COGNITIVE SCIENCE 3, 275–310 (1979)

A Cognitive Model of Planning*

BARBARA HAYES-ROTH
AND
FREDERICK HAYES-ROTH
The Rand Corporation, Santa Monica, California

This paper presents a cognitive model of the planning process. The model generalizes the theoretical architecture of the Hearsay-II system. Thus, it assumes that planning comprises the activities of a variety of cognitive "specialists." Each specialist can suggest certain kinds of decisions for incorporation into the plan in progress. These include decisions about: (a) how to approach the planning problem; (b) what knowledge bears on the problem; (c) what kinds of actions to try to plan; (d) what specific actions to plan; and (e) how to allocate cognitive resources during planning. Within each of these categories, different specialists suggest decisions at different levels of abstraction. The activities of the various specialists are not coordinated in any systematic way. Instead, the specialists operate opportunistically, suggesting decisions whenever promising opportunities arise. The paper presents a detailed account of the model and illustrates its assumptions with a "thinking aloud" protocol. It also describes the performance of a computer simulation of the model. The paper contrasts the proposed model with successive refinement models and attempts to resolve apparent differences between the two points of view.

1. INTRODUCTION

Planning is a familiar cognitive activity. We all have many opportunities to decide how we will behave in future situations. For example, we plan how to get to work in the morning, where and with whom to eat lunch, and how to spend our evenings. We also make longer-term plans, such as what to do on our vacations, how to celebrate Christmas, and what career path to follow. Thus, planning influences many activities, from the most mundane to the most consequential, in everyday life.

We define planning as the predetermination of a course of action aimed at

*ONR Contract N00014-78-C-0039, NR 157-411 supported this research. We thank Bob Anderson, Ed Feigenbaum, Penny Nii, Perry Thorndyke and members of the Rand Cognitive Sciences Brownbag for many valuable discussions of this research. Bob Anderson, Bill Faught, Phil Klahr, Stan Rosenschein, and Bob Wesson provided useful comments on an earlier version of this manuscript. We thank Allan Collins for outstanding editorial assistance. Send reprint requests to: Barbara Hayes-Roth, The Rand Corporation, 1700 Main Street, Santa Monica, California, 90406.

achieving some goal. It is the first stage of a two-stage problem-solving process. The second stage entails monitoring and guiding the execution of the plan to a successful conclusion. We refer to these two stages as *planning and control*. This paper focuses on the planning stage of planning and control. We have two main objectives: to characterize the planning process and to propose a theoretical account of it.

Sacerdoti's (1975) work is probably the best-known previous research on planning. His computer program, NOAH, implements a successive refinement approach to planning. NOAH formulates problems in terms of high-level goals that specify sequences of actions (for example, the monkey should get the bananas and then eat them). NOAH expands each constituent subgoal into additional subgoals, maintaining any indeterminate sequential orderings as long as possible. In this manner, NOAH eventually generates correct plans specifying sequences of elementary actions. When executed, these actions transform initial conditions into a series of intermediate conditions, culminating in the goal state. (See also: Ernst & Newell, 1969; Fahlman, 1974; Fikes, 1977; Fikes & Nilsson, 1971; Sacerdoti, 1974; Sussman, 1973).

While not incompatible with successive-refinement models, our view of planning is somewhat different. We share the assumption that planning processes operate in a two-dimensional planning space defined on time and abstraction dimensions. However, we assume that people's planning activity is largely *opportunistic*. That is, at each point in the process, the planner's current decisions and observations suggest various opportunities for plan development. The planner's subsequent decisions follow up on selected opportunities. Sometimes, these decision-sequences follow an orderly path and produce a neat top-down expansion as described above. However, some decisions and observations might also suggest less orderly opportunities for plan development. For example, a decision about how to conduct initial planned activities might illuminate certain constraints on the planning of later activities and cause the planner to refocus attention on that phase of the plan. Similarly, certain low-level refinements of a previous, abstract plan might suggest an alternative abstract plan to replace the original one.

In general, the assumption that people plan opportunistically implies that interim decisions can lead to subsequent decisions at arbitrary points in the planning space. Thus, a decision at a given level of abstraction, specifying an action to be taken at a particular point in time, may influence subsequent decisions at higher or lower levels of abstraction, specifying actions to be taken at earlier or later points in time.

This view of the planning process suggests that planners will produce many coherent decision sequences, but some less coherent sequences as well. In extreme cases, the overall process might appear chaotic. The relative orderliness of particular planning processes presumably reflects individual differences among planners as well as different task demands.

We have tried to develop a theoretical framework that can accommodate both systematic approaches to planning, like successive refinement, and the more generally opportunistic process described above. The next section of the paper presents a "thinking-aloud" protocol that illustrates the kind of behavior the model must explain. Section 3 describes the proposed planning model. Section 4 shows how the model could produce the thinking-aloud protocol. Section 5 describes a computer implementation of the model and compares its performance to the performance of the human subject. Section 6 addresses questions of theoretical complexity. Section 7 attempts to resolve apparent differences between the proposed model and successive refinement models. Section 8 summarizes our conclusions.

2. PLANNING A DAY'S ERRANDS

The thinking aloud protocol in Figure 1 illustrates the kind of behavior a comprehensive planning model must explain. A college graduate produced it while planning a hypothetical day's errands. We have collected a total of thirty protocols from five different subjects performing six different versions of such errand-planning tasks. The protocol shown in Figure 1 is representative of this set.

The subject began with the following problem description:

> You have just finished working out at the health club. It is 11:00 and you can plan the rest of your day as you like. However, you must pick up your car from the Maple Street parking garage by 5:30 and then head home. You'd also like to see a movie today, if possible. Show times at both movie theaters are 1:00, 3:00, and 5:00. Both movies are on your "must see" list, but go to whichever one most conveniently fits into your plan. Your other errands are as follows:
>
> > pick up medicine for your dog at the vet;
> > buy a fan belt for your refrigerator at the appliance store;
> > check out two of the three luxury appartments;
> > meet a friend for lunch at one of the restaurants;
> > buy a toy for your dog at the pet store;
> > pick up your watch at the watch repair;
> > special order a book at the bookstore;
> > buy fresh vegetables at the grocery;
> > buy a gardening magazine at the newsstand;
> > go to the florist to send flowers to a friend in the hospital

Note that the problem description specifies more errands than the subject could reasonably expect to accomplish in the time available. The subject's task was to formulate a realistic plan indicating which errands he would do, when he would do them, and how he would travel among them.

Figure 2 shows the hypothetical town in which the subject planned his errands. Each of the pictures on the map symbolizes a particular store or other destination. The subject was quite familiar with both the symbology and the layout of the town. In addition, the map was available during planning.

We have numbered small sections of the protocol in Figure 1 to facilitate the discussion. Also, for convenience, we refer to specific errands by the names of the associated stores or other destinations.

1 Let's go back down the errand list. Pick up medicine for the dog at veterinary supplies. That's definitely a primary, anything taking care of health. Fan belt for refrigerator. Definitely a primary because you need to keep the refrigerator. Checking out two out of three luxury apartments. It's got to be a secondary, another browser. Meet the friend at one of the restaurants for lunch. All right. Now, that's going to be able to be varied I hope. That's a primary though because it is an appointment, something you have to do. Buy a toy for the dog at the pet store. If you pass it, sure. If not, the dog can play with something else. Movie in one of the movie theaters. Better write that down, those movie times, 1, 3, or 5. Write that down on my sheet just to remember. And that's a primary because it's something I have to do. Pick up the watch at the watch repair. That's one of those borderline ones. Do you need your watch or not? Give it a primary. Special order a book at the bookstore.

2 We're having an awful lot of primaries in this one. It's going to be a busy day.

3 Fresh vegetables at the grocery. That's another primary. You need the food. Gardening magazine at the newsstand. Definitely secondary. All the many obligations of life.

4 Geez, can you believe all these primaries?

5 All right. We are now at the health club.

6 What is going to be the closest one?

7 The appliance store is a few blocks away. The medicine for the dog at the vet's office isn't too far away. Movie theaters—let's hold off on that for a little while. Pick up the watch. That's all the way across town. Special order a book at the bookstore.

8 Probably it would be best if we headed in a southeasterly direction. Start heading this way. I can see later on there are a million things I want to do in that part of town.

9 No we're not. We could end up with a movie just before we get the car. I had thought at first that I might head in a southeasterly direction because there's a grocery store, a watch repair, a movie theater all in that general area. Also a luxury apartment. However, near my parking lot also is a movie, which would make it convenient to get out of the movie and go to the car. But I think we can still end up that way.

10 All right. Apparently the closest one to the health club is going to be the vet's shop. So I might as well get that out of the way. It's a primary and it's the closest. We'll start . . .

[The experimenter mentions that he has overlooked the nearby restaurant and flower shop.]

11 Oh, how foolish of me. You're right. I can still do that and still head in the general direction.

12 But, then again, that puts a whole new light on things. We do have a bookstore. We do have . . . OK. Break up town into sections. We'll call them northwest and southeast. See how many primaries are in that section. Down here in the southeast section, we have the grocery store, the watch repair and the movie theater. In the northwest section we have the grocery store, the bookstore, the flower shop, the vet's shop, and the restaurant.

13 And since we are leaving at 11:00, we might be able to get those chores done so that some time when I'm in the area, hit that restaurant. Let's try for that. Get as many of those out of the way as possible. We really could have a nice day here.

14 OK. First choose number one. At 11:00 we leave the health club. Easily, no doubt about it, we can be right across the street in 5 minutes to the flower shop. Here we go. Flower shop at 11:05. Let's give oursleves 10 minutes to browse through some bouquets and different floral arrangements.

Figure 1. Thinking aloud protocol from the errand-planning task.

You know, you want to take care in sending the right type of flowers. That's something to deal with personal relationships.

15 At 11:10 we go north on Belmont Avenue to the Chestnut Street intersection with Belmont and on the northwest corner is a grocery.

16 Oh, real bad. Don't want to buy the groceries now because groceries rot. You're going to be taking them with you all day long. Going to have to put the groceries way towards the end.

17 And that could change it again. This is not one of my days. I have those every now and again. Let's go with our original plan. Head to the southeast corner.

18 Still leaving the flower shop at 11:10. And we are going to go to the vet's shop next for medicine for the dog. We'll be there at 11:15, be out by 11:20. The vet's shop.

19 Proceeding down Oak Street. I think it would be, let's give ourselves a little short-cut.

20 Maybe we'll knock off a secondary task too.

21 Proceed down Oak Street to Belmont. Belmont south to the card and gift shop, or rather, to the department store. Cut through the department store to Johnson Street to the newsstand. Pick up our gardening magazine at the newsstand.

22 We're heading this way. We're going to make a definite southeast arrow.

23 Third item will be the newsstand since we are heading in that direction. Often I like to do that. I know buying a gardening magazine is hardly a primary thing to do, but since I'm heading that way, it's only going to take a second. Let's do it. Get it out of the way. Sometimes you'll find at the end of the day you've done all your primary stuff, but you still have all those little nuisance secondary items that you wish you would have gotten done. So, 11:20 we left the vet's office. We should arrive 11:25 at the newsstand. 11:30 we've left the newsstand.

24 Now let's start over here. We're going to be in trouble a little bit because of that appliance store hanging way up north. So we could; appliance store is a primary. It's got to be done.

25 All right, let's do this. This could work out. Market Square, we leave the Market Square exit of the newsstand up to Washington, arrive at the pet store, buy a toy for the dog at the pet store. We're there at 11:35, out at 11:40. Pretty good. 11:40. Proceeding east just slightly, up north Dunbar Street to the appliance store, we arrive there at 11:45, and we leave there, fan belt, leave at 11:50.

