of enthalpy in terms of internal energy, pressure, and volume). The second are the equations of state. An equation of state is a relation expressing any one of the state variables as a function of two of the others. Since, as stated earlier, any state variable can be expressed as a function of any two others, if there are N state variables one can write $N(N - 1)(N - 2)$ distinct equations of state. The equations of state are different for each working substance—water, air, steam, etc.—and they are often given in textbooks in the form of tables or nomographs rather than as explicit analytic equations. The familiar gas laws (e.g., volume is proportionate to temperature divided by pressure) are equations of state for ideal gases.

All of the relations described thus far are algebraic. Some simple applications of the integral calculus may also be required to solve thermodynamics problems. For example, work is measured by force times the distance through which it acts. To determine the work produced by a variable force, the increments of work must be integrated—that is, the integral of $f(x)dx$ must be calculated, where x is the distance traveled, and $f(x)$ is the variable force.

The important point in this description is that chemical engineering thermodynamics problems are basically algebraic in character. A problem is stated by supplying enough givens to determine—with the help of the conservation equations, equations of state, and definitions—the values of all of the remaining input and output variables. In the problem given above as an example, the facts given about the input state allow all of the remaining variables of that state, including the internal energy, to be calculated. Similarly the output energy can be calculated from the information given about the output state. Then, by the law of conservation of energy, the energy transferred out of the system is the simple difference between input and output energies. Although solving this problem requires solving three algebraic equations—two equations of state and the energy conservation equation—the equations need not be solved simultaneously, but can be handled one at a time. This is typical of thermodynamics problems to be found in textbooks.

Some of the information required to make a thermodynamics problem determinate is not given explicitly, but is implicit, the problem solver being supposed to be able to supply it from his store of semantic knowledge. In the sample problem presented above, no mention is made of the respective heights of the intake and output ends of the duct. Hence the change in potential energy, which occurs when a substance is raised or lowered in the gravitational field, cannot be computed. It is assumed by the question writer, and would presumably be assumed by the student working the problem, either that the output is at the same height as the input or, if not, that the potential energy difference is negligible and may be ignored. This convention cannot be deduced from the information given, but must be known.

Learning to solve problems in chemical engineering thermodynamics essentially means learning to handle the information and relations we have just described. It may seem from the description that there really is not very much to be learned, especially since the equations of state do not need to be memorized, but are always available to the problem solver as reference tables. Yet this subject is regarded by most students as one of the more difficult ones in the undergraduate curriculum.

THE EXPERIMENTAL PROBLEMS

The six problems that our subject was asked to solve were selected to explore various facets of the thermodynamics domain. We will describe some of the characteristics of the first three problems. Problem 2, which we used as an example in the last section, is stated at the top of Table 3, Problem 1 at the top of Table 4, and Problem 3 at the top of Table 5.

The first two problems are typical in certain respects of the problems found in textbooks. In each case, the energy equation is not required in its most general form, since only a few kinds of energy are transformed and the remainder can be omitted from the equation. Which energy components are significant for the problem is often signaled by particular terms in the problem text, which we will refer to as *keywords*. In Problem 1, for example, it is stated that water "falls into the river below." The word "falls," or perhaps the whole phrase, serves as a keyword to alert the subject to the fact that the potential energy of the working substance changes.

Another characteristic of the problems is that the givens are selected in such a way that the problem's equations need not be handled simultaneously, but, if taken up in the right order, can be solved one by one. Problem 3 was generated automatically, by a computer program, which selected givens and unknowns arbitrarily. Because of its mode of generation, it is also expressed in more stylized and abstract prose than are the others.

Solving all of the problems called for an understanding of the energy and mass conservation laws and different ways of stating them, of how to use definitions and equations of state, and of how to convert from one set of units to another, as well as the ability to solve simple algebraic equations.

Problem 1. This problem requires the subject to discard irrelevant information (the data on pressure and temperature, and the tables of properties of compressed water with which he was supplied), and conversely, to pay attention to changes in potential energy, which can be ignored in most thermodynamic problems.

Only the energy conservation equation and the definition of efficiency are needed to solve the problem. The subject must select appropriate units for his quantities, and must be able to convert from them to the units in terms of which he is to express the answer. An understanding of the meaning of efficiency is also required to complete the problem correctly.

Problem 2. This problem requires use of the energy conservation equation and state equations (in the form of tables to determine enthalpies from temperatures and pressures). By proper ordering of the work, the three relations can be determined sequentially, without solving simultaneous equations.

Problem 3. In this problem conservation of energy, conservation of mass, and the state equations all come into play. As we shall see, examples like this one, though atypical of the way in which problems generally arise in thermodynamics, give us a great many insights into the subject's representation of his thermodynamic knowledge and skill, since he has to use almost all of it.

METHOD FOR PROTOCOL ANALYSIS

Techniques for analyzing thinking-aloud protocols have been discussed by Newell and Simon (1972), and a system for automatic protocol analysis (PAS-II) has been constructed by Waterman and Newell (1972). In the present study, we have also availed ourselves of the help of the computer in analyzing the subject's problem-solving protocols, but by means of a semiautomated scheme that is aimed at imposing a discipline upon the analysis process while at the same time allowing the human coder to perform a sort of semantic parse and transcribe from the protocol most of the semantic knowledge needed for the encoding.

The SAPA System

The protocol encoding system SAPA (Semi-Automatic Protocol Analysis) may also be viewed as a weak theory of the problem-solving process. That is to say, the component routines of SAPA represent basic, but macroscopic, elements of the process: producing a relevant equation, evaluating a variable, solving an equation, and so on. The theory asserts that the human protocols will be made up of precisely these processes.

