

DOCUMENT RESUME

ED 275 300

IR 012 358

AUTHOR White, Barbara Y.; Frederiksen, John R.
TITLE Progressions of Qualitative Models as a Foundation for Intelligent Learning Environments. Report No. 6277.
INSTITUTION BBN Labs, Inc., Cambridge, MA.
SPONS AGENCY Army Research Inst. for the Behavioral and Social Sciences, Alexandria, Va.; Office of Naval Research, Arlington, Va. Personnel and Training Research Programs Office.
PUB DATE May 86
NOTE 107p.
PUB TYPE Viewpoints (120) -- Reports - Evaluative/Feasibility (142)

EDRS PRICE MF01/PC05 Plus Postage.
DESCRIPTORS *Artificial Intelligence; *Cognitive Processes; *Electric Circuits; *Experiential Learning; *Learning Strategies; *Man Machine Systems; *Models; *Protocol Analysis; *Schemata (Cognition); *Simulated Environment

ABSTRACT

This report discusses the importance of presenting qualitative, causally consistent models in the initial stages of learning so that students can gain an understanding of basic electrical circuit concepts and principles that builds on their preexisting ways of reasoning about physical phenomena, and it argues that tutoring environments must help students to acquire multiple mental models that embody alternative conceptualizations of the domain. An outline of model types and an overview of the learning environment based upon progression of qualitative, causal models are presented; issues related to the design of zero-order, qualitative models for circuit behavior are discussed; different types of possible evolutions of a student's mental model are enumerated; and one path through the space of possible model evolutions, a curriculum for helping students learn troubleshooting, is presented. The learning strategies that such a tutoring system architecture facilitates are then described and some instructional trials of the system are briefly discussed. Finally, an outline of a set of alternative mental models that a student should acquire in order to more deeply understand how circuits work is presented, and the extensibility of this approach to the creation of intelligent learning environments for other subject domains is considered. References and a distribution list are included. (KM)

 * Reproductions supplied by EDRS are the best that can be made *
 * from the original document. *

ED 275 300

BBN Laboratories Incorporated



A Subsidiary of Bolt Beranek and Newman Inc.

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION
CENTER (ERIC)

- This document has been reproduced as received from the person or organization originating it.
- Minor changes have been made to improve reproduction quality.

Report No. 6277

- Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

Progressions of Qualitative Models as a Foundation for Intelligent Learning Environments

Barbara Y. White and John R. Frederiksen

May 1986

Prepared for:

Office of Naval Research and
Army Research Institute

IR 012 358

BEST COPY AVAILABLE

Approved for public release; distributions unlimited

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED		1b. RESTRICTIVE MARKINGS			
2a. SECURITY CLASSIFICATION AUTHORITY		3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release: distribution unlimited. Reproduction in whole or in part is permitted for any purpose of the United States Government			
2b. DECLASSIFICATION / DOWNGRADING SCHEDULE		5. MONITORING ORGANIZATION REPORT NUMBER(S)			
4. PERFORMING ORGANIZATION REPORT NUMBER(S) BBN Laboratories Report#6277					
6a. NAME OF PERFORMING ORGANIZATION BBN Laboratories	6b. OFFICE SYMBOL (if applicable)	7a. NAME OF MONITORING ORGANIZATION Personnel & Training Research Programs Office of Naval Research			
6c. ADDRESS (City, State, and ZIP Code) 10 Moulton Street Cambridge, MA. 02238		7b. ADDRESS (City, State, and ZIP Code) Code 1142 PT Arlington, VA. 22217			
8a. NAME OF FUNDING / SPONSORING ORGANIZATIONS ONR & ARI	8b. OFFICE SYMBOL (if applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-82-C-0580			
8c. ADDRESS (City, State, and ZIP Code)		10. SOURCE OF FUNDING NUMBERS			
		PROGRAM ELEMENT NO. 1721319.WIAE	PROJECT NO. RA042-06	TASK NO. 042-06-02	WORK UNIT ACCESSION NO. NR154-493
11. TITLE (Include Security Classification) Progressions of Qualitative Models as a Foundation for Intelligent Learning Environments					
12. PERSONAL AUTHOR(S) Barbara Y. White and John R. Frederiksen					
13a. TYPE OF REPORT Technical Report	13b. TIME COVERED FROM Sep 82 TO Dec 85	14. DATE OF REPORT (Year, Month, Day) May, 1986	15. PAGE COUNT 88		
16. SUPPLEMENTARY NOTATION					
17. COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)			
FIELD	GROUP	AI and Education Cognitive Modeling			
		Intelligent Tutoring Systems Qualitative Models			
19. ABSTRACT (Continue on reverse if necessary and identify by block number) The design of our intelligent learning environment is based upon a theory of expertise and its acquisition. We find that when reasoning about physical systems, experts utilize a set of mental models. For instance, they may use qualitative as well as quantitative models, and behavioral as well as functional models. The transition from novice to expert status can be regarded as a process of model evolution: students formulate a series of upwardly compatible models, each of which is adequate for solving some subset of problems within the domain. Further, students need to evolve not just a single model, but rather a set of models that embody alternative conceptualizations of the domain. Finally, we claim that in the initial stages of learning, students should focus on the acquisition of qualitative models: quantitative models should be introduced only after the domain is understood in qualitative terms.					
20. DISTRIBUTION / AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS		21. ABSTRACT SECURITY CLASSIFICATION			
22a. NAME OF RESPONSIBLE INDIVIDUAL		22b. TELEPHONE (Include Area Code)	22c. OFFICE SYMBOL		



In this article, we focus primarily on qualitative, behavioral models of electrical circuit operation designed to make the causality of circuit behavior derive clearly from basic physical principles. The constraints on model evolution, in terms of causal consistency and learnability, are discussed and a sequence of models that embody a possible transformation from novice to expert status is outlined.