26 We're looking good. We've knocked off a couple of secondaries that really we hadn't planned on, but because of the locations of some stores that are in the way that could be convenient.

27 Now it's 11:50, right near noontime.

28 And I think one of the next things to do, checking our primaries, what we have left to do, would be to go to the restaurant. And we can be at the restaurant at 5 minutes to noon. We're going to go down Dunbar Street, south on Dunbar Street to Washington east, to the restaurant which is located on the very eastern edge of the map. Meeting our friend there for lunch at 11:55, allowing a nice leisurely lunch. No, oh yeah. An hour, 12:55.

29 Now we've got to start being concerned about a few other things. We can pick up the car from the Maple Street garage by 5:30.

30 It's 12:55, done with lunch. Primary left to do, see a movie, pick up a watch, special order a book, and get fresh vegetables.

31 I would like to plan it so I can see the movie, pick up the vegetables, pick up my car, and then go home. Vegetables would rot.

32 So then with what we have left now to do is special order a book at a bookstore and pick up the watch at the watch repair..

33 So, I think we can make this a very nice trip. We're at the restaurant on Washington Avenue. Let's proceed west one block to Madison, south to Cedar Street. Cedar Street west right there at the intersection of Cedar and Madison is the watch repair. Pick up the watch at the watch repair. We should be at the watch repair by 1:05. Give us a good 10 minutes. 1:05 at the watch repair. Pick up a watch. We're out of there by 1:10.

34 Now I'm going to go just a slight back down Madison to one of the luxury apartments. I arrive at one of the luxury apartments at 1:15. I allow myself 15 minutes to browse. Two bathroom apartment. 1:30. Now I'm leaving that.

35 Next, I'm going to go west on Lakeshore, north on Dunbar, west on Cedar to the bookstore. And I will arrive at the bookstore at 1:35. Special order my book, 1:40.

36 From the bookstore I can go west on Cedar Street just a hair, down Kingsway, to a second luxury apartment. Find out what's happening at that luxury apartment. And I'm there at 1:45, allowing myself another 15 minutes there, 2:00 we're out.

37 We're taken care of checking out 2 out of 3 luxury apartments. We ordered our book.

38 Now we do have a problem. It's 2:00 and all we have left to do is see a movie and get the vegetables. And that's where I think I've blown this plan. I've got an hour left there before the movie.

39 So the best way to eliminate as much time as possible since we are now located at the Cedar Lakeshore apartments. That's not going to be . . .

40 If I go get the groceries now, it's not really going to be consistent with the plans throughout the day because I've been holding off on the groceries for rotting. If I take them to a movie. . . Vegetables don't really perish like ice cream.

41 We leave the luxury apartment on Lakeshore, proceed due east to Dunbar, and we're at the grocery store at 2:05. 2:05 at the grocery store. Hunt around for fresh vegetables, and we can give ourselves 20 minutes there. So we leave there at 2:25.

42 We leave there and we proceed up Dunbar, north to Cedar, Cedar west to the movie theater.

43 We probably arrive at the movie theater at 2:35. 2:35 we arrive at the movie theater which still gives us 25 minutes to kill before the next showing. But that's that. We're going to have to simply do it. I'm going to have to go with it for right now.

44 The plan seems to have worked well enough up until then. We made better time than we had thought. That happens in life sometimes. How did I get here so fast?

45 2:25. We catch the 3:00 showing. We leave there at 5:00. Proceed immediately down Johnson, up Belmont to the parking structure, and we're there at 5:05 at the parking structure. We had to pick it up by 5:30.

46 Got everything done, the only problem being having a little bit of time to kill in that one period.

47 You could have stretched out, to make things fair, you could have said, well, okay, I'll give myself an hour and 15 minutes at lunch, but as I did plan it, I did come up 30 minutes over. 25 minutes there. And that's a little bit of, when that happens you feel bad. You remember the old Ben Franklin saying about don't kill time because it's time that kills us. And I hate to have time to waste. I've got to have things work very nicely.

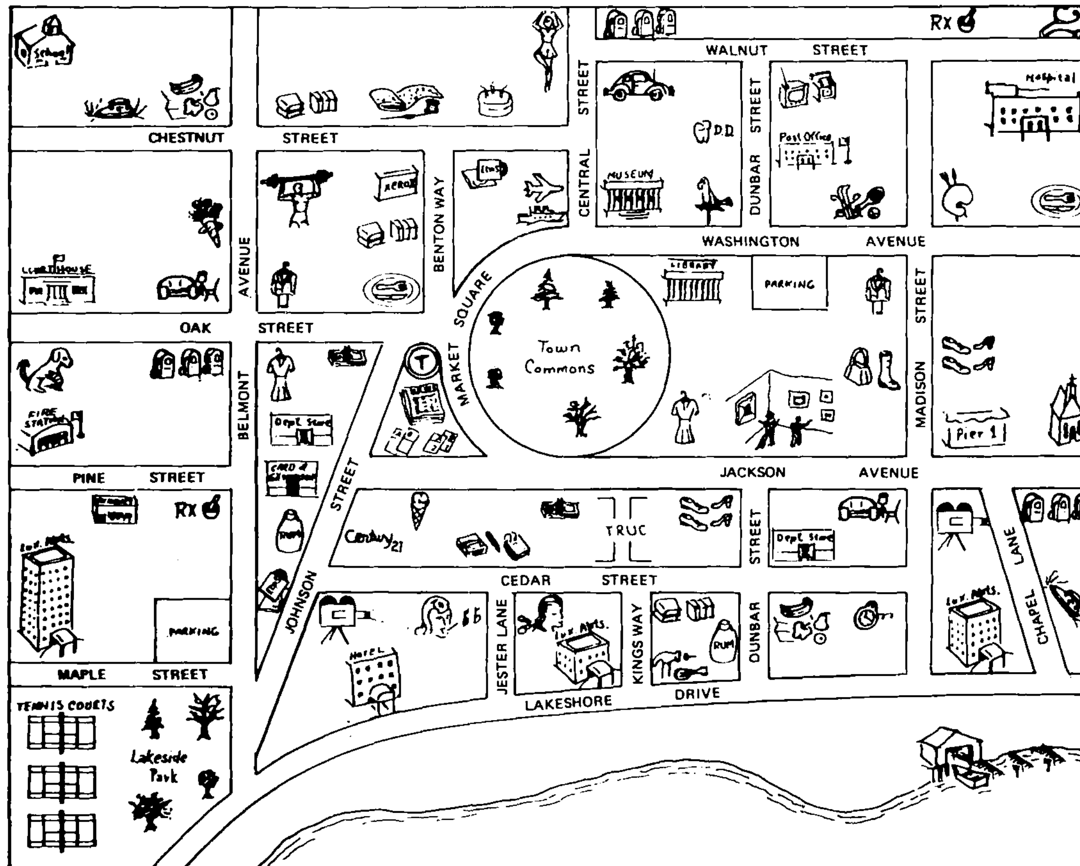

Figure 2. Town map for the errand-planning task.

In sections 1–4, the subject defines his goal and characterizes his task. Thus, in 1 and 3, he uses world knowledge to categorize the errands on his list as either primary errands, which he feels he must do, or secondary errands. In 2 and 4, he infers that, given the time constraints, his goal will be difficult to achieve.

In sections 5–7, the subject begins planning how to go about doing his errands. Notice that he begins planning at a fairly detailed level of abstraction. He has made only one kind of prior high-level decision—defining his goal. He has not considered what might be an efficient way to organize his plan. He has not made any effort to group his errands. He does not take his final location into consideration. Instead he immediately begins sequencing individual errands, working forward in time from his initial location. Thus, he ascertains his initial location, the health club, indicates that he wants to sequence the closest errand next, and begins locating the primary errands on his list, looking for the closest one.

In section 8, the subject changes his level of abstraction. In the course of looking for the closest errand to his current location, he apparently discovers a cluster of errands in the southeast corner of town. This observation leads him to make a decision at a "higher" or more abstract level than he had previously. Thus, he decides to treat the errands in the southeast corner as a cluster. He plans to go to the southeast corner and do those errands at about the same time.

In section 9, the subject modifies his high-level cluster. He discovers that one of the errands in the cluster, the movie, can also be done on the west side of town, near his final destination, the Maple Street parking structure. He changes back to the more detailed level of abstraction. Planning backward in time from his final location, he decides to end his day by going to the movie and then picking up his car. In so doing, he removes the movie from the high-level cluster.

In section 10, the subject begins to instantiate his high-level plan to go to the southeast corner at the lower, errand-sequencing level. Again, he is looking for the closest errand on his way, and he chooses the vet.

At that point, the experimenter interrupts to point out to the subject that he has overlooked several closer errands.

In sections 11 and 12, the subject incorporates the new information into his planning. His first reaction, in 11, is to continue working at the errand-sequencing level, simply considering the newly identified errands among those he might do next. However, additional observation at this level leads him to make a decision at the more abstract level. Again, he decides to treat a group of errands, those in the northwest corner of town, as a cluster. This leads him to revise his high-level plan to include two clusters of errands, the northwest cluster and the southeast cluster.

In section 13, the subject begins instantiating his new high-level plan. He notes the initial time, 11:00, and the presence of a restaurant, another errand in the northwest cluster. These observations lead him to formulate an intermediate level plan regarding how to sequence errands within the northwest cluster. He decides to sequence the errands in that cluster to permit him to arrive at the restaurant in time for lunch.

In sections 14–15, the subject works on instantiating his revised high-level plan at a very detailed level of abstraction. Here, he not only sequences individual errands (the florist and the grocery), he specifies the exact routes he will take among them. In addition, the subject mentally simulates execution of his plan in progress, estimating how long each errand should take and computing the "current" time at each stage of the plan.

In section 16, the subject's mental simulation suggests the inference that his groceries will perish if he picks them up early in the day. This leads him to revise his low-level plan, assigning the grocery a sequential position at the end of the plan.

In section 17, the subject decides to abandon his two cluster high-level plan in favor of his original high-level plan including only the southeast cluster. Presumably he decided that, without the grocery, there were not enough errands in the northwest cluster to occupy him until lunch.

In section 18, the subject begins instantiating his original high-level plan at a more detailed level. Again, he sequences individual errands (the florist and the vet) and specifies exact routes among them, mentally simulating execution of his plan as he formulates it.

In sections 19–23, the subject continues working at the lowest level of abstraction. He works on planning his route from the sequenced errands to the southeast corner, mentally simulating execution of his plan in the process. In so doing, he notices a "short-cut" through the card and gift shop and incorporates it into his plan, later replacing it with one through the department store. He then notices that taking the short-cut will put him very near the newsstand. Although the newsstand is a secondary errand, he decides to incorporate it in his plan because it is so convenient. Thus, a decision at the lowest level of abstraction leads him to make a decision at the next higher level. Note also that this decision implies addition of the newsstand to the subject's definition of his goal.

In sections 24–26, the subject continues working at a low level of abstraction. He notes that his high-level plan does not include any provision for the appliance store, a primary errand. He plans to go there directly, temporarily ignoring his high-level plan to go to the southeast corner. He also notices that another secondary errand, the pet store, is on the way to the appliance store and, because it is so convenient, incorporates that errand into his plan. Again, he plans at the level of sequencing errands and specifying routes and simulates execution of the plan as he goes along. Note that these decisions imply addition of the pet store to the subject's definition of the goal. (Note also that, while the short-cut planned in 19–23 was a short-cut to the southeast corner, it is a detour in the planned route to the appliance store.)

