

# STRUCTURED PLANNING AND DEBUCCING

Mark L. Miller and Ira P. Goldstein  
 Artificial Intelligence Laboratory  
 Massachusetts Institute Of Technology  
 Cambridge, Massachusetts 02139

## Abstract

The SPADE theory uses linguistic formalisms to model the program planning and debugging processes. The theory begins with a taxonomy of basic planning concepts covering strategies for identification, decomposition and reformulation. A handle is provided for recognizing interactions between goals and deriving a linear solution. A complementary taxonomy of rational bugs and associated repair techniques is also provided. SPADE introduces a new data structure to facilitate debugging -- the derivation tree of the program.

SPADE generalizes recent work in Artificial Intelligence by Sussman and Sacerdoti on automatic programming, and extends the theory of program design developed by the Structured Programming movement. It provides a more structured information processing model of human problem solving than the production systems of Newell and Simon, and articulates the type of problem solving curriculum advocated by Papert's Logo Project.

### 1. A Multi-Faceted Approach

The SPADE theory is being developed in three contexts:

1. Education: an editor called SPADEE-0 has been implemented that encourages students to define and debug programs in terms of explicit SPADE design choices, thereby providing a highly structured programming environment.
2. AI: an automatic programmer called PATN has been designed using an augmented transition network embodiment of the SPADE theory. This results in a framework which unifies recent work on planning and debugging by Sacerdoti [75] and Sussman [75].
3. Psychology: a parser called PAZATN has been designed that applies the SPADE theory to the analysis of programming protocols. PAZATN produces a parse of the protocol that delineates the planning and debugging strategies employed by the problem solver. PAZATN extends the series of automatic protocol analysers developed at Carnegie-Mellon University [Waterman & Newell 72, 73; Bhaskar & Simon 76].

Hand-simulations of PATN and PAZATN on elementary programming problems and informal experiments with the SPADEE-0 editor attest to the theory's cogency in accounting for a wide range of planning and debugging techniques [Goldstein & Miller 76a,b; Miller & Goldstein 76b,c,d].

### 2. A Linguistic Analogy

In developing a representation for problem solving techniques, we have been guided by an analogy to computational linguistics, for three reasons.

1. The concepts and algorithms of computational linguistics, though originally intended to explain the nature of language *per se*, supply perspicuous yet powerful descriptions of complex computations in general.
2. Computational linguistics decomposes computations into syntactic, semantic, and pragmatic components. This decomposition clarifies the explanation of complex processes when viewed in the following manner: syntax formalizes the range of possible decisions; semantics the problem description, and pragmatics the procedural relationship between the two.
3. Computational linguistics has undergone an evolution of procedural formalisms, beginning with finite state automata, later employing recursive transition networks (context free grammars), next moving on to augmented transition networks, and culminating in the current set of theories involving frames [Minsky 75, Winograd 75, Schank 75]. Each phase captured some properties of language, but was incomplete and required generalisation to more powerful and

elaborate formalism\*. Following this evolutionary sequence illuminates the complexity of language theory. We have pursued a similar evolutionary approach to clarify the complexity of problem solving processes.

\* To date, our theory of program design has evolved as follows: we first explored context free grammars for planning and debugging, and subsequently their generalisation to ATN's; we then examined the metaphor of protocol analysis as parsing, initially using the planning and debugging grammars to reveal the constituent structure of protocols and later using the derivations produced by the ATN formalism; and, most recently, we have studied the use of a chart-based parser to discover these analyses.

### 3. A Grammatical Theory of Planning

The basis for SPADE is a taxonomy of frequently observed planning concepts (fig. 1). We arrived at this taxonomy by