The learning environment we have constructed lets students solve problems, hear explanations, and perform experiments, all in the context of interacting with a dynamic simulation of circuit behavior. However, unlike most simulations, the underlying model is qualitative not quantitative. Further, the simulation is performed not by a single model, but rather by a progression of models that increase in sophistication in concordance with the evolution of the students' understanding of the domain.

Viewing instruction as producing in the student a progression of models permits a tutoring system architecture with elegant properties. Within our system, the student model, the tutor, and the domain simulation are incorporated within the single model that is active at any point in learning. This model is used to simulate the domain phenomena, is capable of generating explanations by articulating its behavior, and furnishes a desired model of the students' reasoning at that particular stage in learning. The progression of models also enables the system to select problems and generate explanations that are appropriate for the student at any point in the instructional sequence. In order to motivate students to transform their models into new models, they are given problems that the new model can handle but their present model cannot. This evolution of models also enables the system to focus its explanations on the difference between the present model and the new model.

Such a system architecture also permits a variety of pedagogical strategies to be explored within a single instructional system. Since the system can turn a problem into an example by solving it for the student, the students' learning can be motivated by problems or by examples. That is, students can be presented with problems and only see examples if they run into difficulty; alternatively, they can see examples first and then be given problems to solve. Also, by working within the simulation environment, students can use a circuit editor to construct their own problems and thus explore the domain in a more open ended fashion. The system is capable of generating runnable qualitative models for any circuit that the student or instructional designer might create. Further, the learning process can be managed either by the system or by the student. For example, students can be given a map of the problem space and can decide for themselves what class of problems to pursue next or even what pedagogical strategy they want to employ.

**Progressions of Qualitative Models
as a Foundation for
Intelligent Learning Environments**

Barbara Y. White and John R. Frederiksen

BBN Laboratories

Abstract

The design of our intelligent learning environment is based upon a theory of expertise and its acquisition. We find that when reasoning about physical systems, experts utilize a set of mental models. For instance, they may use qualitative as well as quantitative models, and behavioral as well as functional models. The transition from novice to expert status can be regarded as a process of model evolution: students formulate a series of upwardly compatible models, each of which is adequate for solving some subset of problems within the domain. Further, students need to evolve not just a single model, but rather a set of models that embody alternative conceptualizations of the domain. Finally, we claim that in the initial stages of learning, students should focus on the acquisition of qualitative models: quantitative models should be introduced only after the domain is understood in qualitative terms.

In this article, we focus primarily on qualitative, behavioral models of electrical circuit operation designed to make the causality of circuit behavior derive clearly from basic physical principles. The constraints on model evolution, in terms of causal consistency and learnability, are discussed and a sequence of models that embody a possible transformation from novice to expert status is outlined.

The learning environment we have constructed lets students solve problems, hear explanations, and perform experiments, all in the context of interacting with a dynamic simulation of circuit behavior. However, unlike most simulations, the underlying model is qualitative not quantitative. Further, the simulation is performed not by a single model, but rather by a progression of models that increase in sophistication in concordance with the evolution of the students' understanding of the domain.

Viewing instruction as producing in the student a progression of models permits a tutoring system architecture with elegant properties. Within our system, the student model, the tutor, and the domain simulation are incorporated within the single model that is active at any point in learning. This model is used to simulate the domain phenomena, is capable of generating explanations by articulating its behavior, and furnishes a desired model of the students' reasoning at that particular stage in learning. The progression of models also enables the system to select problems and generate explanations that are appropriate for the student at any point in the instructional sequence. In order to motivate students to transform their models into

new models, they are given problems that the new model can handle but their present model cannot. This evolution of models also enables the system to focus its explanations on the difference between the present model and the new model.

Such a system architecture also permits a variety of pedagogical strategies to be explored within a single instructional system. Since the system can turn a problem into an example by solving it for the student, the students' learning can be motivated by problems or by examples. That is, students can be presented with problems and only see examples if they run into difficulty; alternatively, they can see examples first and then be given problems to solve. Also, by working within the simulation environment, students can use a circuit editor to construct their own problems and thus explore the domain in a more open ended fashion. The system is capable of generating runnable qualitative models for any circuit that the student or instructional designer might create. Further, the learning process can be managed either by the system or by the student. For example, students can be given a map of the problem space and can decide for themselves what class of problems to pursue next or even what pedagogical strategy they want to employ.

Table of Contents

1. Introduction	1
1.1. An Overview of the Paper	1
1.2. Mental Models	1
1.2.1. The Importance of Qualitative Reasoning	3
1.2.2. Causal Consistency	4
1.3. Learning as a Process of Model Transformation	6
1.3.1. The Order of a Model	7
1.3.2. The Degree of a Model	8
1.4. An Overview of the Learning Environment	8
2. Qualitative Causal Models of Circuit Behavior	13
2.1. The Instructional Need for Zero Order Models	13
2.2. The Zero Order Models	16
3. Troubleshooting	36
3.1. The troubleshooting algorithms	36
3.2. Facilitating troubleshooting	37
4. Model Evolutions	39
4.1. Types of Model Evolution	40
4.2. The Problem of Modifiability	41
4.3. The Path of Model Evolutions	44
4.3.1. A Zero Order Curriculum	45
4.4. Further Model Evolutions	49
5. The Learning Environment	50
5.1. Problem Types	50
5.2. Problem Selection	52
5.3. Design Philosophy	56
5.4. Learning Strategies	57
5.5. Instructional Effectiveness	58
6. Multiple Alternative Conceptualizations	63
6.1. Zero Order Qualitative Models	63
6.1.1. Models of Functional Interactions Among Devices	65
6.2. First Order Qualitative Models	69
6.3. Quantitative Extensions	77
6.4. Model Similarities	79
6.5. Understanding of a Domain	80
7. The Extendibility of This Approach to Other Domains	82
8. Acknowledgements	84
9. References	84

1. Introduction

1.1. An Overview of the Paper

This paper begins by presenting the theoretical foundations of our approach to constructing intelligent learning environments. In particular, we argue for the importance of presenting, in the initial stages of learning, qualitative, causally consistent models so that students can gain an understanding of basic circuit concepts and principles that builds on their preexisting ways of reasoning about physical phenomena. We argue, in addition, that tutoring environments must help students to acquire multiple mental models that embody alternative conceptualizations of the domain, and an outline of model types is presented. Then, an overview of the learning environment based upon progressions of qualitative, causal models is given.