The remainder of the protocol (sections 27–45) documents the completion of the subject's plan. In the interests of brevity, we simply summarize this part of the protocol. Basically, the subject decides to incorporate the appliance store and the restaurant before finally arriving at the southeast corner. Then he plans all of the remaining errands, including all remaining secondary errands. Figure 9 below shows the subject's final plan.

This protocol illustrates a number of the points made above. First, the subject's plan develops incrementally at various points in the planning space we described. He plans actions at various points in the plan's temporal sequence, and he also plans at different levels of abstraction. Second, the subject appears to plan opportunistically, "jumping about" in the planning space to develop promising aspects of the plan in progress. For example, the planner does not plan strictly forward in time. Instead, he plans temporally-anchored sub-plans at arbitrary points on the time dimension and eventually concatenates the sub-plans. Similarly, the planner does not plan in a systematic top-down fashion across the different levels of abstraction. He frequently plans low-level sequences of errands or routes in the absence, and sometimes in violation, of a prescriptive high-level plan. Finally, decisions at a given point in the planning space appear to influence subsequent decisions at both later and earlier points in the temporal sequence and at both higher and lower levels of abstraction. The protocol contains examples of each of these kinds of influence.

The protocol illustrates another important component of the planning process—the ability to simulate execution of a plan mentally and to use the results of the simulation to guide subsequent planning. Mental simulation can answer a variety of questions for the subject: At what time will I arrive at (or leave) a particular destination? How long will I take to perform a certain action? What sequence of operations will I perform to satisfy a particular sub-goal? How long will it take to execute a plan or partial plan? What effects will my actions produce? What have I accomplished so far? The subject can use this information to evaluate and revise prior planning and to constrain subsequent planning.

The subject performs two kinds of mental simulation corresponding to *time-driven* and *event-driven* processes. Sometimes he simulates his plan by mentally stepping through a sequence of time units for each planned action (e.g., walking, carrying a package, performing an errand). With each successive step, he extrapolates the results of each planned action, updating his understanding of the "current state" accordingly. At other times, the subject performs "event-driven" simulation. In this case, he mentally moves directly from one planned situation to another, ignoring any actions in the intervening temporal interval. He then computes certain consequences arising from the transition.

More importantly, in the present context, the subject simulates execution of plans at different levels of abstraction. Thus, in sections 14–15, he simulates execution of a detailed plan. By stepping through his plan, the subject computes expected times for performing individual errands and traveling specific routes. In sections 24–26, the subject simulates execution of his high-level plan for performing errands in the northwest and then those in the southeast. Here, he performs event-driven simulation, inferring that if he executes his high-level plan, proceeding directly to the southeast corner of town, he will neglect a primary errand.

In the next section, we describe the proposed planning model in detail. The model postulates specific levels of abstraction and a structural organization for the planning space. In addition, it postulates decision-making mechanisms that permit theoretical interpretation of subjects' apparently chaotic progress through the planning space.

3. AN OPPORTUNISTIC MODEL OF PLANNING

Overview

The proposed model assumes that the planning process comprises the independent actions of many distinct cognitive *specialists* (akin to demons in Selfridge's (1959) Pandemonium model). Each specialist makes tentative *decisions* for incorporation into a tentative *plan*. Further, different specialists influence different aspects of the plan. For example, some specialists suggest high-level, abstract additions to the plan, while others suggest detailed sequences of specific actions.

All specialists record their decisions in a common data structure, called the *blackboard*. The blackboard enables the specialists to interact and communicate. Each specialist can retrieve prior decisions of interest from the blackboard, regardless of which specialists recorded them. The specialist combines these earlier decisions with its own decision-making heuristics to generate new decisions.

The model partitions the blackboard into several *planes* containing conceptually different categories of decisions. For example, one plane contains decisions about explicitly planned activities, while another contains decisions about data that might be useful in generating planned activities. The model further partitions each plane into several *levels* of abstraction. These partitionings serve two functions. First, they provide a conceptual taxonomy of the decisions made during planning. Second, they restrict the number of prior decisions each individual specialist must consider in generating its own decisions (see also Englemore & Nii, 1977). Thus, most specialists deal with information that occurs at only a few levels of particular planes of the blackboard.

The proposed model generalizes the theoretical architecture developed by Reddy and his associates (Cf. CMU Computer Science Research Group, 1977; Lesser, Fennell, Erman, & Reddy, 1975; Erman & Lesser, 1975; Lesser & Erman, 1977; Hayes-Roth & Lesser, 1977) for the Hearsay-II speech-understanding system. Others have since applied it to image understanding (Prager, Nagin, Kohler, Hanson, & Riseman, 1977), reading comprehension (Rumelhart, 1976), protein-crystallographic analysis (Nii & Feigenbaum, 1977), and inductive inference (Soloway & Riseman, 1977). The proposed model is, to our knowledge, the first attempt to adapt the Hearsay-II architecture to a "generation" problem. We describe it in detail below.

Specialists

As mentioned above, independent cognitive specialists generate decisions during the planning process. The model operationalizes specialists as condition-action rules.

The condition component describes circumstances under which the specialist can contribute to the plan. Ordinarily, the condition requires the planner to have made certain prior decisions. However, it may also require satisfaction of other, arbitrarily complex criteria. For example, one specialist's condition might require a prior decision to organize the plan by spatial clusters of errands and prior identification of useful clusters.

The action component defines the specialist's behavior. The action may include an arbitrary amount of computation, but always results in the generation of a new decision or modification of a prior decision. For example, one specialist might detect and identify spatial clusters of errands on the map. Another might generate an abstract organizational design for the plan as a whole.

Thus, specialists generalize the symbol-manipulation capabilities of production rules (Newell & Simon, 1972) to more complex, pattern-directed activity (see also: CMU Computer Science Research Group, 1977; Hayes-Roth, Waterman, & Lenat, 1978; Lenat, 1975).

The Blackboard

As discussed above, specialists record their decisions in a common data structure called the blackboard. The blackboard contains five conceptual planes: *plan, plan-abstractions, knowledge-base, executive* and *meta-plan*. We characterize each of these below.

We have already characterized the plan plane in our discussion of the thinking-aloud protocol. Decisions on this plane represent actions the planner intends to take in the world. For example, the planner might decide to travel in a circle around town, performing errands along the way, or to travel from the florist to the vet along Belmont Avenue and Oak Street. Both of these decisions describe explicit actions the planner intends to carry out.

Decisions on the plan-abstractions plane characterize desired attributes of potential plan decisions. Thus, these decisions indicate the kinds of actions the planner would like to take without specifying the actions themselves. For example, the planner might decide to go to the closest errand next. This decision characterizes a desired sequence of errands, but does not identify a particular sequence. Similarly, the planner might decide to organize the plan around spatial clusters of errands. Again, this decision characterizes a desired abstract plan, but does not instantiate it (i.e. does not specify particular spatial clusters).

The knowledge-base contains observations and computations regarding relationships in the world that might bear on the planning process. These computations are useful in suggesting plan-abstractions and instantiating them in the plan. For example, the planner might ascertain that the florist is the closest errand to the health club. That information would permit instantiation of a prior plan-abstraction decision to go to the closest errand next. As a second example, the planner might observe that several errand-sites cluster in close proximity on the map. That observation might suggest a subsequent plan-abstraction decision to organize the plan around several such spatially arrayed clusters.

Plan, plan-abstractions, and knowledge-base decisions determine features of the developing plan. Executive decisions, by contrast, determine features of the planning process. Thus, the executive plane contains decisions about the allocation of cognitive resources during the planning process. These decisions determine which aspect of the plan the planner will develop and which specialist the planner will bring to bear at each point in the process. For example, the planner might decide to determine which errands to include in the plan before working out the details of the plan. As a second example, the planner might decide to focus on working out routes among previously sequenced errands.

The meta-plan plane contains decisions about how to approach the planning problem. These decisions reflect the planner's understanding of the problem, the methods she or he intends to apply to it, and the criteria she or he will use to evaluate prospective solutions.

As mentioned above, the model further partitions each plane of the blackboard into several levels of abstraction. In the following sections, we describe the postulated levels of abstraction for each of the five planes (see Figure 3 on p. 288).

Levels of the Plan Plane. The plan plane has four levels of abstraction. Decisions at the four levels form a potential hierarchy, with decisions at each level specifying a more refined plan than those at the next higher level. Beginning at the most abstract level, *outcomes* indicate what the planner intends to accomplish by executing the finished plan. For the errand-planning task, outcomes indicate what errands the planner intends to accomplish by executing the plan. For example, the planner might decide to accomplish the desired errands at the florist, the vet, and the grocery store. At the next lower level, *designs* characterize the general behavioral approach by which the planner intends to achieve the outcomes. For the errand-planning task, designs characterize the general order in which the planner intends to perform errands. For example, the planner might decide to head toward the southeast cluster. Next, *procedures* specify specific sequences of gross actions. For the errand-planning task, proce-

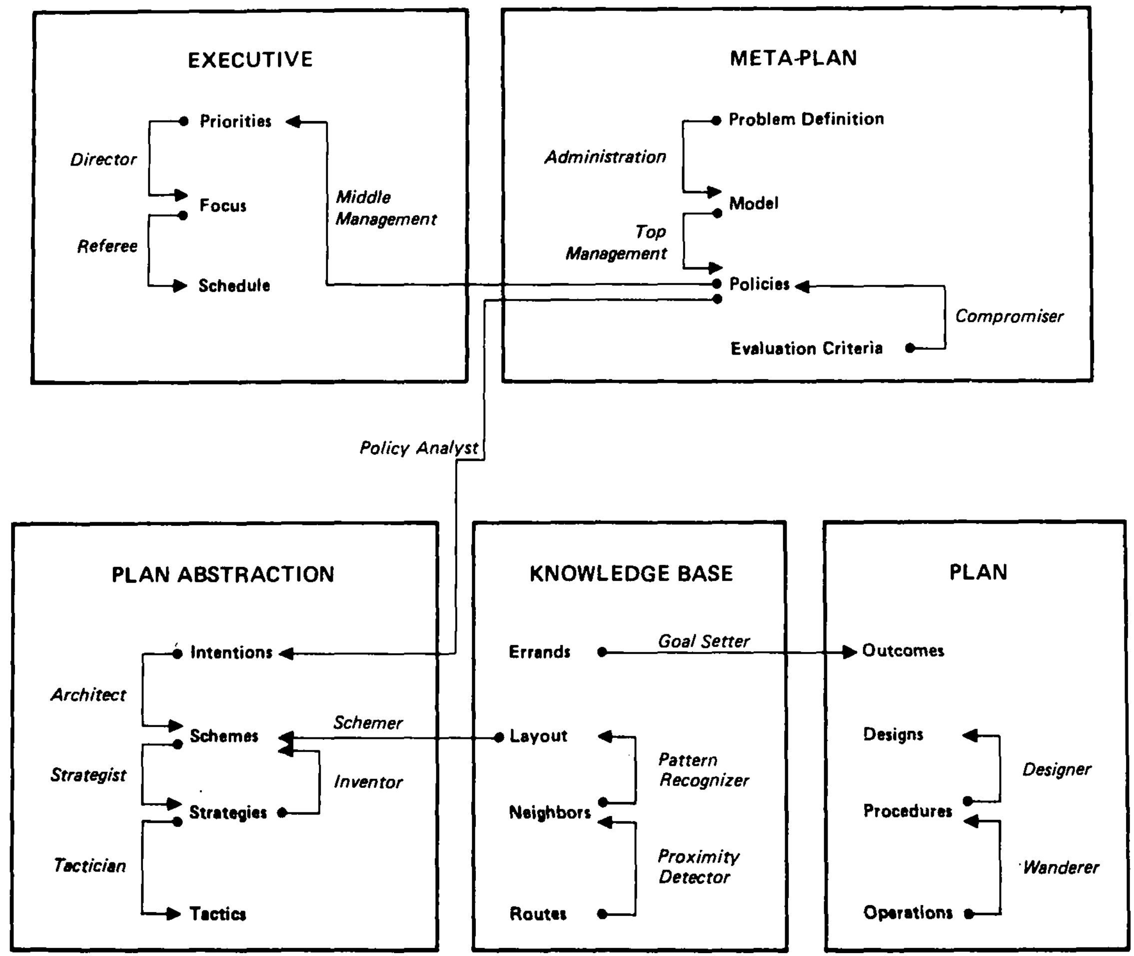

Figure 3. The planning blackboard and the actions of illustrative specialists.