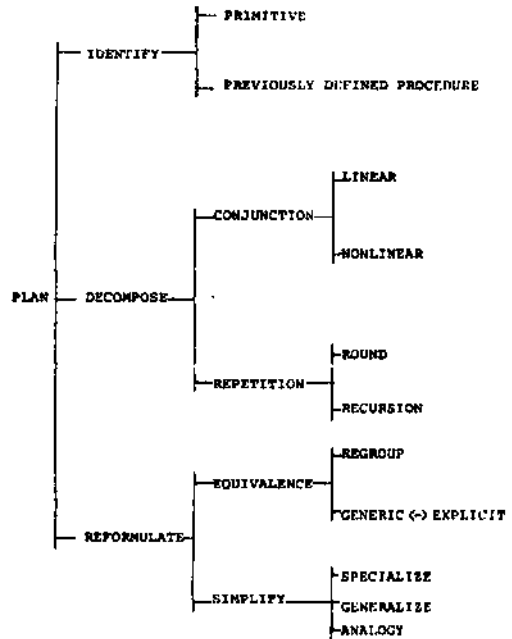


FIGURE 1  
 TAXONOMY OF PLANNING CONCEPTS

introspection, by examining problem solving protocols [Miller & Goldstein 76b], by studying the literature on problem solving [Pglya 57, 65, 68; Newell & Simon 72; Sussman 75; Sacerdoti 75], and by enumerating techniques for finding procedural solutions to problems expressed at predicate calculus formulae [Emden & Kowalski 76]. This last criterion demonstrates that the taxonomy is currently incomplete -- for example, techniques for handling disjunctions have not yet been analysed thoroughly enough to warrant inclusion. However, the taxonomy is adequate for a wide range of elementary programming problems.

There are three major classes of plans in the taxonomy: identification, decomposition, and reformulation. Identification means recognizing a problem at previously solved. Decomposition refers to strategies for dividing a problem into simpler sub-problems. Reformulation plans alter the problem description, seeking a representation which is more amenable to identification or decomposition. The figure indicates how these classes of plans are further subdivided in the SPADE theory.

Planning, according to the theory, is a process in which the problem solver selects the appropriate plan type, and then carries out the subgoals defined by that plan applied to the current problem. From this viewpoint, the planning taxonomy represents a decision tree of alternative plans. The decision process can be modeled by the context free grammar given below. The grammar explicitly states which planning rules involve recursive application of solution techniques to subgoals: setup, interface, mainstep, cleanup, and

parallel.

The grammar is written using the following syntax "I" is disjunction, V is ordered conjunction, "&" is unordered conjunction, "<..>" is iteration, [ ] is optionality, end a lower case phrase in quotation marks (eg, "repeat step"/ describes \$ lexical item which is not further expanded in the grammar.

```

PLAN -> IDENTIFY | DECOMPOSE | REFORMULATE
IDENTIFY -> PRIMITIVE | DEFINED
DEFINED -> "call user subprocedure" & PLAN
DECOMPOSE -> CONJUNCTION | REPETITION
CONJUNCTION -> SEQUENTIAL | PARALLEL
SEQUENTIAL -> [SETUP] + (MAINSTEP + [INTERFACE])* + [CLEANUP]
PARALLEL -> <PLAN>*
SETUP -> PLAN
MAINSTEP -> PLAN
INTERFACE -> PLAN
CLEANUP -> PLAN
REPETITION -> ROUND | RECURSION
ROUND -> ITER-PLAN | TAIL-RECUR
ITER-PLAN -> repeat step" + SEQUENTIAL
TAIL-RECUR -> "itop step" + SEQUENTIAL + "recur step"
  
```

The SPADE theory is not restricted to any particular domain. However, to provide concrete examples, we have concentrated on problems from elementary Logo graphic\* programming [Papert 71]. This domain was chosen because of the availability of extensive student performance data. The grammar rules for primitives in this domain are:

```

PRIMITIVE -> VECTOR | ROTATION | PENSTATE
VECTOR -> (FORWARD | BACK) + "number"
ROTATION -> (LEFT | RIGHT) + "number"
PENSTATE -> PENUP | PENDOWN
  
```




FIGURE 2

A typical task undertaken by beginners in the Logo environment is to draw a wishingwell picture using the computer (fig. 2). Fig. 3 illustrates a solution to the wishingwell problem with its hierarchical annotation according to our planning grammar.

The grammar characterizes the decision process involved in selecting plans from the taxonomy. We illustrate its utility in the next two sections by constructing an editor that embodies the grammar and analyzing debugging in terms of the grammar. Then we show how the grammar can be augmented to include not only the syntax of plans, but their semantics and pragmatic\* a\* well.