Next, the paper discusses issues related to the design of what we have termed zero-order, qualitative models for circuit behavior. In these models, circuit functioning is represented as a series of changes in the qualitative states of devices within the circuit. The models embody basic circuit concepts and principles, and can generate causal accounts of circuit behavior that are compatible with those of higher order models. They are also models of how one wants students to reason at a given stage in learning. The paper goes on to enumerate different types of possible evolutions of a student's mental model and describes one path through the space of possible model evolutions that we have implemented, i.e. a curriculum for helping students learn troubleshooting. The learning strategies that such a tutoring system architecture facilitates are then described and some instructional trials of the system are briefly discussed.

Finally, the paper outlines a set of alternative mental models that a student should acquire in order to more deeply understand how circuits work. The paper concludes with a discussion of the extensibility of this approach to the creation of intelligent learning environments for other subject domains.

1.2. Mental Models

By mental model we mean a knowledge structure that incorporates both declarative knowledge (e.g., device models) and procedural knowledge (e.g., procedures for determining distributions of voltages within a circuit), and a control structure that determines how the procedural and declarative knowledge are used in solving problems (e.g., simulating the behavior of a circuit).

The theoretical framework we adopt is that electrical expertise can be captured by a small set of mental models that embody alternative conceptualizations of circuit operation. For instance, experts utilize qualitative as well as quantitative models, and behavioral as well as functional models. We adopt this viewpoint based upon both empirical and theoretical research. Our models are derived from extensive studies of an expert troubleshooter who teaches in a technical high school (White & Frederiksen, 1984). The initial mental models that we try to give students are also influenced by studies of novices reasoning about circuit problems (e.g., Cohen et al., 1983). Further, the model designs draw upon theoretical AI work on qualitative modelling (Brown & deKleer, 1985; Davis, 1983; deKleer, 1985; Forbus, 1985; Kuipers, 1985; Weld, 1983; Williams, 1985).

We chose mental models as the knowledge structures that we would try and impart to students for several reasons. Firstly, as Brown and deKleer (1985) have argued, such models can embody concepts and laws, can generate causal accounts, and can enable problem solving in a wide range of contexts. For example, the same mental model can be used to make predictions about the behavior of different circuits, to troubleshoot circuits, and to design circuits. This is in contrast with, for example, troubleshooting knowledge in the form of symptom-fix associations which is non-causal, context specific, and is, therefore, of limited use in helping students to understand how circuits work. A further reason for selecting mental models as the knowledge form is that in addition to being efficient and powerful knowledge structures for students to possess, they are also efficient and powerful knowledge structures upon which to base an intelligent learning environment. At any given point in the student's knowledge evolution, a single model can provide not only a model of how one wants the student to reason, but also can provide an interactive simulation of domain phenomena. Further, the simulation is capable, by simply reasoning out loud, of generating causal accounts for the behavior that the student is observing and creating. For instance, the student can close a switch and see a light turn on and, at the same time, hear an explanation for why the light turned on. Thus, we argue that mental models enable both the instructional system and the student to reason from general principles and to generate causal accounts of circuit behavior.

In this article, we focus on the design of an intelligent learning environment that is based upon qualitative, behavioral models of circuit operation. We view the role of instruction as developing in students a progression of increasingly sophisticated

mental models for reasoning about circuit behavior. We argue that these models should initially be qualitative and able to generate qualitative, causal accounts of the sequences of changes in circuit states that occur during the operation of a circuit.

In addition, we claim that the form of qualitative models employed should facilitate learning alternative conceptualizations of how circuits work. The concepts and reasoning processes employed in qualitative models should, for example, be compatible with quantitative models of circuit behavior and with functional accounts of system operation. This is important not only for facilitating the learning of multiple conceptualizations, but also for reasoning using multiple conceptualizations in the course of solving problems.

1.2.1. The Importance of Qualitative Reasoning

When novices and experts reason about physical domains, their approach to solving problems has something in common: Both employ primarily qualitative reasoning. Experts reason qualitatively about the phenomena before they resort to quantitative formalizations (Chi et al., 1981; Larkin et al., 1980), whereas, novices are only capable of qualitative, and often incorrect, reasoning (White, in preparation). If, however, one looks at less naive novices, such as people who have had one or two years of physics instruction, their reasoning is primarily quantitative and involves searching for equations that contain the givens in the problem (Chi et al., 1981; Larkin et al., 1980). This discrepancy is due, in part, to the emphasis placed, in most physics instruction, on learning quantitative methods and on solving quantitative problems. Experts, like beginning novices, make extensive use of qualitative reasoning. In the domain of electricity, for example, deKleer (1985) observes that, "an engineer does not perform a quantitative analysis unless he first understands the circuit at a qualitative level (p.275)".