dures specify sequences of errands. For example, the planner might decide to go to the vet after the florist. Finally, *operations* specify sequences of more minute actions. For the errand-planning task, operations specify the details of performing individual errands and the routes by which the planner will proceed from each errand to the next. For example, the planner might decide to travel from the vet to the florist via Belmont Avenue and Oak Street.[1]

Levels of the Plan-Abstractions Plane. The plan-abstractions plane contains four levels. Each level characterizes types of decisions suggested for incorporation into the corresponding level of the plan plane. For example, the planner might indicate an *intention* to establish all of the "critical" errands as the outcome of the plan. At the next lower level, the planner might generate a *scheme* that suggests generating a design featuring spatial clusters of errands. At the next level, the planner might develop a *strategy* to go to the closest errand next, characterizing a desirable procedure level decision. Finally, the planner might adopt a *tactic* to search for a short-cut between one errand and the next, characterizing a desirable operation level decision.

Levels of the Knowledge Base Plane. The knowledge base also has four levels of abstraction. Each level contains observations and computations useful in suggesting decisions at the corresponding level of the plan-abstractions plane or instantiating them at the corresponding level of the plan plane. Because the levels of the knowledge base contain problem-specific information, we have given them problem-specific names. At the *errand* level, for example, the planner might determine the relative importance of each desired errand. At the *layout* level, the planner might observe that several errands form a convenient spatial cluster. At the *neighbor* level, the planner might observe that two planned errands are near one another. At the *route* level, the planner might detect a short-cut.

Levels of the Executive Plane. The executive plane has three levels of abstraction. Decisions made at the three levels on this plane form a hierarchy, with decisions at each level potentially refining ones at the level above. Beginning at the top, *priority* decisions establish principles for allocating cognitive resources during the entire planning process. These decisions generally indicate preferences for allocating processing activity to certain areas of the planning blackboard before others. For example, by approaching the errand-planning task as a resource-limited scheduling problem, the planner might decide to determine which errands to do before working out the details of the plan. At the next lower

[1]Obviously, partitioning plan decisions into four discrete categories is arbitrary and probably over-simplified. However, we find these categories intuitively appealing and they provide a convenient terminology for discussion. In addition, Hayes-Roth and Thorndyke (1979) have shown that theoretically naive subjects group statements drawn from planning protocols in exactly these four categories.

level, *focus* decisions indicate what kind of decision to make at a specific point in time. For example, the planner might decide to focus attention on generating an operation-level refinement of a previously generated procedure. Finally, *schedule* decisions resolve any remaining conflicts among competing specialists. If, given current priorities and focus decisions, more than one specialist can make a contribution to the plan, the planner must make schedule decisions to decide among them. Schedule decisions select specialists on the basis of relative efficiency, reliability, etc. (Hayes-Roth & Lesser, 1977).

Levels of the Meta-Plan Plane. The meta-plan plane has four levels: *problem definition, problem-solving model, policies,* and *evaluation criteria.* Unlike the levels on the other four planes, these levels do not produce a neat hierarchy of decisions. However, they emphasize different aspects of the subject's approach to the planning problem and relate in systematic ways to the other planes of the blackboard.

Beginning at the top, problem definition decisions characterize the planner's own formulation of the task. These include descriptions of goals, available resources, possible actions, and constraints. For the errand-planning task, the problem definition would reflect the subject's understanding of the list of errands, contextual information, and associated instructions.

The chosen problem-solving model indicates how the planner intends to represent the problem and generate potential solutions. For example, the planner might view the errand-planning task as an instance of the familiar "traveling salesman" problem (Christophides, 1975) and approach the problem accordingly. Problem-solving models can also consist of general problem-solving strategies, such as "divide and conquer," "define and successively refine" (Aho, Hopcroft, & Ullman, 1974), etc. The planner presumably chooses a particular problem-solving model from known alternatives in response to specific problem characteristics. The problem-solving model, in turn, bears directly on subsequent executive decisions. For example, adoption of the traveling salesman model should lead to basically "bottom-up" executive decisions. That is, the planner should focus attention on the procedures and operations levels of the plan plane and on corresponding levels of the plan-abstractions and knowledge-base planes.

The planner's policies specify global constraints and desirable features for the developing plan. For example, the planner might decide that the plan must be efficient or that it should minimize certain risks. Some policy decisions derive implicitly from particular problem-solving models. For example, the traveling salesman model naturally implies a route-efficiency policy. Other policies are model-independent. In either case, policy decisions bear directly on subsequent plan-abstractions decisions. Particular policy decisions make particular plan-abstractions more or less desirable. For example, the route efficiency policy favors a strategy to go to the closest errand next. By contrast, it inhibits an intention to achieve only the most important errands.

Finally, solution-evaluation criteria specify how the planner intends to evaluate prospective plans. For example, the planner might decide to speculate on what could go wrong during execution and insure that the plan is robust over those contingencies. Again, some of these decisions derive implicitly from particular problem-solving models, while others are independent. Obviously, the planner brings these criteria to bear on the developing plan and uses them to determine which plan decisions to preserve and which to change.

Control of the Planning Process

Under the control of the executive, the planning process proceeds through a series of "cycles" during which various specialists execute their actions. At the beginning of each cycle, some number of specialists have been invoked—that is, their conditions have been satisfied. The executive selects one of the invoked specialists to execute its action—that is, to generate a new decision and record it on the blackboard. The new decision invokes additional specialists and the next cycle begins. This process ordinarily continues until: (a) the planner has integrated mutually consistent decisions into a complete plan; and (b) the planner has decided that the existing plan satisfies important evaluation criteria. Under certain circumstances, the process might also terminate in failure.

4. ANALYSIS OF THE PLANNING PROTOCOL UNDER THE OPPORTUNISTIC MODEL

In this section, we use the proposed model to "parse" sections 1–10 of the protocol. We intend this exercise to demonstrate the descriptive power of the model. Of course, the psychological validity of the model rests on more than this informal sufficiency test. In subsequent sections of the paper, we discuss a more formal sufficiency test based on a computer simulation and summarize several empirical tests of the model's assumptions.

Figures 4–8 show blackboard representations of sections 1-10 of the protocol as individual decisions. They also show how individual specialists respond to the presence of particular decisions on the blackboard by generating other decisions and recording them at appropriate locations on the blackboard. Each arrow represents the invocation and execution of a specialist. Thus, an arrow from one decision to another indicates that the former decision invoked a specialist that recorded the latter decision. In order to clarify the flow of activity, we have numbered decisions in Figures 4–8 according to their presumed order of occurrence. Note, however, that arrows need not connect consecutively numbered decisions. Occasionally, an early decision invokes a specialist that is not scheduled until after one or more other specialists have been scheduled and added their decisions to the blackboard.

We have omitted only one kind of decision from these illustrations—

schedule decisions. As discussed above, at each point in the sequence of recorded decisions, a schedule decision selects one of the currently invoked specialists to execute its action. We have omitted these decisions from Figures 4–8 for simplicity. However, it is appropriate to assume that a schedule decision selected each of the indicated specialist actions (noted by arrows).

Figure 4 shows the blackboard representation of sections 1–4 of the protocol. In sections 1 and 3, the subject works through the list of errands, assigning binary importance values (primary versus secondary) to each one. In sections 2 and 4, the subject remarks that the large number of primary errands implies that he will have a busy day. According to our assumptions, a specialist calculates importance values for individual errands and records these at the errands level of the knowledge base. However, we assume that a considerable amount of activity, unstated in the protocol, preceded and motivated this action. Figure 4 shows the blackboard representation of this implicit activity.

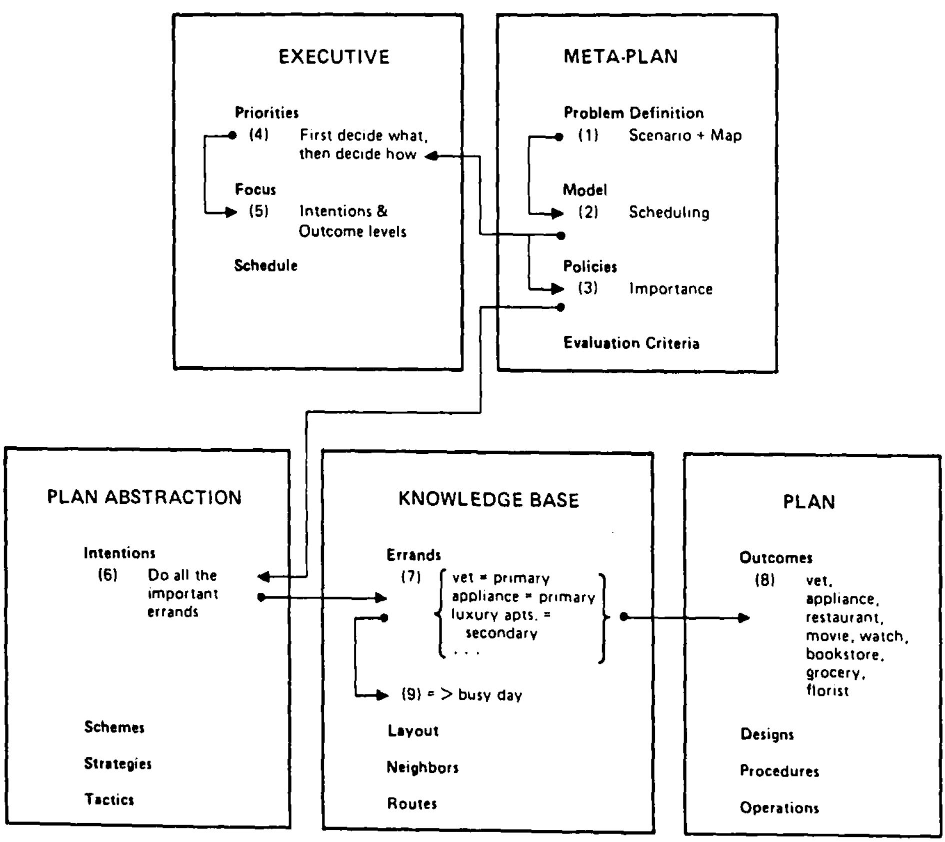

Figure 4. Blackboard representation of sections 1–4 of the protocol.

The subject begins the task with a problem definition (1), including the scenario and map provided by the experimenter. The protocol suggests that the subject categorizes the problem as a resource-limited scheduling problem (2). In other words, the subject apparently views the task as one in which he cannot do all of the things he wants to do and, therefore, must decide which things to do and then how to do them. The appearance of this problem-solving model on the blackboard presumably invokes two other specialists, One generates and records a useful policy (3), emphasizing the importance of individual errands. The other generates and records an appropriate set of priorities (4). The priorities, in turn, motivate a decision to focus on the intentions and outcomes levels of the plan-abstraction and plan planes (5). Given this focus and the errand-importance policy, a specialist records an intention to do all the important errands (6). This intention presumably invokes the specialist described above that calculates the errand-importance values actually stated in the protocol (7). This activity implies another unstated decision, that the intended outcomes should include the designated primary errands (8). Finally, the statements in sections 2 and 4 of the protocol imply that the errand-importance calculations invoke another specialist that infers: "It's going to be a busy day" (9).