#### 4. SPADEE-0. A Planning Assistant

One reason for calling our theory of planning and debugging structured is to emphasise the link between our research and the Structured Programming movement. Dahl, Dijkstra, and Hoare [72] call for a style of programming which reflects coherently structured problem solving; but a detailed formalization of what this style entails is lacking. Our efforts in this direction, therefore, supplement the work of Dijkstra and others. How can we judge, though, whether a particular grammar of plans captures the planning decisions involved in solving problems for some domain? One methodology is to incorporate the grammar into an editor (SPADEE-0) whose purpose is to augment and direct the capabilities of a human user. The critical question then becomes the extent to which the editing system aids or hinders the user.

Suppose a problem solver is defining a Logo program for drawing the wishingwell shown earlier. In SPADEE-0, this is accomplished by applying the planning grammar in generative mode:

- 1a. What is the name of your procedural
- 1b. >WV
- 2a. The rule is: PWK -> IDENTIFY | DECOMPOSE | REFORMULATE  
What now?
- 2b. >DECOMPOSE




FIGURE 3. PATH'S HIERARCHICAL DERIVATION TREE FOR THE WISHINGWELL TASK

- 3a. The rule is: DECOMPOSE -> CONJUNCTION | REPETITION.
- 3b. >CONJUNCTION
- 4a. The rule is: CONJUNCTION -> SEQUENTIAL ...  
-> [SETUP] + <MAINSTEP+[INTERFACE]>\* + [CLEANUP].  
Do you want to define the optional SETUP?
- 4b. >Later

SPADEE-0 thus encourages users to articulate their design decisions in top-down order. At the same time, the system allows the user to escape from this strict discipline if an alternative solution order seems preferable. This was illustrated by the user's "later" instruction, which suspends the current goal for subsequent solution.

SPADEE-0 was implemented by assigning an interpretive procedure to each grammatical operator. In essence, the editor is a bookkeeper for the user's goal tree. Though simple, the editor serves three useful purposes.

1. From an educational standpoint, the editor encourages students to articulate their problem solving strategies. The fundamental hypothesis of the Logo Project, as presented by Papert [71], is that such articulate problem solving is beneficial to the learner. SPADEE-0, with its extreme form of articulation, provides an experimental vehicle for evaluating Papert's claim. Our experiment will be to test whether students exposed to SPADEE-0 learn Logo faster than controls whose problem solving is more tacit.
2. From an AI standpoint, its use will indicate whether the planning grammar is adequate, or whether certain plans are not present that competent problem solvers feel are necessary.
3. From a psychological standpoint, we will collect transcripts of individuals using the editor and formulate personal grammars based on the particular rules usually employed by each user. The personal grammar will model the problem solving skills of that individual. In the past we have manually analysed protocols from standard Logo. SPADEE-0 protocols, with their explicit planning choices, should be





SQUARE X;  
 TRIANGLE Y;  
 INSIDE X Y;

then a decomposition that draws the square and triangle independently and then attempts to fit them together to achieve the inside relation will fail. However, a problem description of the following form allow\* a successful decomposition:

SQUARE X, WITH SIDE \* 100;  
 TRIANGLE Y, WITH SIOE - 300;  
 CENTER OF X ■ CENTER OF Y.

The INTRAKCTIONS predicate is a conjunction of tests on the model register. Each test is responsible for detecting a given non-linearity. A corresponding action modifies the model, adding new statements to make the interaction explicit. The REFINKMKNT loop is the repository for what Sussman [75] calls the Critics Gallery. The theoretical progress of PATN is to integrate the Critics Gallery concept into a theory of planning. In Sussman's HACKER, the critics gallery and library of programming techniques were separate modules: there was no integrated theory.

Of course, at any point in time the system may be unaware of a given type of non-linearity. In such cases, the absence of an interaction test will lead to a sequential decomposition that ultimately fails. The design of a program for debugging such failures is the subject of the next section.

### 7. DAPR — An ATN for Debugging

PATN can make mistakes. That is, PATN will sometimes introduce what we term *rational hug\** into its plans, due to making arc transitions with imperfect knowledge of subtleties or interactions in the task domain. Hence, PATN must be equipped with a complementary debugging module, DAPR (fig. 7).