We therefore argue that students should initially be exposed to qualitative, causal reasoning in order (1) to make connections with their naive intuitive models of physical phenomena, and (2) to enable them to acquire this important problem solving skill that evidence has shown they lack. Quantitative reasoning should only be introduced after students have been given a qualitative, intuitive conception of the domain, and the form of quantitative reasoning then taught should be a logical extension of the qualitative reasoning they have acquired. Further, the form of qualitative, causal reasoning should build upon students' naive but accurate intuitions and thus help to override their naive but inaccurate intuitions. In this regard, it

should be compatible with reasoning employed in other physical domains, such as mechanics, about which students' may have knowledge and experience that can be drawn upon during learning. It should also be compatible with students' intuitions about the causal nature of the world, such as: *changes in states have precipitating causes*.

This initial emphasis on qualitative thinking requires that students be given problems that necessitate qualitative reasoning for their solution. For instance, in the domain of electrical circuits, circuit design and troubleshooting problems can have this property. Problems of this type are thus useful in motivating the development of qualitative reasoning skills.

1.2.2. Causal Consistency

Conventionally, electrical theory is taught by presenting a series of laws which describe fundamental relations among voltage, current, and resistance in a circuit (e.g., Ohm's law, and Kirchhoff's voltage and current laws). The laws are presented as algebraic equations, which can be manipulated as to form (e.g., $I=V/R$, $V=IR$, $R=V/I$). Instruction then focuses on how to apply those equations in their various forms to the analysis of problems involving circuits of varying degrees of complexity, and the resulting constraints on voltages and currents in a circuit are used to develop quantitative solutions for the unknown quantities in the problem. (cf., Riley, 1984). Note that by using such constraint-based reasoning, the causal relations among voltage, current, and resistance are not made explicit, and the implicit causal model is actually inconsistent. Thus, at times the current flowing through a fixed resistance is viewed as determining the voltage, and at other times applied voltages are viewed as determining the current through a resistance.

It is also the case that qualitative theories are not necessarily consistent about the basic causal relations between voltage, current, and resistance. For example, deKleer's EQUAL (1985) infers that an increase in current out of a node causes a decrease in the voltage at that node (using the Kirchhoff current law or what is termed the KCL heuristic). At other times, an increase in voltage across a component causes the current through the component to increase (Ohm's law). Thus, the qualitative reasoning makes inferences about the effects of changes in current on voltage, and it also allows inferences about the effects of changes in voltage on current flow.

Our view is that mental models should be consistent in the assumed direction of causality among resistance, voltage, and current. In particular, current through a component, as Steinberg (1983) has argued, is determined by the voltage or electric force applied to the component. Voltages applied to a component within a circuit are, in turn, determined by resistances within the circuit. Viewing electric force as causing current flow also allows one to explain electrical phenomena that cannot be explained by current flow alone (for example, the behavior of capacitors; see Steinberg, 1983).

With electrical forces viewed as the causal agent, to understand a circuit's behavior, one needs to understand how changes in the conductivity (resistance) of circuit components alter the distribution of voltages applied to components within the circuit. Thus, our models employ a qualitative rule relating resistance to voltage (the $R \rightarrow V$ rule), and a qualitative version of Kirchhoff's voltage law. For example, the $R \rightarrow V$ rule states that a decrease in resistance of a component causes a decrease in voltage across the component (except if the component is directly connected to a voltage source). It further states that if the resistance of a component is zero (such as a switch when closed), there is no voltage drop across the component. To propagate the effects of that change in voltage, the underlying concept employed is one of physical systems attaining states of equilibrium. The instantiation of that general concept in the domain of electrical circuits is Kirchhoff's voltage law, which states that the electrical forces (voltages) around any loop within a circuit must balance one another, that is, sum to zero. For example, if a switch is closed, then any series of resistive components connected in parallel with the switch can have no voltage drop across them. In analyzing a circuit, one reasons using rules such as these to determine the distribution of charges within the circuit after a change in the state of a device has occurred, and the effects of those changes on the states of other devices within the circuit. Qualitative reasoning is thus based initially upon a subset of the constraints available in quantitative circuit theory, chosen for their causal consistency.

Simulating circuit behavior through the use of such qualitative models will reveal the sequence of device state changes that occur during the operation of the circuit and the reasons for those state changes. Thus, the student can see how changes in the state of a circuit precipitate other changes in the state of the circuit. For example, if a switch is suddenly closed, it may cause a capacitor to start discharging,

which in turn could cause a light to go on. The behavior of devices is causally determined by changes in other devices' states.

This sequence of device behaviors could equally well be constructed by a quantitative or qualitative model. However, qualitative models, in addition to being able to simulate the propagation of state changes within a circuit, can generate causal explanations for why the devices change state. For instance, they can describe how closing a switch completes a circuit, causing a voltage to be applied to a light bulb and thereby causing the light to go on. This is achieved by embedding within the simulation the basic electrical concepts of conductivity, resistance, and voltage drop, and by having the simulation utilize basic circuit principles relating to, for instance, how changes in conductivity and resistance can produce changes in voltage drops.

Understanding the causality of circuit behavior thus motivates the need to understand basic circuit concepts such as conductivity, resistance, and voltage and also basic circuit principles such as Kirchhoff's voltage law. These are non-trivial concepts and laws to master, so we take the approach of introducing them gradually, starting with simple circuits that can be reasoned about with simple forms of qualitative reasoning and progressing to more sophisticated circuits that require more sophisticated forms of qualitative reasoning for their analyses.

1.3. Learning as a Process of Model Transformation

A view of learning that follows from the mental models approach is that, in the process of acquiring an expert model, the student formulates a series of models each of which is adequate for some subset of problems (White and Frederiksen, 1985). These models are transformed into increasingly more adequate models in response to the demands of more complex problems undertaken by the student. Thus, the primary learning construct is one of model transformation. Transformation may involve the elaboration of model features, addition of features, generalization of features, differentiation among features, or even the construction of alternative models for representing the relations among and functions of devices within the domain. The representation of the learner's current knowledge state is a description of the model he or she currently has evolved. This representation, in turn, characterizes the types of problems that the learner can currently solve.