Figure 5 shows the blackboard representation of section 5 of the protocol. In section 5, the subject states: "All right. We are now at the health club." This statement conveys a procedure-level specification of the initial location (13). Figure 5 shows the implicit sequence of activity that produced this statement, given the prior state of the blackboard shown in Figure 4. First, having decided what to do (8), the subject proceeds to his second priority, deciding how to do it. Accordingly, he changes focus to the lower levels of the blackboard (10). Given this focus, a strategy-generating specialist records its decision to plan forward from the initial location (11). This decision motivates another specialist to identify the initial location (12) which, in turn, motivates a specialist to record the initial location at the procedure level of the blackboard (13).

Figure 6 shows the blackboard representation of sections 6–8 of the protocol. In section 6, the subject asks, "What is going to be the closest one?" This question indicates a strategic decision to plan to perform the closest errand next in the procedure sequence (14). The appearance of this strategy on the blackboard invokes a specialist that evaluates the relative proximities of other primary errands to the initial location, the health club (15). Section 7 of the protocol describes these evaluations.

Section 8 of the protocol reflects a discontinuity in the planning process. The preceding statements aim toward recording the second errand in the procedural sequence. Instead, however, the subject states in section 8: "Probably it would be best if we headed in a southeasterly direction. Start heading this way. I can see later on there are a million things I want to do in that part of town." This statement expresses a higher-level design, recorded on the blackboard as a decision to perform the errands in the southeast cluster, performing whatever other

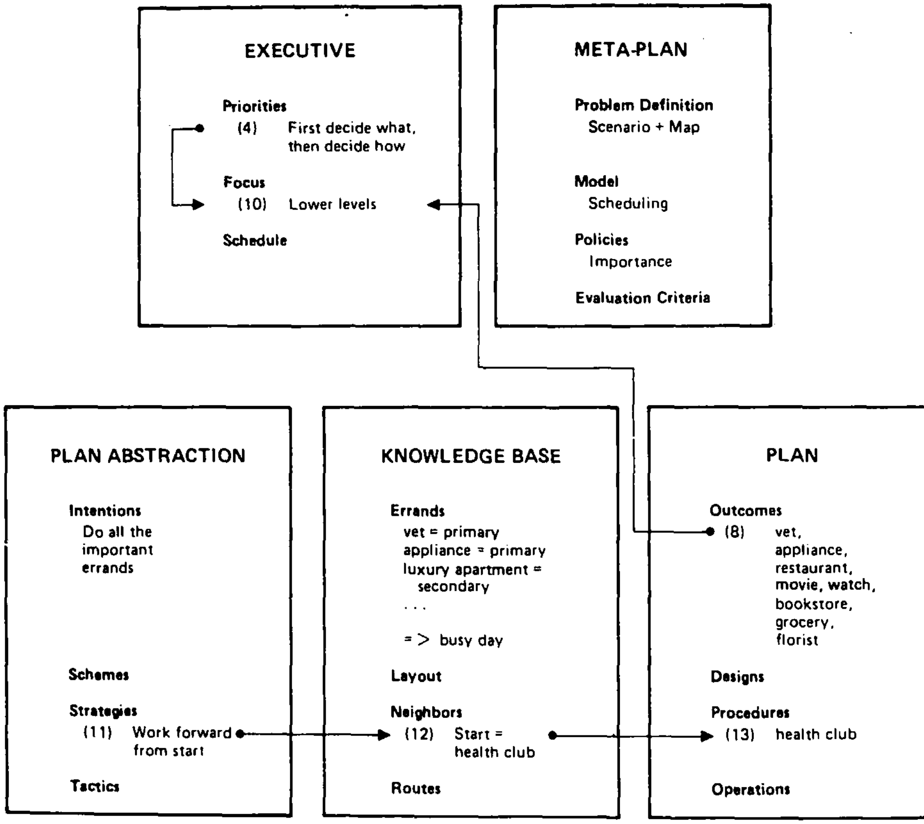

Figure 5. Blackboard representation of section 5 of the protocol.

errands occur along the route from the initial location to the southeast cluster (18).

Let us consider how the subject might have arrived at this design. The subject's immediately-preceding overt activity, evaluation of proximities, requires him to locate each errand in the list. In doing so, the subject locates (at least) three consecutive errands, the movie, the watch repair, and the bookstore, in the southeast corner of town. Apparently, this sequence of visual observations invokes a specialist that identifies clusters of errands and records the identity of the detected cluster at the layout level of the knowledge base (16). The appearance of the cluster on the blackboard invokes another specialist that generates schemes. It suggests exploiting the spatial cluster of errands by organizing a design around it (17). Another specialist responds to the new scheme and the identified cluster by recording the appropriate design on the blackboard (18).

Figure 7 shows the blackboard representation of section 9 of the protocol. In section 9, the subject indicates a procedure decision to sequence the movie

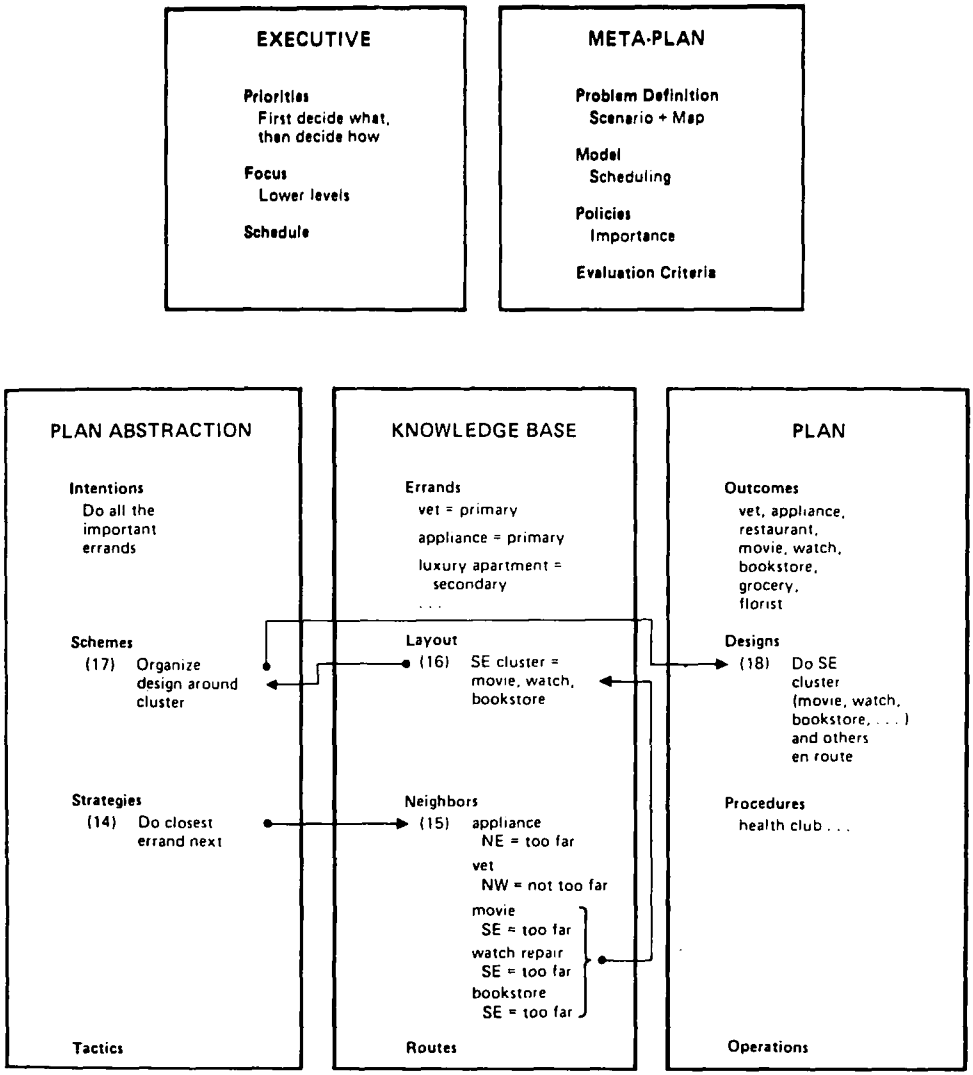

Figure 6. Blackboard representation of sections 6—8 of the protocol.

right before picking up his car at the end of the day (21). He tells us explicitly that, in so doing, he is removing the movie from the previously-defined southeast cluster (22). He also tells us why he has made this decision: because it would be "convenient to get out of the movie and go to the car" (20).

Figure 7 models these decisions, beginning with the subject's prior definition of the southeast cluster (16). Presumably, attention to one of the errands in the cluster, the movie, invokes a specialist that notices another movie on the west

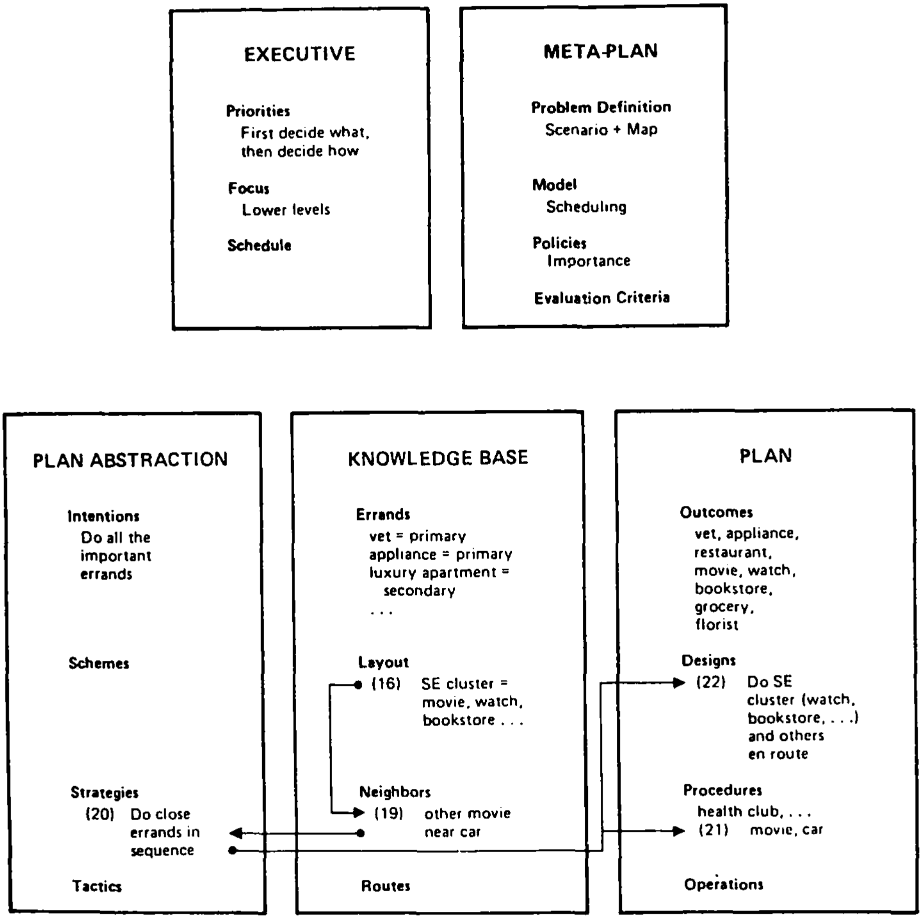

Figure 7. Blackboard representation of section 9 of the protocol.

side of town close to the parking structure (19). The proximity of these two errands invokes a specialist that suggests a more general strategy to perform two proximate errands in sequence (20). This new strategy invokes another specialist that records the suggested sequence, movie–car, at the procedure level of the plan plane (21) and amends the prior design accordingly (22).