In SAPA, these basic processes are organized in a general, but flexible, strategy, in the form of a sequence of steps that is followed unless it is overridden by the coder, who interacts with SAPA while sitting at the computer terminal with the protocol before him. The coder has periodic options to depart from the sequence in order to imitate the actual sequence of the protocol. Moreover, the program does not itself perform algebraic manipulations, solve equations, or read the problem text, but returns control to the terminal so that these processes can be performed by the human coder.

Thus, the program is a theory and simulation of problem-solving processes in the sense that it postulates that the protocol can be encoded in terms of a definite fixed set of basic processes, and that these processes will usually follow one another in certain sequences. The program is an encoding scheme to the extent that it permits the interacting human coder to introduce the semantic information he finds in the protocol, and to modify, under certain conditions, the sequence of processes, to match those actually followed in a particular case.

To the extent that the processes embodied in the program are incomplete or inaccurate, this will be revealed by the difficulties encountered in encoding the protocol. To the extent that the normal process sequence is not followed by the subject, this will be revealed by the choices the coder is forced to make. Protocol and encoding can generally be compared sentence by sentence. We now give a more detailed description of the program.

The program is written in a version of the SNOBOL programming language known as SITBOL. Figure 1 shows the flow of control of the program in terms of its high-level subroutines.

The routine INQUIRY, not shown in the figure, returns control to the terminal whenever data are to be entered by the coder, including answers to questions that may cause changes in the flow of control.

The remaining routines of SAPA are executed in order from left to right, unless that order is changed by one of the replies in INQUIRY. SYSTEM is the first routine executed. It requests a description of the thermodynamic system involved in the problem and the question to be asked about it.

The second routine executed is ENERGYEQUATION, which obtains from the coder the energy equation in the form in which it is written down by the subject. At both the beginning and end of this routine, keywords may be entered, and on the basis of these, the form of the energy equation may be changed. For example, if the subject working Problem 1 noticed the keyword "fall" after writing down an initial energy equation, this keyword could be noticed, and a new energy equation, incorporating a potential energy term, then entered. Two routines called WORD and KEYWORD are used in these processes.

The third routine, SOLVEQUATION, is the executive routine for a whole set of procedures. This routine takes the energy equation, evaluates each of the variables that appear in it, and solves for the unknown variable. First, it calls subroutine ASSIGNVALUES, which presents the variables in the energy equation, in an order determined by the coder, and assigns values to each successively. Branch points in this routine return control to the coder, allowing him to interject comments, or to change the form of the energy equation. The values are not assigned by ASSIGNVALUES itself, but by a subroutine called ASSUMPTION. If the value of a variable is mentioned in a problem statement, this value is simply provided to ASSUMPTION by the coder and becomes its output. If the value is not stated, ASSUMPTION will attempt to supply it by

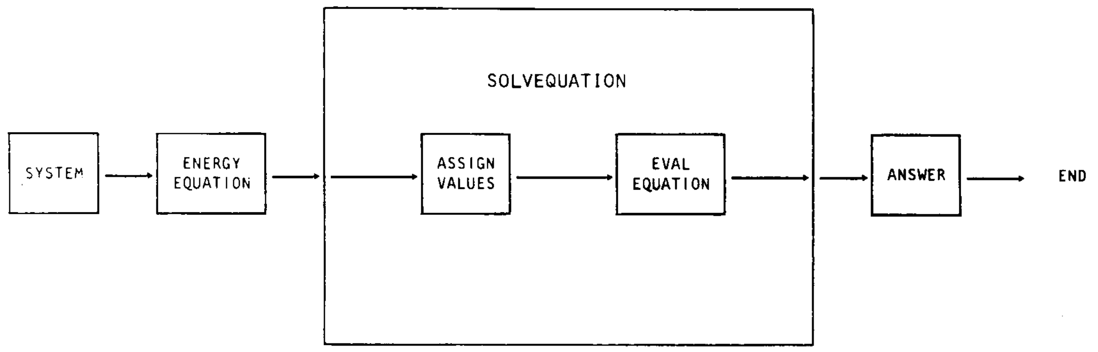

FIG. 1 Flow diagram of SAPA program.

asking the coder to look up tables representing equations of state, or to solve an appropriate subsidiary equation. The subsidiary equation may be solved numerically, or symbolically to provide an expression that can be substituted back into the main equation.

After ASSIGNVALUES returns a set of values to SOLVEQUATION, the latter process calls EVALEQUATION. First, a decision is made whether to check for appropriateness of the units or for keywords before finding the solution. Then the equation is solved, and a decision is made as to whether the entire problem has been solved, or whether additional equations remain to be solved. The system then outputs the ANSWER and terminates.

A Sample Output of SAPA

Table 1 shows a sample output, from the encoding of the protocol for Problem 2, which we used as an illustration earlier. The lines beginning with ** are requests by SAPA to the terminal for input; the lines without asterisks are the inputs provided by the coder in response to the requests. In Line 1, the coder is asked to describe the system (answer: duct, from Line 2 of the protocol, Table 3). In Line 2, he is asked if he wants to evoke keywords (No, an inference from Line 4 of the protocol). In Line 3, he is asked for the energy equation, which he supplies. In Line 4, he is asked again whether he wants to evoke a keyword (No). In Line 5, he is asked to type out the list of variables from the equation of Line 3. In Line 6, he is asked whether he wants to interpolate a comment (Yes). On Line 7, the comment is solicited. Line 8 asks whether he wants to continue with the current equation (Yes). He is next asked for the values of the successive variables in the equation. In the first three cases, the coder provides them, in the fourth (Line 18), he executes a carriage return, whereupon (Line 19), he is given the alternatives of calling the routine for deriving values from tables of state equations (TABLE), or the routine for solving subsidiary equations (SUBEQUATION). In this case, the coder asks for the tables. Line 19 asks what tables are to be consulted. Line 21 asks for the value read from the table. Line 23 returns to asking for the value of the next variable, which again (Lines 24,25,26) leads to a table lookup. Values having been provided for all variables in the equation, Line 28 asks whether the coder is ready to provide a solution. If not (Line 29), the coder is given an opportunity to branch elsewhere in the program. If, instead, he gives a solution, he is then asked (Line 30) whether the problem is solved. By an examination of the actual tape-recorded protocol in Table 3, it is possible to see that the coding requires only a relatively straightforward semantic parse of the protocol by the coder.