The form of mental model that we attempt to teach novices is not simply a subset of more sophisticated expert models. For example, students may learn to

reason about discrete changes in states of devices on the basis of the voltages that are present within a circuit. Later, they may learn to reason about incremental changes in voltages and how they influence device states. These alternative models represent different ways of reasoning about a circuit, which share some concepts but differ in others. Another example involves changes in a model's control structure. For instance, initially we focus on the behavior of a single device, such as a light bulb, in a circuit, and how one reasons about the behavior of the light bulb as it is effected by changes in the circuit. Later in the model progression, we focus on how one reasons forward from a change in the circuit, such as closing a switch, to the effect on all of the devices in the circuit.

Our work has focused on creating a progression of increasingly sophisticated models for reasoning about the behavior of electrical circuits. These models furnish learning objectives for different stages in instruction. They also represent different aspects of circuit behavior and are useful in their own right in reasoning about those particular aspects of a circuit's behavior. We define two dimensions on which models may vary: their order and their degree.

1.3.1. The Order of a Model

We distinguish models that reason on the basis of the mere presence or absence of resistance, voltage, current, which we call "zero order models", from those that reason on the basis of changes in resistance, voltage, or current, which we call "first order models". Zero order models can reason about binary states of devices and can answer questions of the form, "Is the light in this circuit on or off?" First order models on the other hand reason on the basis of qualitative (first-order) derivatives and can answer questions such as, "Is the light getting brighter?" Whereas, second-order models reason about the rate at which a variable is changing, such as, the rate at which the light is getting brighter. Each of these is distinguished from quantitative models that can answer questions of the form, "How much brighter is the light or how bright is the light?" All of these orders of model are thus useful for answering questions about circuit behavior of a particular sort¹.

¹Zero order models, for example, are sometimes taught as a basis for learning to troubleshoot electrical circuits (White and Frederiksen, 1985).

1.3.2. The Degree of a Model

Over the course of instruction, models developed increase in what we term their "degree of elaboration". This is determined by the number of qualitative rules used in propagating the effects of changes in state of circuit components on the behavior of other components. The initial qualitative models employ principles for determining voltages applied to components based upon only two basic rules: the $R \rightarrow V$ rule and a qualitative version of Kirchhoff's voltage law. These constraints are sufficient to understand and simulate the qualitative behavior of a large class of circuits, even though they are based upon only a subset of the available constraints of circuit theory. In subsequent models, a qualitative version of Ohm's law is introduced in order to relate changes in voltages across components to current through those components when their resistance is fixed. In later models, additional constraints are again introduced into the student's repertoire, namely qualitative rules based upon Kirchhoff's current law and a second constraint based upon Ohm's law, relating resistance to current. Finally, in the most sophisticated models a third constraint based upon Ohm's Law is introduced relating changes in current to changes that can be inferred to have occurred in voltage. In introducing this third constraint based upon Ohm's Law, we do not present the constraint as a causal relation between current and voltage (which would violate the causal consistency principle). Rather, we present the constraint as an example of backwards reasoning, where one infers the voltage change that precipitated a change in current.

The purpose of presenting models of increasing degrees of constraint elaboration is to teach students to reason flexibly using the full set of constraints available to them, however redundant they may be for the purposes of qualitative reasoning about simple circuit behavior. This is important if one seeks to then introduce quantitative reasoning as a natural extension of qualitative reasoning. When reasoning quantitatively, there are circuit problems that will require students to apply the full set of constraints available in circuit theory, and for students to reason "algebraically" in finding and applying multiple constraints.

1.4. An Overview of the Learning Environment

The instructional system we have built addresses the evolution of the unelaborated, zero-order model which is described in more detail in the following section. It enables students to learn how to reason qualitatively about device states using general circuit principles based upon the $R \rightarrow V$ rule and Kirchhoff's voltage

law. To this end, it develops the basic concepts of voltage and resistance and incorporates device models for devices commonly encountered in circuits. Students learn how to apply a knowledge of device models and circuit principles to simulate the operation of a circuit. They also learn strategies for troubleshooting that apply those principles.

The learning environment is based upon a decomposition of the knowledge domain into a sequence of increasingly sophisticated, qualitative models that correspond to a possible evolution of a learner's mental model. The progression of models constitutes a series of instructional goals for the student, namely, mastery of the model that is currently driving the simulation environment. Based upon the student's current mental model and a knowledge of possible model evolutions, students develop a model transformation goal (i.e., they choose which level of model they want to master next). The method of bringing about such a transformation is to instantiate it in problems for the student to work out. The instructional system presents to the student those problems that can be solved under the transformed model but not under the untransformed model. The students are thus motivated to revise their current model.

In order to facilitate this model transformation, the system can turn any problem into an example for the student by reasoning out loud while it solves the problem. As models become more sophisticated, they also become more verbose. The mechanism for pruning explanations is to focus the explanations on the difference between the transformed and the untransformed model. Reasoning of the transformed model that was present in the untransformed model either does not articulate itself or, if it is necessary to support the model increment, is presented in summary fashion.

Looking at the difference between the transformed model and the student's current model also helps to define what aspects of the problem solving process should be represented to the student. For instance, if students are learning about determining when there is or is not a voltage drop across a device, the system illustrates paths to voltage sources. However, later in the model progression, when it is assumed that students already know how to determine the presence of a voltage drop, the paths are no longer displayed.

Thus looking at the difference between the transformed and the untransformed model in the progression of models enables one to determine (1) what problems to present to the student, (2) what aspects of circuit behavior to articulate verbally, and

(3) what aspects of circuit behavior and of the problem solving process to visually display to the student.