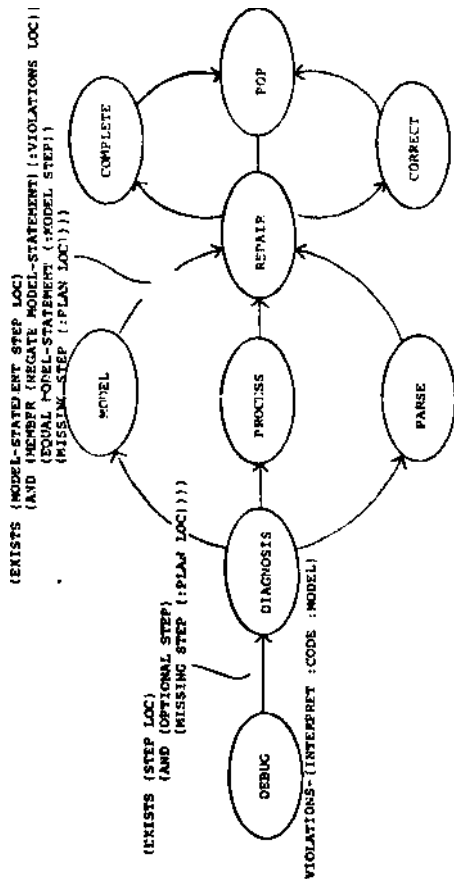


FIGURE 7 DAPR: PATH'S DEBUGGING ATN

DAPR's task is easier than that of RAID: DAPR must analyze the closed set of bug types to which PATN is subject, whereas RAID is intended to assist human programmers in finding and correcting a wide assortment of bugs. DAPR employs three diagnostic techniques:

model, process, and plan diagnosis. Model diagnosis is the basic technique. It amounts to comparing the effects of executing a plan to a formal description of its goals, to determine if, and in what fashion, the plan has failed. Another DAPR technique, based on Susctnan's HACKKR [75], is examining the state of the process at the time of the error manifestation. Plan diagnosis, a DAPR first, involves examining the *caveat\** left by the planner as various nodes were constructed.

DAPR will also be used to provide additional guidance to RAID. This illustrates the synergism possible when educational, psychological and AI facets of a cognitive theory are studied in an integrated fashion. This integration is further exemplified in the next section when we apply the SPADK theory to protocol analysis.

### 8. PAZATN, a Protocol Analyser

As soon as one has an heuristically adequate theory of program design, it is natural to ask, "Can the theory provide an account of how *people* design programs?". An experimental technique we employ for answering this question is the analysis of protocols collected during problem solving sessions. By adopting this methodology we follow the precedent established in seminal studies conducted at Carnegie Mellon University [Newell & Simon 72; Waterman & Newell 72, 73; Rhaskar & Simon 76]. Our work extends their approach along three dimensions.

1. With the exception of the recent Rhaskar & Simon effort, the CMU studies have been restricted to very limited domains such as cryptarithmic. Rather than limiting the task domain, we limit their range of responses. Typically protocols are transcriptions of think-aloud verbalisations; we focus on the more restricted interactions arising from a problem solving session at a computer console. The analysis task in this setting is to interpret user actions — editing, executing, tracing, etc. — in terms of the SPADK theory of planning and debugging.

2. The CMU theory centers on the *production system model*. Although productions are Turing universal, they encourage a less hierarchical, less local program organization than the linguistic formalisms of the SPADK theory. In PATN, each arc transition, consisting of a predicate and an action, can be thought of as a production. However, PATN organizes these productions into local contexts, each of which consists of the arcs exiting from a given node. Not all of the arc productions are present at any moment in time; an arc is present only when the problem solver is at the relevant node. In the production systems discussed in *Human Problem Solving* [Newell & Simon 72], all of the productions are always present and are tested in serial order.

3. CMU analyses are based on the *problem behavior graph*. Pursuing an analogy to computational linguistics, we define an interpretation of a protocol to be a *parse tree* supplemented by semantic and pragmatic annotation. The parse tree characterizes the constituent structure of the protocol. Semantic and pragmatic annotation — variables and assertions attached to nodes of the parse tree — formalize the problem description and the rationale for particular planning choices. Annotated parse trees closely reflect the local structure of PATN's linguistic problem solving machinery, leading more directly to inferences regarding individual differences than is evident from problem behavior graphs.