Table 2 shows the encoded protocol produced by this particular interchange. The words and phrases in capital letters paraphrase the questions that the system puts to the coder. The lines marked "COMMENT" are literal reproductions of comments provided by the coder. The remaining lines represent the information provided by the coder, as we have already seen in Table 1, embedded in

TABLE 1
Examples of Interactive Protocol Encoding—Problem 2

[1] **Describe the system.
(2) the system will be the duct
[2] **Do you want to look for keywords, Y or N?
N
[3] **What is the energy equation?
(5) $Q + Ws = m(h2-h1)$
[4] **Do you want to look for keywords, Y or N?
N
[5] **Please type in the list of variables
Q Ws m h2 h1
[6] **Do you want to make a comment, Y or N?
Y
[7] **Please make your comment now.
(6–7) I'm ignoring the changes in kinetic and potential energy.
[8] **Type C or R if you want to continue or return.
C
[9] **You can now give the value of *Q*, or you can type ⟨cr⟩.
(8–9) unknown, dependent variable.
[10] **Do you want to make a comment, Y or N?
Y
[11] **Please make your comment now.
(10) Since it is a duct, there is no shaft work.
[12] **You can give the value now of *Ws*, or you can type ⟨cr⟩.
(11) 0
[13] **Do you want to make a comment, Y or N?
Y
[14] **Please make your comment now.
(12) now $Q = m(h2-h1)$
[15] **Type C or R if you want to continue or return.
C
[16] **You can give the value now of *h*2, or you can type ⟨cr⟩
(13) 100 lb/min
[17] **Do you want to make a comment, Y or N?
N
[18] **You can give the value now of *h*2, or you can type ⟨cr⟩
#
[19] **Do you want to look up the value in a TABLE or by solving a subsidiary equation, LEQ? Type T or L.
T
[20] **What table do you have to look at to get the value of *h*2?
nitrogen tables
[21] **And what is the value you get?
(16–17,19–23,25–27), 126.443 BTU/lbm
[22] **Do you want to make a comment Y or N?
N
[23] **You can give the value now of *h*1 , or you can type ⟨cr⟩
#
[24] **Do you want to look up the value in a TABLE or by solving a subsidiary equation, LEQ? Type T or L
T
[25] **What tables do you have to look at to get the value of h1?
nitrogen tables
[26] **And what is the value you get?
(24,28–30) 187.408 BTU/lbm.
[27] **Do you want to make a comment, Y or N?
N
[28] **Do you want to go directly to a solution, or reevaluate the equation? Type D or R ⟨cr⟩
D
[29] **Do you want to return a solution right now, or do you want to check dimensions and equation? If you want to give a solution, type it in now, else type ⟨cr⟩
(35–37) −6096.5
[30] **If you think the problem has been solved, type S, otherwise type F
S

Note: Numbers in parentheses refer to line numbers in the protocol, Table 3. Numbers in square brackets refer to the actual coding sequence.

TABLE 2
Encoded Protocol—Problem 2

```
CHOOSE SYSTEM.
  The system is: (2) the system will be the duct.
  EQUATION: (5) Q + Ws = m(h2-h1).
  The variables are; Q Ws m h2 h1
COMMENT: (6-7) I'm ignoring the changes in kinetic and potential energy.
FIND Q.
  VALUE Q: (8-9) unknown, dependent variable
COMMENT: (10) Since it is a duct, there is no shaft work
FIND Ws
  VALUE Ws: (11) 0
COMMENT: (12) now Q = m(h2-h1)
FIND m.
  VALUE m: (13) 100 lb/min
FIND h2.
  READ TABLE.
  nitrogen tables: (16-17,19-23,25-27) 126.443 BTU/lbm
  VALUE h2: (16-17,19-23,25-27) 126.443 BTU/lbm
FIND h1.
  READ TABLE.
  nitrogen tables: (24,28-30) 187.408 BTU/lbm
  VALUE h1: (24,28-30) 187.408 BTU/lbm
SOLVE: (5) Q + Ws = m(h2-h1)
CHECK UNITS.
SOLVE: (5) Q + Ws = m(h2-h1)
  SOLUTION: (35-37) -6096.5
END
```

stereotyped English sentences to indicate the questions he was answering with each item.

Table 3 shows the raw protocol for comparison with Table 2.

SAPA as a Problem-Solving Theory

The structure of the subject's problem-solving effort postulated in the SAPA program was derived partly from the theory of human problem solving (Newell & Simon, 1972), and partly from an examination of the subject's protocol for Problem 2. The basic scheme that directs the subject's efforts is a form of means–ends analysis.

Means–ends analysis involves detecting a difference between the present state of the problem and the desired goal state, applying an operator that is relevant to reducing a difference of this kind, and repeating the cycle until the goal is reached. In problems like the ones before us, the goal is to determine the value of one or more variables that are initially unknown. Two main kinds of differences can be detected: (1) the value of a variable is unknown, and (2) an equation is

TABLE 3
Protocol for Problem 2

Problem 2

Nitrogen flows along a constant area duct. It enters at 40°F and 200 psi. It leaves at atmospheric pressure and at a temperature of −210°F. Assuming that the flow rate is 100 lb/min, determine how much heat will be transferred to the surroundings.