Basing an instructional system on a progression of qualitative, causal models thus enable the system to:

1. Simulate circuit behavior. Each model is able to accurately simulate the behavior of a certain class of circuits. (The models can, in fact, simulate the behavior of any circuit, however, the simulation will not be accurate for all circuits.)
2. Model the students. The students are assumed to have the current model when they can correctly solve problems that the current model can solve but the previous model could not.
3. Tutor the students. By reasoning out loud, the models can generate qualitative, causal explanations for circuit behavior.

Each model can serve as a student model, a circuit simulator, and a tutor. All of the functions of the instructional system are thus performed, at a given point in the learning progression, by a single model.

The instructional system provides students with a problem-solving environment within which circuits can be built, tested, and modified. The student can select circuit components from a list of devices that includes batteries, resistors, switches, fuses, light bulbs, wires, transistors, and capacitors. The student then places the device on the screen in the desired location and indicates its connections to other devices. At the same time, as the student is constructing a circuit diagram on the screen, the system is constructing a qualitative model of the circuit. The student can request that the model "run" in order to obtain a visual representation of circuit behavior and, if desired, a verbal explanation for the circuit's behavior (presented via computer generated speech and in written form on the display screen). Thus, students can, for example, use a circuit editor to create circuits and experiment with them by changing the states of devices, inserting faults, and adding or deleting components.

The objective is to be able to have the simulation describe the behavior of a circuit in both verbal and graphic terms. There are graphic icons for each device in the circuit which are represented on the display screen with the appropriate connections. When a fault is introduced into the circuit, both the device model and the graphic representation of the device change appropriately. For instance, shorts

to ground alter the connectivity of the circuit, while opens alter the conductivity of the circuit. Similarly, when a device changes state, either as a result of an externally introduced change or as a result of the functioning of the circuit itself, the icon associated with that device can depict the new state. Furthermore, when search processes operate, they can leave a visible trace of the path they are currently pursuing so that, for example, when the simulation determines that there is a path with no resistance from a port of a device to ground, that path can be illustrated graphically on the display screen.

In addition to being able to construct and modify circuits, the system makes available a progression of problem sets for the student to solve based upon the progression of mental models. Circuit problems given to students include (1) making predictions about circuit behavior, and (2) troubleshooting or isolating faults within circuits. Corresponding to each of these two types of problems are two tutoring facilities: (1) the qualitative, causal model of electrical circuits that underlies the simulation and can illustrate principles for reasoning about circuits; and (2) an "expert" troubleshooter that can demonstrate a strategy for isolating faults within circuits and that incorporates the same type of reasoning as that involved in predicting circuit behavior. The troubleshooting expert operates in interaction with the circuit model as it diagnoses faults.

When solving problems, students can call upon these programs to explain reasoning about circuit operation or troubleshooting logic. The qualitative simulation utilizes a model appropriate for the student at a given stage in learning and thus can articulate its reasoning at an appropriate level of explanation. When circuits with faults are introduced, the circuit model can explain to students the operation of such circuits in either their faulted or unfaulted condition. Explanations of troubleshooting logic produced by the troubleshooting expert are also coordinated in level of complexity with the explanations of circuit behavior offered by the circuit simulation. The students can thus see a map of the learning space, as defined by the progression of circuit behavior and troubleshooting models, and can utilize this map to select problem sets.

By using these tools provided by the learning environment, the students can manage their own learning. For instance, they can choose to create their own problems using the circuit editor, and/or they can attempt problem sets and sequences of problem sets defined by the model progression. Further, they can ask to

see the behavior of a circuit simulated and can ask to hear explanations generated by the resident qualitative model. All of these learning tools are enabled by the qualitative model that is driving the learning environment at a given point in time and by the model progressions.

2. Qualitative Causal Models of Circuit Behavior

2.1. The Instructional Need for Zero Order Models

The pioneering work of deKleer (1979) and others (in Bobrow (Ed.), 1985) has shown how models can be developed that enable a computer to reason qualitatively about a physical domain. Further, these researchers have demonstrated that such models can be adequate to solve a large class of problems (e.g., deKleer, 1985). Our work on the design of qualitative models for instructional purposes has focused on creating models that (1) enable decompositions of sophisticated models into simpler models that can, nonetheless, accurately simulate the behavior of some class of circuits, and (2) enable the causality of circuit behaviors for the simpler models to be clear and at the same time compatible with that for more sophisticated models.

DeKleer (1985, p. 208) argues that: "Most circuits are designed to deal with changing inputs or loads. For example, ... digital circuits must switch their internal states as applied signals change The purpose of these kinds of circuits is best understood by examining how they respond to change." DeKleer's behavioral circuit model reasons in terms of qualitative derivatives obtained from qualitative versions of the constraint equations ("confluences") used in quantitative circuit analysis. These enable it to analyze the effects of changing inputs on circuit behavior.

The difficulty with utilizing such a model, at least at the initial stage of instruction, is that novices typically do not have a concept of voltage or resistance, let alone a conception of changes in voltages or resistance (Collins, 1985; Cohen et al., 1983). For example, as part of a trial of our instructional system, we interviewed seven high school students who had studied physics as part of a middle school science course, but who had not taken a high school physics course. They all initially exhibited serious misconceptions about circuit behaviors. For example, when asked to describe the behavior of the light in the circuit shown in Figure 1 as the switches are opened and closed, only one of the seven students had a concept of a circuit. The other students predicted that the bulb would light if only one of the switches were closed. A typical remark was the following, "If one of the switches on the left is closed, the light will light. It does not matter whether the switches on the right are open or closed." Further, they said, "if you close both switches on the left, the light will be twice as bright as if you close only one of them". In addition to this lack of a basic circuit concept, all seven of the students predicted that when you close the

switch in Figure 2, the light would still light -- the statement that the switch was not resistive when closed did not matter. In fact, five of the students stated that they did not know what was meant by the term "not resistive". They thus had no conception of how a non-resistive path in a circuit could affect circuit behavior.