Ruvcn Brooks [75] applied the CMU approach to the programming domain, developing a model of *coding* — the translation of high level plans into the statements of a particular programming language — and testing the model by analyzing protocols. His model is a set of production rules whose conditions match the patterns of plan elements and whose actions generate code statements. Protocols are analyzed manually, with the experimenter attempting to infer the plan which is then expanded by the production system into code paralleling that of the protocol. The processes of understanding the problem, generating the plan, and debugging are not formalized. SPADE goes beyond this in that it can be used to parse protocols and that the parse constitutes a formal hypothesis regarding not only the coding knowledge but also the planning and debugging strategies employed by the problem solver.

[Miller & Goldstein 76b] provides an example of such analysis being performed by hand. The example is a segment from a protocol several hundred lines long in which a high school student uses Logo



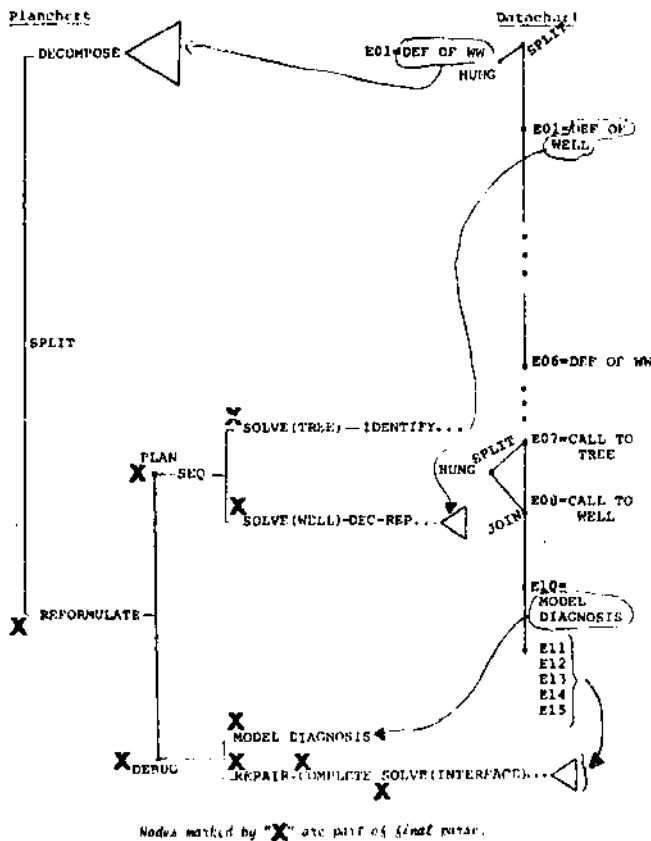


FIGURE 11 DATACHART LINKING PROTOCOL EVENTS TO PLANCHART LEAVES

a hypothetical dialogue with SPADEE-2, representing the original 'SPADEE-O augmented by both PATN and PAZATN.

- 1a. Solving for WISHINGVEIL. Pending subgoals are:  
 ROOF, POLE, WELL, interfaces. What now?  
 1b. >SQUARE
- 2a. OK. WELL has been solved by a call to SQUARE.  
 SQUARE has already been solved. What now? \*

PAZATN will increase the editor's flexibility in handling ambiguous events, and in alleviating what might seem to some users to be an excessive allocation of time and effort to the planning phase

### 9. Conclusions

The use of tools from computational linguistics — grammars, ATN's, derivation trees, parsing algorithms, charts — has led to a perspicuous representation for a theory of planning and debugging. Computational linguistics is also responsible for suggesting the propitious decomposition of problem solving processes into components involving syntactic, semantic and pragmatic knowledge.

Our multi-faceted approach — studying problem solving in the three distinct contexts of AI, education, and psychology — holds out the possibility of a synergistic effect. But proof of this must await further experimentation. Although all of the programs have been designed and hand-simulated, as of this writing only the SPADKK-0 editor has been implemented. Furthermore, the theory has not yet been exercised in enough contexts to prove its generality. However, at least for the three domains in which the theory has been explored — Logo, the Blocks World, and elementary calculus — it has provided a unified treatment of plans and bugs, a significant stride for a theory of program design.