Protocol

1. OK, the first thing I'm going to do is pick a system,
2. that is, the system will be the duct.
3. OK, I draw that like this,
4. and I'm going to write the first law on this duct,
5. as Q plus Ws will equal m times $h2$ minus $h1$,
6. where I'm ignoring the changes in kinetic and potential energy.
7. And this is probably a pretty good assumption.
8. OK, I'm asked to determine how much heat will be transferred to the surroundings.
9. OK, that will be the Q term here.
10. Since we just have a duct here, there will be no shaft work,
11. so Ws will equal zero.
12. Q then will simply equal m times $h2$ minus $h1$.
13. OK, m I know as 100 lb/min,
14. and $h2$ minus $h1$,
15. in order to determine that, I will need some physical properties for nitrogen.
16. So let me look these up.
17. Found these.
18. OK, so let me put m is 100 lb/min.
19. OK, $h2$, 2 is downstream,
20. so $h2$ is the enthalpy at one atmosphere and −210°F.
21. So let me look that up.
22. −210°, OK, this is in degrees Rankine, −210°, ah . . .
23. OK, is 250 degrees Rankine, this is $T2$.
24. While I'm at it, I'll just note that $T1$ will be 500°R.
25. OK, so 14.7 lbf/cu. in.,
26. at a temperature of 250 degrees Rankine.
27. I read h as 126.443 BTU/lbm, that's $h2$.
28. Now as $h1$ I have 200 psia and 500 degrees Rankine.
29. Let me look that up,
30. and I read $h1$ as 187.408 BTU/lbm.
31. OK, I'm simply going to do the calculation.
32. I see that the pound masses cancel as they should,
33. and my final answer will be in BTUs/min,
34. which is what I'd expect.
35. Let me do the calculation now.
36. 126.443 − 187.408
37. and when I multiply that by 100 and I get −6096.5,
38. the negative sign is as it should be,
39. because it indicates that heat is being transferred out of the system.
40. That's it.

unsolved. When the value of a variable is unknown, the relevant operators are three in number: (1) to find the value among the givens in the problem statement, (2) to determine the value from an appropriate table or nomograph, or (3) to find an equation containing the variable that can be solved for it. When an equation is unsolved, the relevant operators are: (1) to replace one of the independent variables in the equation with its value, and (2) if all independent variables have been evaluated, to solve the equation for the dependent variable.

It can be seen that this means–ends scheme corresponds closely to the main components of the SAPA program. ASSIGNVALUES detects the independent variables in an equation requiring evaluation, and calls on ASSUMPTION to determine the value of each. ASSUMPTION either finds the value directly among the givens, or calls upon TABLES or SUBEQUATION to determine it. SUBEQUATION selects an equation containing the unevaluated variable, and calls on SOLVEQUATION recursively.

However, this machinery is not quite adequate for handling thermodynamics problems. There must also be some means for supplementing the values of variables given explicitly in the problem statement with values of others that are determined implicitly by the conditions of the problem. The initial routines, SYSTEM and ENERGYEQUATION, perform this function (in part). From the nature of the system to which the problem refers, it usually can be inferred that certain variables may be set equal to zero, or assigned certain "default" values (e.g., if the system is a duct, the change in potential energy may be assumed to be zero; if the discharge conditions of a device are not mentioned, a pressure of one atmosphere and a normal air temperature may be assumed). These assumptions determine the particular form of the energy equation that should be postulated.

The SAPA program has one additional mechanism for detecting and correcting errors: the routine for checking the consistency of the units in which the variables are expressed. This mechanism lies outside the main means–ends scheme, and it would never have to be evoked if the system were entirely error-free.

While this program may seem "self-evident," it is not actually deducible, in this form, from the idea of means–ends analysis alone. Within the general framework of means–ends analysis, there are many other procedures a subject might follow in attacking these problems. Let us just enumerate a few of the major and minor variants that are conceivable:

1. The subject might construct at the outset a plan for his solution, determining in what order he is going to solve for the various variables and what equations he is going to use, before actually attempting the numerical solution.
2. The problem might be worked backward from the variable whose value is demanded in the problem statement. An equation (not necessarily the energy

equation) involving this variable would be evoked, and the other variables in this equation would be taken as independent variables. Thus, the procedure would not have to start with the energy equation, as SAPA does.

3. Instead of searching at the outset for the assumptions, implicit in the problem statement, that are required to make it solvable, the subject might make these assumptions sequentially, as he comes to deal with the relevant variables and equations.

4. SAPA is ambiguous about the form in which the energy equation is written initially. The subject might either write it in a standard form, and then specialize it to the particular assumptions of the problem, or he might create a tailor-made version of the energy equation that already incorporates the special assumptions of the problem.

5. SAPA permits the variables in an equation to be evaluated in various orders: for example, in order of their appearance in the equation, or in the order in which they are mentioned in the problem statement.

6. Subsidiary equations, instead of being solved numerically, might be solved algebraically and the resulting expression substituted fot the dependent variable in the main equation.

7. The subject might check his units at each step, instead of making a special check as a last step in solving the main equation.

When we come to examine the details of our subject's behavior, we will see that he sometimes diverges from the scheme of SAPA in ways suggested by these variants.

It is not as easy to conjure up possible procedures for solving these kinds of thermodynamics problems that make no use at all of the generalized means–ends analysis we have described above. The most radical departure we have been able to imagine is a scheme in which the subject writes down as a separate equation each fact in the problem statement, adds equations for the assumptions implicit in the statement, adds the equations of thermodynamic theory (energy, continuity, and equations of state), and then uses some standard systematic method to solve the resulting system of simultaneous equations. (We do not take into account schemes that would require inordinate amounts of search, such as schemes for generating possible solutions at random, and then testing them.)