Figure 8 shows the blackboard representation of section 10 of the protocol. In section 10, the subject decides to go to the vet after the health club because it is the closest primary errand. Thus, section 10 conveys a procedure-level decision (26) and the strategy that motivated it (24). We assume that the presence of a modified design on the blackboard motivates a narrowing of the focus to aim at instantiating the design at the procedure level (23). In accordance with this focus, the design also invokes a specialist that generates a stragegy to do the closest errand in the right direction (24). This strategy invokes a specialist that evaluates the proximities of individual errands at the neighbors level of the knowledge base (25). Finally, the observation that the vet is the closest errand to the initial location, the health club, invokes a specialist that records the vet as the next errand in the procedural sequence (26).

We can analyze the remainder of the protocol in much the same fashion. However, we conclude the analysis at this stage for brevity.

The analysis reinforces the main points made in section 2. The subject plans at different points in the planning space along both temporal and abstractness dimensions. In particular, the subject appears to make decisions at each of

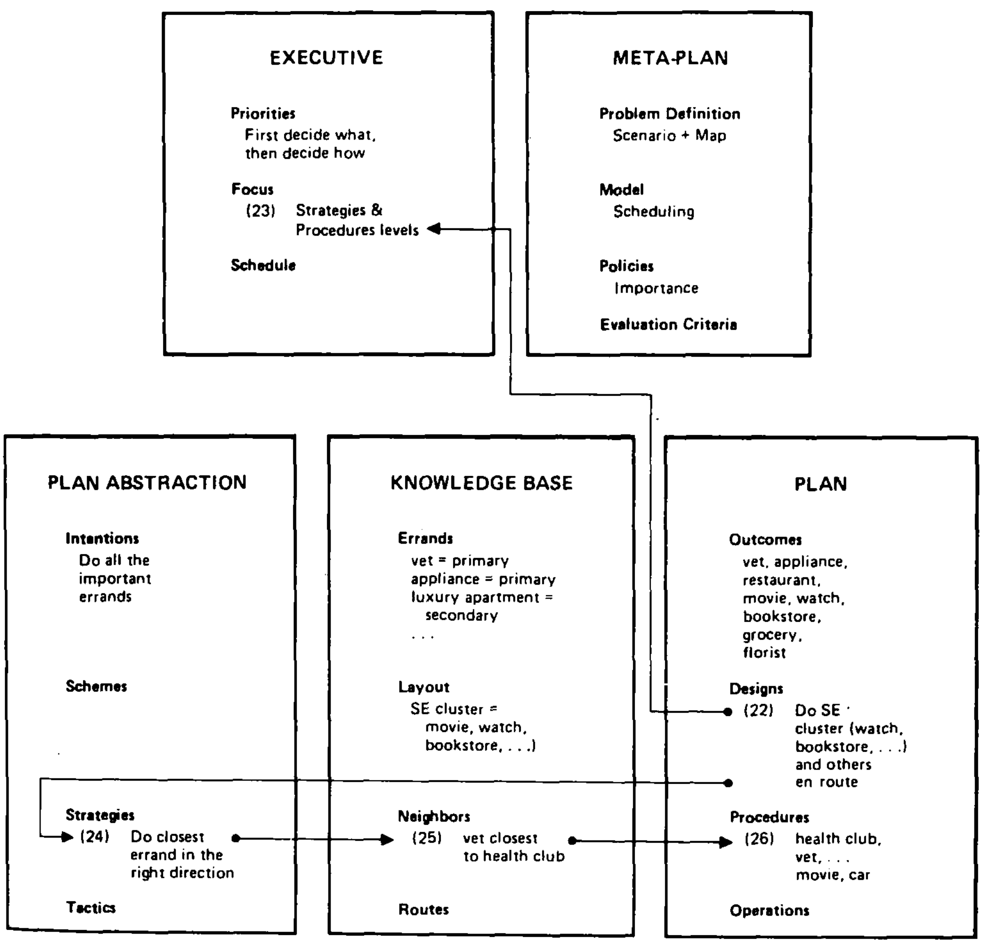

Figure 8. Blackboard representation of section 10 of the protocol.

the postulated levels on all five planes of the blackboard. Further, the subject makes decisions opportunistically. Rather than working systematically through the levels along either of the two dimensions, he enters the planning space at various points and moves about freely within it. The subject's observations and computations on the available data (the map and the scenario) exert a powerful influence on the point in the planning space at which he makes each successive decision. This produces a strong "bottom-up" component to the planning process. However, prior decisions at both higher and lower levels influence the subject's decisions, as assumed by the model.

5. COMPUTER SIMULATION

We have implemented a computer simulation of the planning model described above. The simulation is written in INTERLISP. It contains an internal representation of the map shown in Figure 1, a blackboard structure to organize planning decisions, and about forty specialists. (See Hayes-Roth, Hayes-Roth, Rosenschein, & Cammarata (1979) for a detailed description of the simulation.)

We designed the specialists to model some of the knowledge in our subject's protocol. Following the reasoning used in section 4, we postulated condition–action rules for producing many of the decisions in the protocol, as well as rules for producing the necessary intermediate decisions. The specialists generalize these rules. For example, in section 8 of the protocol, the subject notices that certain errands appear in close proximity in the southeast part of town. Based on this section of the protocol, we designed a specialist whose condition requires that at least three errands have been located on the map and that they appear in the same region (northwest, northeast, southwest, or southeast). Its action is to identify as a cluster any set of errands that satisfies its condition. Thus, this specialist can identify not only the particular cluster the subject noticed, but other clusters as well.

The simulation includes specialists for most of the condition-action rules inferred from the protocol. However, we did not attempt to capture all of the subject's idiosyncracies. For example, although the subject used several slightly different navigation rules, the simulation has only one. Thus, the simulation represents an approximate model of the subject's knowledge.

We can evaluate two aspects of the simulation's performance: the plans it produces and the process by which it produces them. We discuss each of these below.

Figure 9 shows the plan the subject produced for the problem discussed above. Figure 10 shows the plan produced by the simulation. The two plans are quite similar. Both plans include all primary errands and at least some of the secondary errands. While the subject included all secondary errands, the simulation included only one very convenient secondary errand. While the simulation and the subject planned different routes, both routes are fairly efficient, though clearly sub-optimal. Both the simulation and the subject planned to arrive at

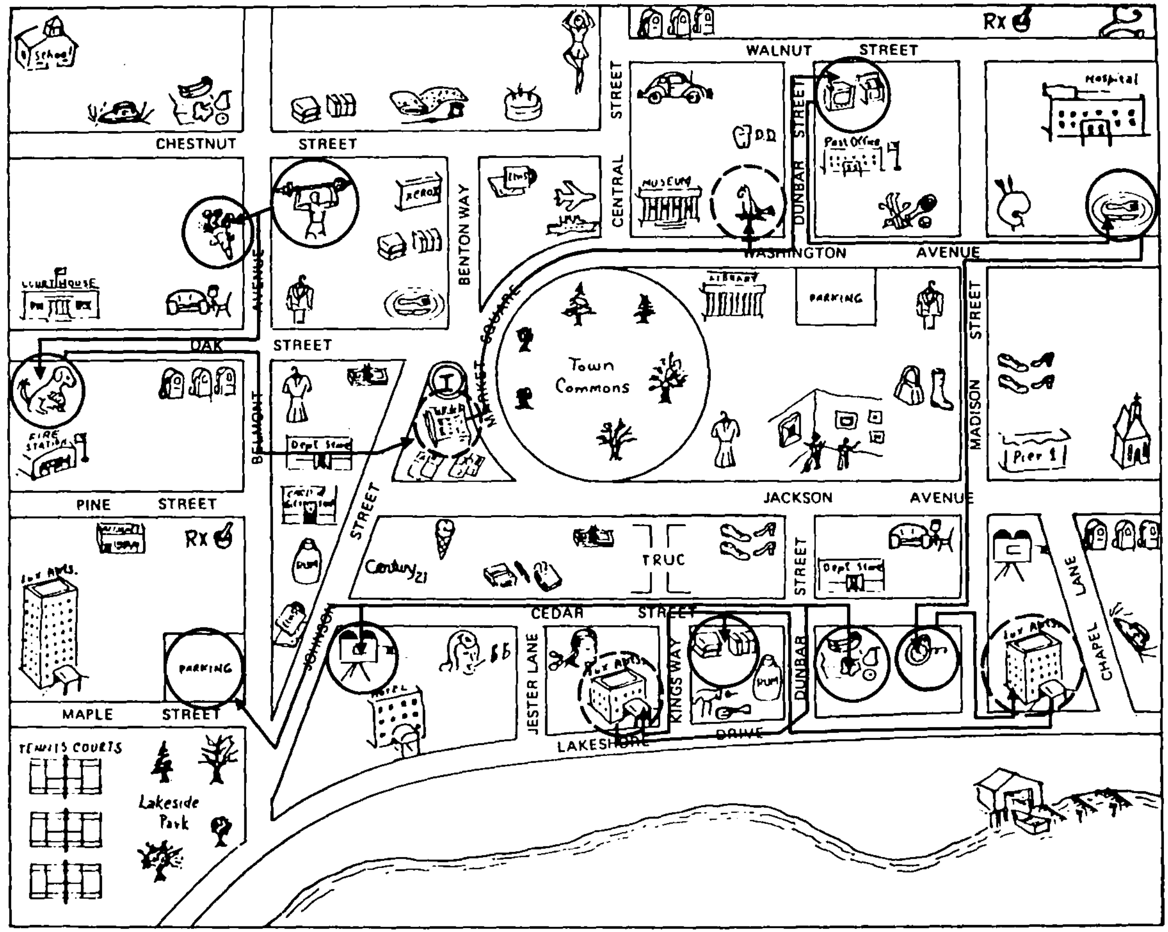

Figure 9. Plan produced by the human subject.