Figure 1.

Novices such as these, who do not have accurate models of when a voltage is applied to a device in a circuit, could not possibly understand what is meant by a change in voltage across a device. Thus, we argue that students should initially be taught a progression of zero order, qualitative models that reason about gross aspects of circuit behavior. This type of model can accurately simulate the behavior of a large class of circuits, and can be utilized to introduce fundamental ideas about circuit behavior.






Figure 3.



(a)



(b)

Figure 2.

The knowledge embedded in the zero order models has been shown to be the type of knowledge that even college physics students lack (Cohen et al., 1983), and is also crucial knowledge for successful troubleshooting. For example, consider an elementary form of troubleshooting such as trying to locate an open in the circuit shown in Figure 3. Imagine that a test light is inserted into the middle of the circuit as shown in the figure. In order to make an inference about whether the open is in the part of the circuit in series with the test light or the part in parallel with it, one needs to know that if switch #1 were open, the light would not be on even if the circuit had no fault. Similarly, one needs to understand that if switch #2 were closed, the test light would not be on even if the circuit were unfaulted. Thus, even for performing the most elementary type of electrical troubleshooting, one needs a "zero order understanding" of circuit behavior.

Once basic aspects of circuit behavior have been understood, students can then progress to analyzing more subtle aspects of circuit behavior. For example, they can learn to determine how increasing the resistance in a branch of a circuit increases and decreases voltages within the circuit. Such an analysis requires a more sophisticated form of qualitative reasoning that utilizes qualitative derivatives. The qualitative model used in tutoring the students can no longer simply reason about whether or not there is a voltage applied to a device, rather, it must determine whether the voltage is increasing or decreasing. This type of analysis is necessary when analyzing, for instance, the occurrence of feedback within a circuit. Thus the progression of qualitative models must evolve to incorporate "first order reasoning", that is, reasoning about qualitative derivatives.

The first order models utilize many of the features of the zero order models and will be described in more detail later in the paper. This section describes the design and operation of the zero order models.

2.2. The Zero Order Models

The progression of zero order behavioral models incorporate knowledge of the topological structure of the circuit, the behavior of the devices within the circuit, and basic electrical principles relating to the distribution of voltages within the circuit. These principles enable the model to reason about the effects of changes in the conductivity of circuit components. The instructional system also includes a progression of general troubleshooting algorithms for localizing faults within a circuit

described in the next section. These algorithms utilize the behavioral models as part of their problem solving process. Both the behavioral models and troubleshooting algorithms can articulate their thinking, both visually and verbally, when simulating the behavior of a given circuit or when troubleshooting.

Circuit topology. The topology of the circuit is represented by the set of devices included in the circuit, together with the set of interconnections between designated ports of those devices. Thus, each instantiation of a device type within a circuit includes a table containing, for each of its ports, the electrical node to which it is connected.

Device models. The behavioral models contain device models for devices typically found in circuits. The devices modelled are batteries, switches, resistors, bulbs, diodes, fuses, capacitors, transistors, test lights, and wires (wires are explicitly introduced as devices). Device models include rules for determining a device's state, based upon the circuit environment of the device. For example, if there is a voltage drop across the two ports of a light bulb, the light bulb will be in the "on" state; otherwise it is in the "off" state. When a device's state changes, the device model activates additional rules which reevaluate a set of variables associated with the device. These variables include (1) the conductivity of the device (is it purely conductive, conductive but resistive, or nonconductive), and (2) whether or not the device is a source of voltage. For example, when a capacitor is in the charged state, it is nonconductive and a source of voltage. Finally, the device models include fault states, which include rules for altering the device variables to make them consistent with a particular fault, and which override the normal states for the device. For example, when a light bulb is faulted "open", it becomes non-conductive and its state will be "off". Some illustrations of device models² are given below.

²The devices are modelled as ideal components. Thus, for example, the battery is modelled as purely conductive because an ideal battery has no resistance, even though, real world batteries are resistive.

BatteryStates: Charged or Discharged

If the battery is discharged and if it has a voltage applied to it, then it becomes charged; otherwise it remains discharged.

If the battery is charged and if there is a path with no resistive elements across the battery, then it becomes discharged; otherwise it remains charged.

Internal Conductivity: Purely-ConductiveVoltage Source:

If the battery is charged, then it is a source of voltage; otherwise it is not.

Fault Example: Permanently Discharged

If the fault is permanently discharged, then set its status as a voltage source to negative.

For relatively complex devices such as capacitors, it is unrealistic to expect students at the outset to acquire the most sophisticated device models. Students are therefore introduced to a progression of increasingly sophisticated and adequate models for such devices.³ The initial capacitor model is illustrated below. The conditions for the rules that determine device states are written in such a way that only one of them can be true at a given point in time and they are evaluated in parallel, so that, on a given evaluation, only one of the rules will be executed.

³The initial capacitor model only incorporates two discrete states: charged and discharged. One limitation of such a capacitor model is that it does not take into account the non-steady states of charging and discharging. Furthermore, a capacitor is not just "charged", rather it is "charged to a given voltage". So, for example, if it is being charged by a small battery, it becomes charged to a low voltage, whereas, if it is being charged by a large inductor, it becomes charged to a high voltage. The consequence is that when a capacitor is charged to a given voltage, it is conductive-resistive to voltage sources higher than that voltage and is non-conductive to lower voltage sources. Thus the internal conductivity and resistance of the capacitor, which can affect the behavior of other devices in the circuit, can only be determined by knowing the level to which the capacitor is charged. For circuits with only one voltage source and for certain circuits with multiple voltage sources, circuit behavior can be accurately simulated without making this distinction. However, more complex circuits require the distinction to be made and thus learning about capacitors can motivate the need to understand more complex aspects of circuit behavior. They also can be used to introduce the limits of qualitative models and motivate the need for quantitative models. For example, consider a case where there are two low level batteries in series. The model now needs a rule saying that two voltage sources in series add together, but, what is LOW + LOW? Even further, what is LOW + HIGH? This illustrates a fundamental limitation of models that utilize category scales.