What SAPA provides for us is a framework—a zeroth-order approximation to the subject's behavior—that permits us to detect where that behavior follows the scheme and where it deviates, and to diagnose the causes of the deviations. When we come to examine the subject's protocols in detail in the next section, we are not "testing" SAPA as a theory of the subject's behavior, but using it as a tool to help us induce what his actual program is. Hence, we will not be concerned with statistical tests of goodness of fit, or other quantitative measures of deviation. Rather, we will be interested in the qualitative features that the subject's behavior reveals.

THE SUBJECT'S BEHAVIOR

We are now ready to use the SAPA program as our vehicle for describing the subject's behavior while he solved the first three problems. We will begin with Problem 2, which was used in developing SAPA and which we have already looked at briefly. We will then take the remaining two problems.

Problem 2: Nitrogen Flow in a Duct

Table 2 is the encoded protocol for Problem 2. Since this protocol was used as a guide in developing the coding scheme, it is not surprising that its content is captured quite accurately and completely in the encoding, with little departure from the basic structure of SAPA. We need make only a number of comments on matters of detail.

The subject immediately writes down the energy equation (5) in his "standard" form, explicitly (6–7) omitting the kinetic and potential energy terms. Since the problem asks him to find the amount of heat transferred, Q, and since this variable appears in the energy equation, this is the equation one would write down first, even if working backward by strict means–ends analysis.

He then proceeds to evaluate the independent variables in the energy equation, working from left to right exactly in the order in which he has written them. When he comes to determine the value of W, the work performed, he reviews the problem statement ((10) "we just have a duct here"), and infers from the lack of mention of a work-performing device that the work is zero—a good example of implicit information.

He finds (13) the value of the mass flow in the problem statement. When he comes (14) to the enthalpies ($h2$ and $h1$), from his knowledge of thermodynamics he is aware that, given two thermodynamic properties of the input state (temperature and pressure) and two of the output state, he can determine $h2$ and $h1$ from standard thermodynamic tables for nitrogen. It can be conjectured that his associations flow about as follows: enthalpy to be determined; this is a variable of state; values of what other variables are known for the same state; check problem statement; temperature and pressure are known; use tables to find enthalpy as function of temperature and pressure.

At each step of his calculations, the subject carries along the units in which his variables are expressed. Not noted in the encoding is his conversion of temperature from Fahrenheit to Rankine units in order to use the nitrogen tables, which are stated in terms of the latter units. Similarly, he checks whether the units of the quantities he has substituted into the energy equation are consistent before he solves the equation numerically. These checks of dimensionality are a principal means he uses for detecting errors in his equations.

One detail of the protocol (Table 3) not captured in the encoding is worth mention. In Line 13 of the protocol, the subject mentions that he knows the value

TABLE 4
Encoded Protocol—Problem 1

Problem 1

2000 lb of compressed water at 200 psi and at 200°F fall through a tube 50 ft long every minute. It turns a wheel at the bottom and falls into the river below. If the wheel has a mechanical efficiency of 99% and it delivers power to a generator with an efficiency of 80%, how much power will the generator deliver?

CHOOSE SYSTEM
The system is: (2) water, (3) a wheel
NOTICE (11) enthalpy
ASSOCIATE TO (12) adiabatic
EQUATION: (14) $Ws = m(h2-h1)$
The variables are: m $h2$ $h1$
COMMENT: (15) I've assumed potential and kinetic energy negligible
REVISE EQUATION
NOTICE (16) falling
ASSOCIATE TO (16) potential energy
WRITE EQUATION
EQUATION: (18–20) $Ws = m(h2-h1) + gH$
The variables are: m $h2$ $h1$ H
FIND m
VALUE m: [known]
FIND $h2$
READ TABLE
(25) steam tables: (23–26) unknown, insufficient information
VALUE $h2$: (23–26) unknown, insufficient information
COMMENT: (27–28) therefore assume enthalpy change is zero
REVISE EQUATION
NOTICE (34) potential energy per lbm
ASSOCIATE TO (35) flow rate
EQUATION: (36) $Ws = mgH$ [which he interprets as $Ws = VdgH$—see (98–103)]
The variables are: m d g H
FIND m
VALUE m: (38) known, 2000 lb/min
FIND d
READ TABLE
(39) steam tables: (41–65) 1/0.0166 lbm/cu.ft
VALUE d: (41–65) 1/0.0166 lbm/cu.ft
FIND g
VALUE g: (66–67) 32.17 ft/sec²
FIND H
VALUE H: (69) 50 ft
SOLVE: (36) $Ws = mgH$
CHECK UNITS
CHECK FAILS
ASSOCIATE TO lbf vs. lbm
SOLVE: (36) $Ws = mgH$
CHECK UNITS
CHECK FAILS
ASSOCIATE TO (100) I interpreted Ws as volumetric flow rate instead of mass flow rate
SOLVE: (36) $Ws = mgH$
SOLUTION: (98) rewrite energy equation
SOLVE (99) energy equation for Ws
The variables are: m $g(c)$ H
FIND m
VALUE m: (106) 2000
FIND g
VALUE g: (107) 32.17
FIND $g(c)$
VALUE $g(c)$: (107) 32.17
FIND H
VALUE H: (107) 50
SOLVE:
SOLUTION: (108–109) 100,000 ft lbf/min
SOLVE (111–114) $W = Ws \times Em \times Eg$
The variables are: Ws Em Eg
FIND Ws
VALUE Ws: (117) 100,000
FIND Em
VALUE Em: (114) 99%
FIND Eg
VALUE Eg: (115) 80%
SOLVE:
SOLUTION: (118) 79,200 ft lbf/min
END

of m (from the problem statement). In Line 18, he interrupts his search for the values of $h2$ and $h1$ in order to substitute this numerical value of m in the equation. This kind of event occurs fairly frequently in the protocols, and there is nothing in the system to predict exactly when he will begin writing down information, as distinct from noting that he has it.