The automatic problem solving aspect was supported by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract NO0014-75-C-0643, the educational aspect by the National Science Foundation under grant C40708X, and the protocol analysis aspect by the Bolt Beranek & Newman Intelligent Instructional Systems Group under contract MDA 903-76-C-0108 jointly sponsored by Advanced Research Projects Agency, Air Force Human Resources Laboratory, Army Research Institute, and Naval Personnel Research & Development Center.

### References

- Bhasker, R., and H Simon, 1976 "Problem Solving in Semantically Rich Domains: An Example, from Engineering Thermodynamics" Carnegie-Mellon U, CIP Working Paper 314
- Brooks, R., 1975 A Model of Human Cognitive Behavior in Writing Code for Computer Programs Carnegie-Mellon U, Report AFOSR-TR-1084
- Dahl, O J., E. Dijkstra and C.A.R. Hoare 1972. Structured Programming London, Academic Press
- Emden, M Van, and R Kowalski, 1976 "The Semantics of Predicate Logic as a Programming Language" JACM 23 4, pp 733-742
- Goldstein, I, and H Miller, 1976a AI Based Personal Learning Environments MIT AI Memo 384
- Goldstein, I, and M Miller, 1976b Structured Planning and Debugging A Linguistic Theory of Design MIT AI Memo 387
- Kaplan, R., 1973 "A General Syntactic Processor" in R Rustin (ed), Natural Language Processing, NY, Algorithmics Press, pp 193-241
- Key, M., 1973. "The MIND System" in R Rustin (ed), Natural Language Processing, NY, Algorithmics Press, pp 155-186
- Lesser, V, R Fennell, L Emen and DR Reddy, 1975 "Organization of the Heersoy II Speech Understanding System" IEEE Transactions on Acoustics, Speech, and Signal Processing Assp-23., pp. 11-24
- Miller, M, and I Goldstein, 1976b Parsing Protocols Using Problem Solving Grammars MIT AI Memo 385
- Miller, M, and I Goldstein, 1976c SPADE: A Grammar Based Editor For Planning and Debugging Programs MIT AI Memo 386
- Miller, M, and I Goldstein, 1976d PAZATN: A Linguistic Approach To Automatic Analysis of Elementary Programming Protocols MIT AI Memo 388
- Minsky, M, 1975 "Frame-Systems: A Framework for Representation of Knowledge" in P Winston (ed), The Psychology of Computer Vision, NY, McGraw-Hill
- Newell, A, and H. Simon, 1972. Human Problem Solving N.J., Prentice-Hall
- Papert, S, 1971 Teaching Children Thinking MIT AI Memo 247
- Paxton, W, and A. Robinson, 1975. "System Integration and Control in a Speech Understanding System" AJCL 5, pp. 5-16
- Polya, G, 1957 How to Solve It NY, Doubleday Anchor Books
- Polya, G, 1965 Mathematical Discovery (Vols 142) NY, Wiley and Sons
- Polye, G, 1968 Mathematics and Plausible Reasoning (Vols t&2) NJ, Princeton U. Press
- Sacerdoti, E, 1975 "The Nonlinear Nature of Plans" AIJCAI, Tbilisi, Georgia. USSR, pp 206-218
- Schank, R, 1975. "Using Knowledge to Understand" in R Schank & B Nesh-Webber, Theoretic Issues in Natural Language Processing, pp 117-121
- Sussman, G, 1975 A Computational Model of Skill Acquisition NY, American Elsevier
- Waterman, D, and A Newell, 1972 Preliminary Results with a System For Automatic Protocol Analysis Carnegie-Mellon U, CIP Working Paper 211
- Waterman, D, and A Newell, 1973 "PAS-II: An Interactive Task-Free Version of An Automatic Protocol Analysis System" SIJCAI, Stanford, Ca, pp 431-445
- Winograd, T., 1975. "Frame Representations and the Declarative-Procedural Controversy" in O. Bobrow & A. Collins, Representation and Understanding, Academic Press, pp. 185-210.
- Woods, W, 1970 "Transition Network Grammars for Natural Language Analysis" CACM 1310, pp 591-606