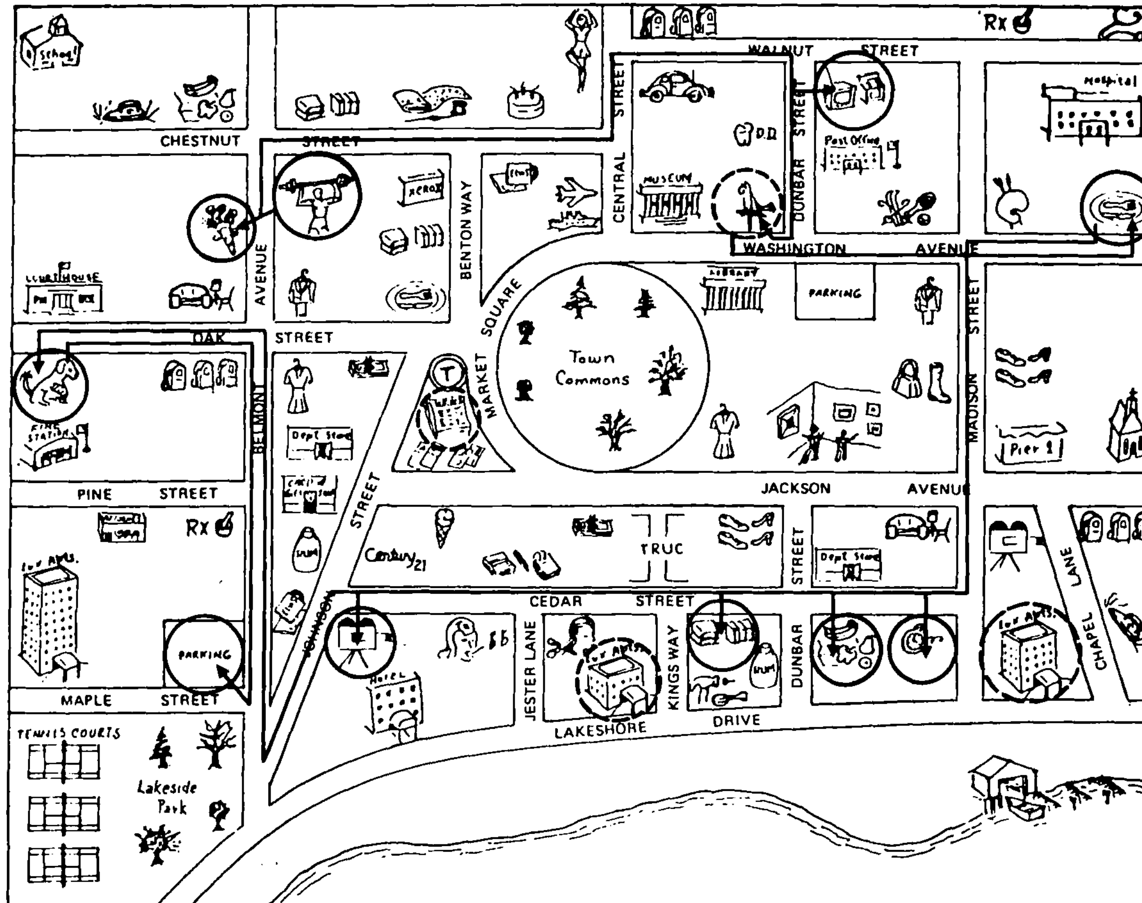

Figure 10. Plan produced by the computer simulation.

time-constrained destinations (e.g. the restaurant and the movie) at reasonable times. The major difference between the two plans lies in their relative "realism." The subject's plan is quite unrealistic—one could not execute the complete plan in the time available for doing so. The simulation's plan is somewhat more realistic, primarily because it omits many of the secondary errands.

1–9	[Omitted]
10	Work forward for the starting location.
*11	The starting location is the Health Club which is on Belmont Avenue.
12	Begin at the Health Club.
*13	Go to the closest errand next.
*14	The vet is on Oak Street.
*15	The appliance store is on Walnut Street.
16	The restaurant is on Washington Avenue.
*17	The watch repair is on Cedar Street.
*18	The bookstore is on Cedar Street.
19	The grocery is on Cedar Street.
*20	The watch repair, the bookstore, and the grocery are all in a cluster in the southeast part of town.
21	Organize the plan around spatial clusters.
*22	Proceed from the Health Club toward the southeast cluster.

Figure 11. Excerpt from the simulation's protocol.

Like the subject, the simulation produced a planning protocol—the series of decisions underlying the final plan. Figure 11 shows an excerpt from this protocol—decisions 10–22. (The protocol actually produced by the simulation is a series of decision "nodes" in list notation. Figure 11 translates the protocol into standard English for clarity.)

Decisions 10–22 of the simulation's protocol correspond quite closely to sections 5–8 of the subject's protocol. Recall that, in those sections, the subject identified the health club as the starting location, decided to schedule the closest errand to the health club next, located several errands on the map in his search for the closest errand, detected a cluster of errands in the southeast part of town and, accordingly, decided to head in that general direction.

Now consider the simulation's protocol. Decision 10 establishes a strategy to work forward from the starting location. Decisions 11 and 12 identify the starting location as the health club and establish it as the first errand in the procedure. Decision 13 establishes a strategy to go to the closest errand next. Decisions 14–19 locate individual errands on the map in a search for the closest errand to the health club. Decision 20 detects the cluster of errands in the southeast part of town. Decision 21 suggests organizing an overall design for the plan around the spatial cluster and decision 22 does so.

In addition to performing essentially the same functions the subject performed, the simulation made many of the decisions explicitly declared in the subject's protocol. Common decisions are preceded by "*" in Fig. 11. Of course, the remainder of the simulation's protocol does not always mirror the subject's protocol as closely as the section in Figure 11. Two factors produce the divergences.

First, as mentioned above, the simulation does not contain the entire set of specialists used by the subject. Thus, the simulation occasionally uses a specialist that is slightly different from the one the subject uses. This produces differences in both the protocols and the resulting plans. For example, in section 18 of the protocol, the subject decides to go from the health club to the vet on the way to the southeast part of town. At the same point in its protocol, the simulation decided to go from the health club to the appliance store. Both were trying to find the closest errand along the way, but they used slightly different specialists, and, as a consequence, chose different errands.

Second, the simulation's executive is incomplete. On some cycles, two or more invoked specialists are equally attractive and the simulation chooses randomly among them. Frequently, this random choice fails to select the specialist the subject used at that point. In such cases, the protocols also diverge.

In our opinion, it would be unproductive to model the subject's performance at a level of detail sufficient to counteract the effects of these two factors. Therefore, we look for the same general features in the two protocols, rather than exact replication. Such commonalities are readily apparent. Both the simulation and the subject made decisions at various levels of abstraction on each of the five planes of the planning blackboard. Both exhibited many coherent decision sequences in which each decision appeared to build on its predecessors. However, both also frequently "jumped about" the planning blackboard, rather than working systematically along any particular dimension. In particular, both the simulation and the subject occasionally redirected or dramatically changed their own activity in response to fortuitous observations or computations on the available data (e.g., the map).

6. THEORETICAL COMPLEXITY

The opportunistic model seems, at first glance, fairly complex. It postulates five different conceptual "planes" of decisions and several levels of abstraction within each of those planes. It postulates numerous planning specialists whose simultaneous efforts to participate in the planning process require the supervision of a fairly sophisticated executive. Although a number of complex models have proved fruitful in the last few years (Cf., Anderson, 1976; Anderson & Bower, 1973; Rumelhart, Lindsay, & Norman, 1972; Winograd, 1972), most of us still adhere to the law of parsimony, preferring simpler models to complex models.

In fact, the proposed model is computationally quite simple. It postulates a uniform decision mechanism, the specialist, to perform all of the varied decision-making functions planners perform. By modeling executive decisions with the same mechanism, it can account for a wide range of distinct planning styles without additional assumptions. (See Hayes-Roth (1979) for a discussion of executive "flexibility.")

Most of the apparent complexity in the model derives from the details of the blackboard structure. However, the blackboard partitions provide another important computational efficiency. Each specialist gets invoked whenever a new decision on the blackboard satisfies its condition. If the blackboard were simply an unstructured collection of decisions, each specialist would have to examine every new decision to determine whether the decision satisfied its condition. This would require an enormous amount of computation, much of it unnecessary. The blackboard partitions reduce the amount of computation required by permitting each specialist to restrict its "attention" to only those new decisions that occur at particular levels.

The blackboard structure also permits the model to capture an important psychological feature—interruptability. People have the power to interrupt their own cognitive processing at arbitary points. After performing some more or less related processing, they may or may not continue the interrupted task. This interruptability appears throughout our protocols and we believe it is a salient feature of cognitive processing in general. In the proposed model, interruption can occur only between individual decisions. Thus, the blackboard structure embodies our view of possible loci for interruption.

7. COMPARISON WITH SUCCESSIVE REFINEMENT MODELS

As discussed in the introduction to this paper, our view of planning differs somewhat from earlier views of planning as a process of successive refinement. This section explores several differences between the two views and attempts to resolve the differences.

Top-Down versus Multi-Directional Processing

While earlier work has assumed that planning is a top-down process, the proposed model treats planning as a multi-directional process. The diverse observations people make while planning often guide subsequent planning. Some of these observations arise from planning at an abstract level and guide subsequent planning at a more detailed level. The errand-planning protocol illustrates this kind of top-down processing in section 10, where the subject begins to instantiate a previously planned design at the lower procedure level. However, observations also arise from planning at a low level and guide subsequent planning at a more abstract level. The protocol illustrates this kind of bottom-up processing in section 8 where the subject formulates a design based on observations related to previous decisions at the lower procedure level. Many other examples of both top-down and bottom-up processing appear throughout this protocol and the others we have collected.

The sample protocol confirms the more general assumption of multi-directionality in another way. If the subject were operating in a top-down fashion, he would begin planning at the highest (most abstract) level of the planning space. He could plan at a lower level only if he had already planned that particular subtask at all higher levels. The errand-planning protocol disconfirms this presumption repeatedly. The subject begins forming his actual plan at a relatively low level, the procedure level. Thus, he plans at this level in the absence of any corresponding high-level plans. Similar instances of planning a subtask at a low level without having previously planned it at higher levels appear throughout this protocol and our others. These findings follow directly from the multi-directional assumption. (See Hayes-Roth and Thorndyke, 1979, for additional evidence on this point.)

Complete versus Incremental Planning

A second difference between the earlier view of planning and the proposed model concerns the relative completeness attributed to abstract plans. The earlier work assumes that, while initial plans may be abstract, they will be complete and fully integrated. Under a breadth-first processing assumption, this requires that complete plans at each level must precede any planning at the next lower level. Under a depth-first processing assumption, it requires only that the highest-level plan must be complete before planning activity can proceed at lower levels. Under either assumption, the earlier view presupposes that complete plans will eventually exist at all levels of abstraction.

By contrast, we assume that planning is incremental and, therefore, will rarely produce complete plans in the systematic fashion described above. We assume that people make tentative decisions without the requirement that each one fit into a current, completely integrated plan. As the planner relates each new decision to some subset of his previous decisions, the plan grows by incremental accretion. Further, the developing plan need not grow as a coherent integrated plan. Alternative subplans can develop independently either within or between levels of abstraction. The planner can incorporate these sub-plans into the final plan as she or he wishes.

The sample protocol provides evidence for these assumptions. For example, in section 9, having established only his initial location at the procedure level, the subject plans a sequence of two errands with which to conclude. In the following several sections of the protocol, he intermittently plans alternative designs (none of which covers the planned concluding sequence) and initial sequences of errands (none of which he concatenates with the planned concluding sequence). Similar partial plans appear throughout the protocol as well as in the other protocols we collected. These findings confirm our assumption that specialists record tentative decisions in various locations on the blackboard in response to relevant prior decisions.

Hierarchical versus Heterarchical Plan Structures

Earlier conceptions of plans as hierarchical structures responded to the appealing simplicity of hierarchically structured programs and the successive refinement method. None of our observations denies the putative merits of these hierarchical approaches. Of course, one can always interpret a sequence of actions as a hierarchy with some number of levels. Therefore, one must perform some more informative analysis to contrast hypothesized hierarchical plans with more complex plan structures. More importantly, a satisfactory theory of planning must describe all decisions made during the planning process as well as those that appear in completed plans.

Our efforts to model the planning process sugggest that people make many decisions that do not fit a simple hierarchical structure. Under the proposed model, one might attempt to construe the final set of decisions on the plan plane as a hierarchical structure, but our protocols do not provide strong evidence for such a structure. For example, the design maintained throughout most of the sample protocol dictates that errands on the way to the southeast cluster should be performed first followed by those errands within the cluster itself. However, much of the subject's planning at lower levels concerns errands not covered by this design (e.g., the newsstand, the pet store, the appliance store, and the restaurant).