Problem 1: Water Power

Table 4 shows the problem statement and the encoded protocol. The subject found himself "tricked" by this problem in several respects. First, he initially wrote down his "standard" energy equation (14), ignoring the change in potential energy of the falling water which is the core of the problem. Second, he was misled by the mention of temperature and pressure into thinking that there would be a significant change in enthalpy. He very quickly noticed the omission of potential energy (16), and revised the equation (18–20) to accommodate it. Only after he had attempted to calculate the enthalpy change (23–26) did he discover that insufficient facts were mentioned in the problem statement. He then decided to assume that the change was negligible. (Alternatively, he could have assigned "default" values to the output temperatures and pressure of the water, and there is nothing in the protocol to indicate why he made the other choice. The amount of the change would in fact have been negligible.)

Having dealt with these difficulties, the subject now begins to write down the simplified energy equation (36), but falls into a new error. He assumes, contrary to the explicit language of the problem, that he is given the volumetric flow rate instead of the mass flow rate, and that he must therefore convert the given volume of water into its mass. Since the problem speaks of "compressed" water, he looks up (39) the density of the water at the given temperature and pressure in the tables. He discovers his mistake (100) only after he has substituted all numerical values in the equation and is checking the dimensionality of the result. As soon as he notices the error, he arrives at the correct result.

The subject's behavior on this problem provides strong evidence for the general validity of the SAPA simulation. The problem is exceedingly simple if approached as a problem in mechanics rather than thermodynamics. The work done by falling water is simply the product of the weight of the water times the distance through which it falls. The subject, however, goes through his standard "thermodynamics" routine, writing an equation containing an irrelevant term and omitting a relevant one, then adding the omitted term and deleting the unneeded one. The program does not, however, predict the error he will make in units of measurement.

Problem 3: Air Flow

The problem statement and coded protocol are shown in Table 5. The problem, which was generated automatically, is an intricate one. The variable to be solved for is the outlet pressure, which does not appear in the energy equation.

TABLE 5
Encoded Protocol—Problem 3

Problem 3

The working fluid of a flow system is air. The work done is 592 lbf ft/sec. The inlet temperature is 46.1 °F. The outlet velocity is 15.8 ft/sec. The heat input is 154.97 BTU/sec. The inlet area is 4.6 sq ft. The inlet pressure is 16.8 psia. The outlet area is 36 sq ft. The outlet specific volume is 0.32 cu.ft/lbm. What is the outlet pressure?

CHOOSE SYSTEM
The system is: (4–5) flow system, air
NOTICE (10) outlet velocity
ASSOCIATE TO (12–15) kinetic energy
WRITE EQUATION
NOTICE (18) air
ASSOCIATE TO (19) potential energy
EQUATION: (16,21)$Q + Ws = m(h2-h1) + m(u2**2-u1**2)/2g$
The variables are: *Ws u*2 *Q P*1 *m h*2 *h*1 *u*1
COMMENT: (22) OK, so there are a few things that I know
FIND *Ws*
VALUE *Ws*: (23) 592 ft lbf/sec
COMMENT: (24) Let me underline in this equation to indicate I know that
FIND *u*2
VALUE *u*2: (26) known
FIND *Q*
VALUE *Q*: (27) known
FIND *P*1
VALUE *P*1: (28) known
COMMENT: (29) the rest are physical. I know some dimensions and some physical properties
FIND m
FIND EQUATION FOR IT
(31,31.1,33) equation of continuity: (34–35) $m = d1A1u1 = d2A2u2$
SOLVE IT
The variables are: *u*2 *A*1 *P*1 *A*2 *d*1 *u*1 *d*2
COMMENT: (36) Let me underline the terms I know in here
FIND *u*2
VALUE *u*2: known
FIND *A*1
VALUE *A*1: (38) known
FIND *P*1
VALUE *P*1: (39) known
COMMENT: (40) that'll be physical properties
FIND *A*2
VALUE *A*2: (41) known
FIND *d*1
FIND EQUATION FOR IT
(45–46) ideal gas law: (47) $d1 = P1M/RT1$
SOLVE IT
The variables are: *P*1 *M R T*1
FIND *P*1
VALUE *P*1: (44,48) 16.8 psia
FIND *M*
VALUE *M*: (49) 29
FIND *R*
VALUE *R*: (50) 10.73
FIND *T*1
VALUE *T*1: (43,51) 46.1 *F* = 506.1°R
SOLVE: (45–46) ideal gas law
SOLUTION: (53–54) .09 lb/cu.ft
VALUE *d*1: (53–54) .09 lb/cu.ft
FIND *u*1
VALUE *u*1: (63) unknown
FIND *d*2
VALUE *d*2: (56,63) unknown
COMMENT: (57–59) Don't know *h*2, so can't compute *d*2 from State2
REVISE EQUATION
No value computed for *m*
VALUE *m*:
COMMENT: [(62,63) can't yet solve continuity equation]
REVISE EQUATION
NOTICE (60,64) *h*1
ASSOCIATE TO (65–66) reference basis for *h*2
WRITE EQUATION
NOTICE relations
ASSOCIATE TO (82) substitute for *u*1 in terms of *d*2
EQUATION: (85) $u1 = d2A2u2/d1A1$
The variables are: *u*1
FIND *u*1
VALUE *u*1: (82) substitute for *u*1 in terms of *d*2
COMMENT: (84) substitute (85) in energy equation

TABLE 5 (*continued*)

REVISE EQUATION
NOTICE (86–89) expressions for $h2$ and m
ASSOCIATE TO (87–89) in terms of $h1$, which is known, and $d2$
EQUATION: (91–94) $Q + Ws = (d2A2u2)Cp(T2-T1) + (d2A2u2)(u2{*}{*}2-u1{*}{*}2)/2g$
The variables are: $T1$ $T2$ $u1$ $d2$
FIND $T1$
VALUE $T1$: (94) known
FIND $T2$
FIND EQUATION FOR IT
(96) ideal gas law, $d2 = P2M/RT2$: (97–98) $T2 = P2M/Rd2$
SUBSTITUTE INTO SUPERORDINATE EQUATION
The superordinate equation with substitution is: (99–100) $Q + Ws = (d2A2u2)Cp(P2M)/Rd2-T1) + (d2A2u2)(u2{*}{*}2-u1{*}{*}2)/2g$
$T2$ eliminated by substitution
FIND $u1$
VALUE $u1$: (100.4) unknown
COMMENT: I've got one expression, and let me write down the other one I have
FIND $d2$
VALUE $d2$: (105–106) $d2$ known
COMMENT: (107–111) Since $d2$ is known, I can get $u1$ from continuity relation.
REVISE EQUATION
NOTICE (116) outlet pressure
ASSOCIATE TO (116) outlet specific volume
WRITE EQUATION
NOTICE (123) only unknown will be $T2$
ASSOCIATE TO (124) ideal gas law, solve for $P2$
EQUATION: (121–121.3) $Q + Ws = (d2A2u2)Cp(t) + (d2A2u2)(u2{*}{*}2-(d2A2u2/d1A1){*}{*}2)/2g$
The variables are: Q Ws $d2$ $A2$ $u2$ Cp t $u2$ $u1$ g
FIND Q
VALUE Q: (127) 154.97
FIND Ws
VALUE Ws: (129–131) −592 ft lbf/sec
COMMENT: (131) divide Ws by 778 to convert to BTUs
FIND $d2$
VALUE $d2$: (132) 1/.32 lbm/cu.ft
FIND $A2$
VALUE $A2$: (133) 36 sq. ft
FIND $u2$
VALUE $u2$: (133.1) 15.8 ft/sec
FIND Cp
VALUE Cp: (133.2) 7
COMMENT: (134–135) divide by 29 to get BTU/lbm °F
FIND t
VALUE t: (136–136.1) $T2-T1$
FIND $u2$
VALUE $u2$: (136.2) 15.8 ft/sec
FIND $u1$
FIND EQUATION FOR IT.
(136.3) continuity equation: [$u1 = m/d1A1$]
SOLVE IT
The variable are: $u1$
FIND $u1$
VALUE $u1$: [will solve later (136.3)]
SOLVE: (136.3) continuity equation
SOLUTION: [will solve later]
VALUE $u1$: [will solve later]
FIND g
VALUE g: known
COMMENT: (138) compute $u1$
REVISE EQUATION
EQUATION: (127–136.4)
The variables are: $u1$ g t
FIND $u1$
VALUE $u1$: (140–142) 4293
COMMENT: (142–144) A high number, check against change in cross section, OK
FIND g
VALUE g: (146–146.1) 32.17
FIND t
VALUE t: (149–151) 1193505
COMMENT: (149,152) a very large number
SOLVE: energy equation
SOLUTION: (154) $T2 = 1193550$
SOLVE (155) ideal gas law FOR (115) $P2$
The variables are: $T2$ R $d2$ M
FIND $T2$
VALUE $T2$: (157) known
FIND R
VALUE R: (159) .73
FIND $d2$
VALUE $d2$: (161) 1/.32
FIND M
VALUE M: (162) 29

TABLE 5 (*continued*)

COMMENT: (163) final answer is in atmospheres	SOLVE: SOLUTION: (164) 93922 atmospheres END

Ultimately, it will have to be obtained by solving an equation of state for the output—in this case the equation for an ideal gas that expresses pressure in terms of temperature and density (or specific volume). The output specific volume is given, but the temperature can be obtained only by solving the energy equation. That equation is presumably solvable, because the input state is known (the pressure and temperature are given), and the work done and the heat input are both given. However, the mass flow must also be calculated from the equation of continuity; which can be done, since output area, velocity, and specific volume are all given.

It does not appear from the protocol that the subject has such a comprehensive plan in mind when he starts to solve the problem. Instead, he begins as usual by writing down the energy equation (16), deciding, on the basis of the nature of the system and the givens, that he can neglect potential energy changes (19) but must include kinetic energy terms (12–15). His next step is to check through the energy equation (22–29), not attempting to substitute numerical values for the variables (28) that appear in it, but simply checking to determine which of the variables are given. The order in which he does this shows that he is working *from* the problem statement *to* the equation, and not vice versa, for he considers the variables in the order in which they are mentioned in the statement and he mentions variables that appear in the statement, but not in the equation. As he checks off each variable, he underlines it in the equation.

When he has finished this step, it is clear that the energy equation cannot be solved immediately, for four variables—mass flow, input enthalpy, output enthalpy, and input velocity—are unknown. He turns (31) to the first of these in the order of appearance in the equation, mass flow, and proposes to find its value by solving the equations of continuity (34–35). Again (36–41) he checks off the variables in the order in which they appear in the problem statement, and underlines the givens in the equations. When he comes to the input specific volume, which is unknown, he interrupts his checking (45–46) in order to evaluate this variable by solving (47) the input equation of state for an ideal gas (pressure and temperature are given). He fails to return to checking and underlining the givens in the continuity equation, and as a result, mistakenly supposes (56,63) that both input velocity and output specific volume are unknown—in fact the value of the latter is given.

As a result of this oversight, the subject goes through the process of solving the continuity equations for mass flow and input velocity in terms of the supposedly unknown output specific volume (82–89), and substituting the resulting expressions in the energy equation (91–94). He uses the ideal gas laws to find temperature in terms of specific volume (96–98), but is left with an energy equation with two unknowns, specific volume and output temperature (99–100.4). He now realizes (105–106) that output specific volume is known, and quickly backtracks. This time, before proceeding, he states a plan (107–124): to use the known value of the specific volume to compute the mass flow and input velocity from the continuity equations, then solve the energy equation (123) for the outlet temperature, then use that and the specific volume (with the ideal gas law) to compute output pressure (124). He then carries out this plan.