The assumption of hierarchical plan structure becomes more tenuous if we consider the many other kinds of decisions our subject made while planning. We have observed four categories of decisions that do not describe what the subject actually plans to do at all. These correspond to the four remaining planes of the planning blackboard. Thus, the subject makes decisions about data—how long errands should take, how important individual errands are, what the consequences of a particular action might be, etc. He makes decisions about abstract features of plans—what *kinds* of plan decisions might be useful. He makes meta-planning decisions—how to approach the problem and how to constrain and evaluate his plan. Finally, the subject makes executive decisions about how to allocate his cognitive resources during planning. While all of these decisions contribute to the planning process, they do not exhibit a single hierarchical planning structure. For these reasons, we prefer to think in terms of heterarchical plan structures.

Relative Advantages of Hierarchical versus Opportunistic Planning

We might also speculate on the relative merits of hierarchical versus opportunistic planning. The orderly, systematic nature of the top-down process and the simplicity of its hierarchical structure argue in its favor. The recent emphasis on structured programming, a top-down approach to software engineering, reflects these merits (Cf., Dahl, Dykstra, & Hoare, 1972). One might also argue that top-down processes would minimize memory load (Cf., Thorndyke, 1978). The planner could restrict attention to a single area of the hierarchy, rather than attending intermittently to several different areas of the planning space.

On the other hand, planning in tasks fraught with complexity and uncertainty might benefit from less of the discipline imposed by a top-down process. In such complex tasks, general, a priori solutions or problem-solving methods may not exist or may be computationally intractable. Even if some general approach were available, opportunistic planning would free the planner of the burden of maintaining a structurally integrated plan at each decision point. Instead, the planner could formulate and pursue promising partial plans as opportunity suggested.

More importantly, a multi-directional process might produce better plans. It certainly permits more varied plans than a strictly top-down process does. If the planner always began with a fixed high-level plan, she or he could refine it in only a limited number ways. The bottom-up component in multi-directional processing provides a potentially important source of innovation in planning. Low-level decisions and related observations can inspire novel higher-level plans. We observed this in the errand-planning protocol, for example, when the subject generated a high-level design based on observations and decisions made at the lower procedure level. Similarly, Feitelson and Stefik (1977) observed that their expert geneticist deliberately exploited the potential for innovation in bottom-up processing:

> Thus, not only is the planning process largely event driven but sometimes steps are taken somewhat outside the plan of the experiment to make a possibly interesting observation. This kind of behavior reflects the convenience of making certain interesting observations while the equipment is set up. Often this is done to verify the successful completion of an experimental step, but sometimes the observations seem to correspond more to fishing for interesting possibilities. (p. 31)

Resolving the Two Points of View

Although the preceding discussion argues for the proposed opportunistic model in favor of successive refinement models, we would not "reject" either model in favor of the other. Obviously, both models have merit and can best explain different situations. We can suggest three variables which might influence a planner's approach to a particular problem: problem characteristics, individual differences, and expertise.

Problem characteristics could have a major impact on the approach a planner takes. For example, planners might usefully exploit a top-down approach to planning whenever the problem at hand exhibited an inherent hierarchical structure. A study by Byrne (1977) supports this conjecture. His subjects planned dinner menus. As one might expect, subjects appeared to plan menus by deciding on type of dinner (e.g., Chinese dinner, Christmas dinner), main course (e.g., roast beef, turkey), and accompaniments (e.g., cranberry sauce, mashed potatoes). This is a nice example of a hierarchical planning structure. In addition, Byrne's subjects appeared to make decisions within this structure in a top-down fashion.

By contrast, the errand-planning problems discussed in this paper did not exhibit any obvious hierarchical structure. In such circumstances, planners might reasonably resort to more opportunistic methods. Apparently, this is what our subjects did.

A study by Hayes-Roth (1979) provides more direct evidence for the influence of problem characteristics.She successfully induced alternative planning approaches by manipulating the amount of time available for plan execution. For problems that imposed severe time constraints, most subjects adopted a top-down approach. For problems that imposed minimal time constraints, most subjects adopted a bottom-up approach.

The Hayes-Roth (1979) study also provides evidence for the impact of individual differences on planning methods. Many of her subjects exhibited a strong proclivity to adopt a bottom-up approach regardless of problem characteristics. Even with explicit instruction, some subjects persisted in using the bottom-up approach. Other subjects were more flexible, adopting an appropriate approach in response to problem characteristics or instruction.

Finally, planning expertise might influence which planning model a planner brings to bear on particular problems. A practiced planner working on a familiar, constrained problem may possess well-learned, reliable abstract plans for dealing with the problem. This extensive experience may support the application of standard methods for systematically refining abstract plans. On the other hand, a practiced planner working on an unconstrained problem can also exploit opportunistic methods to advantage. Feitelson's and Stefik's (1977) study of the experiment-planning of an expert molecular geneticist provides a nice illustration:

> The experiments described here reflect a combination of goal driven behavior and event driven behavior. . . If there were no goals, behavior might seem very erratic and follow no general course. If there is no event driven component to the planning process, then the experimental procedure must admit no feedback or changes of plans as a result of observations. Thus, no advantage will be made of fortunate observations. What is being suggested here is that the planning in this experiment involved far more exploitation of events and changes of plan according to the events than the authors had anticipated. (p. 30)

One resolution of the apparent conflict between the two models would simply incorporate the top-down model as a special case of the opportunistic model. We have discussed the importance of the problem-solving method a planner brings to bear on a task. This decision can have a major impact on subsequent executive decisions and, consequently, on the planner's progress through the remaining levels of the blackboard. For example, a planner might adopt a "define and successively refine" problem-solving method. Given strict adherence to this method, the planner's formulation of decisions on the plan plane would indeed proceed in a systematic top-down fashion. These are exactly the decisions modeled in the earlier work on top-down planning.

Note that "define and refine" is only one of many problem-solving methods adoptable in the framework of the opportunistic model. Thus, the question is no longer which model is correct, but rather, under what circumstances do planners bring alternative problem-solving methods to bear? We have suggested problem characteristics, individual differences, and expertise as important factors. We should also ask which problem-solving methods work best for different kinds of problems.

8. CONCLUSIONS

The opportunistic model draws on earlier theoretical work in cognitive psychology and artificial intelligence. It incorporates the strongest points of these models with its own assumptions regarding multi-directionality, opportunism, and incrementation in a heterarchical plan structure. We believe that the model is flexible enough to handle the complexity and variability of people's planning behavior. Yet, it is vulnerable to data. We hope the opportunistic model will provide a useful framework for future investigations of the planning process.

REFERENCES

Aho, A. V., Hopcroft, J. E., & Ullman, J. D. *The design and analysis of computer algorithms.* Reading, Mass.: Addison-Wesley, 1974.

Anderson, J. R. *Language, memory, and thought.* Hillsdale, N.J.: Lawrence Erlbaum Associates, 1976.

Anderson, J. R., & Bower, G. H. *Human associative memory.* Washington, D.C.: V. H. Winston, 1973.

Byrne, R. Planning meals: Problem-solving on a real data-base. *Cognition,* 1977, *5,* 287–332.

CMU Computer Science Research Group. Summary of the CMU five-year ARPA effort in speech understanding research. Technical Report, Carnegie-Mellon University, 1977.

Christophides, N. *Graph theory: An algorithmic approach.* New York: Academic Press, 1975.

Dahl, O. J., Dykstra, E. W., & Hoare, C. A. R. *Structured programming.* New York: Academic Press, 1972.

Engelmore, R. S., & Nii, H. P. A knowledge-based system for the interpretation of protein x-ray crystallographic data. Report No. STAN-CS-77-589, Stanford University, 1977.

Erman, L. D., & Lesser, V. R. A multi-level organization for problem solving using many diverse cooperating sources of knowledge. *Proceedings of the Fourth International Joint Conference on Artificial Intelligence,* Tbilisi, USSR, 1975, 483–490.

Ernst, G. W., & Newell, A. *GPS: A case study in generality and problem solving*. New York: Academic Press, 1969.

Fahlman, S. E. A planning system for robot construction tasks. *Artificial Intelligence*, 1974, *5*, 1–49.

Feitelson, J., & Stefik, M. A case study of the reasoning in a genetics experiment. Heuristic Programming Project, Working Paper 77–18, Department of Computer Science, Stanford University, April 1977.

Fikes, R. E. Knowledge representation in automatic planning systems. In A. K. Jones (ed.), *Perspectives on computer science*. New York: Academic Press, 1977.

Fikes, R. E., & Nilsson, N.J. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 1971, *2*, 189–203.

Hayes-Roth, B. Flexibility in executive strategies. N:1170, The Rand Corporation, Santa Monica, California, 1979.

Hayes-Roth, B., Hayes-Roth, F., Rosenschein, S., & Cammarata, S. Modeling planning as an incremental, opportunistic process. *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*, Tokyo, Japan, 1979.

Hayes-Roth, B., and Thorndyke, P. Decision-making during the planning process. N:1213, The Rand Corporation, Santa Monica, California, 1979.

Hayes-Roth, F., & Lesser, V. R. Focus of attention in the Hearsay-II speech understanding system. *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, Boston, Mass., 1977, 27–35.

Hayes-Roth, F., Waterman, D. A., & Lenat, D. E. Principles of pattern-directed inference systems. In. Waterman D. A. & Hayes-Roth F. (eds.), *Pattern-directed inference systems*. New York: Academic Press, 1978.

Lesser, V. R., Fennell, R. D., Erman, L. D., & Reddy, D. R. Organization of the Hearsay-II speech understanding system. *IEEE Transactions on Acoustics, Speech and Signal Processing*, ASSP-23, 1975, 11–23.

Lesser, V. R., & Erman, L. D. A retrospective view of the Hearsay-II architecture. *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, Boston, Mass., 1977, 790–800.

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

Nii, H. P., & Feigenbaum, E. A. Rule-based understanding of signals. In D. A. Waterman & F. Hayes-Roth (eds.), *Pattern-directed inference systems*. New York: Academic Press, 1978.

Prager, J., Nagin, P., Kohler, R., Hanson, A., & Riseman, E. Segmentation processes in the VISIONS system. *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, Boston, Mass., 1977.

Rumelhart, D. E. Toward an interactive model of reading. Technical Report 56, Center for Human Information Processing, University of California at San Diego, La Jolla, Ca., 1976.

Rumelhart, D., Lindsay, P. H., & Norman, D. A. A process model for long-term memory. In E. Tulving & W. Donaldson (eds.), *Organization of memory*. New York: Academic Press, 1972.

Sacerdoti, E. D. Planning in a hierarchy of abstraction spaces. *Artificial Intelligence*, 1974, *5*, 115–135.

Sacerdoti, E. D. A structure for plans and behavior. Technical Note 109, Standford Research Institute, Menlo Park, California, August, 1975.

Selfridge, O. Pandemonium: A paradigm for learning. *Symposium on the mechanization of thought*. London: HM Stationery Office, 1959.

Soloway, E. M., & Riseman, E. M. Knowledge-directed learning. *Proceedings of the Workshop on Pattern-Directed Inference Systems*, special edition of the *SIGART Newsletter*. New York: Association for Computing Machinery, June, 1977.

Sussman, G. J. A computational model of skill acquisition. AI TR-297, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Mass., 1973.

Thorndyke, P. W. Pattern-directed processing of knowledge from text. In D. A. Waterman & F. Hayes-Roth (eds.), *Pattern-directed inference systems*. New York: Academic Press, 1978.
Winograd, T. A program for understanding natural language. *Cognitive Psychology*, 1972, *3*.