This analysis is interesting in demonstrating that the subject's basic strategy is one of working forward from the energy equation, not working backward from the goal. He takes up the variables of the equation to see whether all but one are known; when this condition is not satisfied, he looks for subsidiary equations that he might solve for the unknowns. When he finds himself with two unknowns in a subsidiary equation, instead of one, he expresses one in terms of the other, and substitutes the resulting expressions back in the main equation. Near the very end, when he has discovered his important oversight of a given, he shifts strategy and works out a complete plan before again trying to solve any equations.

Of course, checking off the known and unknown variables in an equation before trying to insert numerical values is itself a form of (partial) planning. Neither of the planning procedures the subject uses are anticipated by the SAPA program, although it is not difficult to accommodate to them in encoding the protocol. Planning of the one or other kind appears to be evoked by the subject when the problem before him crosses some threshold of difficulty—presumably when the strategy of eliminating variables successively appears to fail.

Summary

Some feeling for the degree of fit between the method postulated by the SAPA program and the protocols can be obtained from Tables 2, 4, and 5. To the extent that the processes defined in the program account for the activities evident in the protocols, most of the lines in the protocol will be mentioned explicitly in the encoding, and comments will be relatively infrequent, occurring mainly when the subject makes his assumptions explicit. If the sequence prescribed by SAPA is generally followed by the subject, then the lines of protocol will occur in orderly sequence in the encoding. If the subject verbalizes fairly completely what he is doing, then not many items will appear in the encoding that cannot be referred to specific lines in the protocol. Examination of the encodings in terms of these criteria shows that SAPA does indeed provide a zeroth-order approximation to the subject's behavior, and that the encoding procedure does permit us to detect clearly whatever digressions occur from the predicted sequence.

The program does not predict the subject's errors, and it is mainly in connection with these errors that we see deviations from the general procedure. The subject does inject several new strategies when the problems become especially difficult (e.g., in Problem 3). In particular, before solving an equation, he checks which of the variables are already known and which are unknown. Several times midway in the more difficult problems, he also sets up an explicit final goal, and announces a more or less explicit plan for reaching it. Otherwise, he mainly follows a working forward procedure of working from the energy equation (and the equations of continuity, when necessary), and solving successively for the values of variables until the equations themselves can be solved for a single dependent variable.

CONCLUSION

Our analysis of the problem-solving behavior of a subject solving problems in chemical thermodynamics shows that he follows a consistent pattern of approach that might be described as a form of means–ends analysis modified by his knowledge of the central role that the conservation of energy equation plays in such problems.

The method of means–ends analysis that he uses in this context is not specialized to thermodynamics problems, but is a content-free scheme for solving systems of algebraic equations—and, in particular, systems in which the equations can be evaluated one by one for a single dependent variable, and the value of that variable substituted in the remaining equations. When the goal is to find the value of a variable, then the means employed is to find an equation containing that variable and to solve it for the variable as the dependent variable. When the goal is to solve an equation, the means employed is to evaluate numerically each of the independent variables in turn, then evaluate the equation. The subject sometimes applies this scheme recursively, evaluating an independent variable by setting up and solving a subsidiary equation, then returning to his place in the original equation and continuing.

The subject does not use means–ends analysis quite consistently, however. At the outset, he does not first establish the goal of writing an equation having the quantity to be found as dependent variable. Instead, he routinely writes down the energy conservation equation. He appears to have a "standard form" for this equation, which he modifies, adding or eliminating terms on the basis of information given in the problem statement.

The subject's knowledge of thermodynamics enters into the solution process in several ways. First, he is able to deduce that certain variables can be ignored or set equal to zero on the basis of the language of the problem statement. "Key words" in the statement appear to evoke this information from memory. Second, when faced with a variable to be evaluated, he generally is able to evoke from memory an equation containing that variable which might be solved for it.

The subject uses checks on the dimensionality of equations as an important means for detecting and correcting errors in writing equations. This requires that he have associated with each variable the units in terms of which the variable is expressed.

Most of the subject's departures from the simple scheme of means–ends analysis described above can be attributed either to his efforts to recover from errors, or to planning activities he undertakes when the simple variable-elimination scheme appears not to be working. Planning may consist simply of checking ahead to see whether a sufficient number of values is known to solve an equation, or may consist in checking out a solution procedure for the remainder of a problem.

In general, the subject's problem-solving program resembles closely the programs that have been observed in task environments having less rich semantic content—for example, in solving cryptarithmetic puzzles (Newell & Simon, 1972). The semantic information takes the form both of data structures in long-term memory (e.g., known thermodynamic equations) or of procedures (e.g., procedures for modifying the energy equation, or for looking up a thermodynamic variable in a table). An important part of the equipment for making use of this semantic information is a recognizing or evoking mechanism that retrieves the information from long-term memory at appropriate times during the problem-solving process.

On the methodological side, the analysis demonstrates the value of a semiautomated protocol analysis system, SAPA, that provides a zeroth-order approximation to a theory of the subject's behavior, at the same time that it provides a framework to guide and formalize the encoding process.

ACKNOWLEDGMENTS

We are undebted to John R. Hayes for many hours of patient discussion, and to Allen Newell for comments on an earlier draft.

REFERENCES

Newell, A., & Simon, H. A. *Human problem solving*. Englewood Cliffs, N.J.: Prentice-Hall, 1972.

Waterman, D. A., & Newell, A. PAS-II: An interactive task-free version of an automatic protocol analysis system. Pittsburgh: Department of Computer Science, Carnegie-Mellon University, 1972.