Capacitor

State: Charged or Discharged.

If it has a voltage applied to it, then its state is charged.

If it does not have a voltage applied to it and if its state is discharged, then it remains discharged.

If it does not have a voltage applied to it and if its state is charged and if it has a conductive path across it, then its state becomes discharged.

If it does not have a voltage applied to it and if its state is charged and if it does not have a conductive path across it, then its state remains charged.

Internal Conductivity:

If it is charged then it is non-conductive.

If it is discharged then it is purely conductive.

Voltage Source:

If it is charged, then it is a source of voltage.

If it is discharged, then it is not a source of voltage.

Fault Example: Internally Shorted

If the capacitor is internally shorted, then set its internal conductivity to purely conductive and its status as a source of voltage to negative.

When a particular device, such as a light bulb, is employed within a particular circuit, a data table is created for the specific instantiation of that device in that circuit. This table is used to record (1) the present state of the device, (2) whether it is presently a voltage source, (3) its internal conductivity (what possible internal conductive paths exist among its ports and whether they are presently purely conductive, resistive, or nonconductive), (4) the device polarity, as well as (5) its connections to other devices in the circuit, and (6) its fault status.

When the student is performing a mental simulation of a particular circuit, the student must also keep track of this information. Device connections are already given by the circuit diagram and thus do not need to be included in the student's device data table. However, the rest of the information related to the state of the device and its polarity must be recorded, either above the device in the circuit diagram or in a device data table, as illustrated in Figure 7.

A mental model for a device in the form illustrated for batteries and capacitors, enables the student to determine the state of the device regardless of the circuit environment in which it is placed.⁴ Information related to the state of the device, such as its internal conductivity and whether or not it is a source of voltage, will in turn affect the behavior of other devices in the circuit. Such a device model will thus form the basis for understanding the causality of circuit behavior in terms of showing how a change in state of one device can produce a change in state of another device within the circuit. It does not, however, provide the student with a "complete" understanding of how a battery works or how a capacitor works. For example, the capacitor model cannot generate an explanation for why a capacitor becomes non-conductive when it is charged. One ultimately needs to introduce, in addition to behavioral models, physical models for devices.

Circuit Principles. When simulating a particular circuit, the only information that the qualitative simulation requires is information about the structure of the circuit, that is, the devices and their interconnections. All of the information about circuit behavior, as represented by a sequence of changes in device states, is inferred by the qualitative simulation as it reasons about the circuit. To reason about device polarity and state, the device models utilize general qualitative methods for circuit analysis. For instance, when attempting to evaluate their states, device models can call upon procedures to establish voltages within the circuit. In the case of the zero order models, these procedures determine, based upon the circuit topology and the states of devices, whether or not a device has a voltage applied to it.⁵ The most sophisticated zero order voltage rule is based on the concept that, for a device to have a voltage applied to it, it must occur in a circuit (loop) containing a voltage source and must not have any non-resistive paths in parallel with it within that circuit. More formally, the zero order voltage rule can be stated as:

If there is at least one conductive path to the negative side of a voltage source from one port of the device (a return path), and if there is a conductive path from another port of the device to the positive side of that

⁴It should be noted that the behavior of the device will be accurate within the limits of the adequacy of the device model. Thus for complex circuits, a more sophisticated capacitor model may be required, as discussed later in the paper.

⁵In the case of the first order models, these procedures reason about whether the voltage drop across a device is increasing or decreasing as a result of changes in its resistance and the resistance of other devices in the circuit.

voltage source (a feed path), with no non-resistive path branching from any point on that "feed" path to any point on any "return" path, then, the device has a voltage applied to that pair of ports.⁶

Changes in a circuit, such as closing a switch, can alter in a dramatic way, the conductivity of the circuit and thereby produce changes in whether or not a device has a voltage applied to it. To illustrate, when the switch is open in the circuit shown in Figure 2(a), the device model for the light bulb calls upon procedures for evaluating voltages in order to determine whether the light's state is on or off. The procedure finds a good feed path and a good return path and thus the light bulb will be on. When the switch is closed, as shown in Figure 2(b), the procedure finds a short from the feed to the return path and thus the light bulb will be off.

Causal explanations. Simply having the model articulate that when the switch is closed, the light will be off because there is a non-resistive path across it, is not a sufficient causal explanation for students who have no understanding of (1) what is meant by non-resistive, or (2) what affect such a path can have on circuit behavior. First of all, students need definitions for concepts such as voltage, resistance, current, device state, internal conductivity, series circuit, and parallel circuit. Further, they need a "deeper" causal explanation of the circuit's behavior. For instance, there are two alternate perspectives on the causality of circuit behavior -- a current flow perspective and a voltage drop perspective. To illustrate, the following are explanations that (1) a current flow model, and (2) a voltage drop model could give as to why the light is off when the switch is closed for the circuit shown in Figure 2.

(1) The current flow model could state:

In order for the bulb to light, current must flow through it. There is a device in parallel with the bulb, the switch. In parallel paths, the current is divided among the paths. More current flows through the path with the least resistance. If one of the paths has no resistance, all of the current will flow through it. Since the bulb has resistance and the switch does not, all of the current will flow through the switch. Since there is no current flow through the bulb, it will be off.

⁶By "voltage applied to a device", we mean the qualitative version of the open circuit (or Thevenin) voltage, that is, the voltage the device sees as it looks into the circuit. In the case of the zero order voltage rule, this is simply the presence or absence of voltage.

(2) Whereas, the voltage drop model could state:

In order for the bulb to light, there must be a voltage drop across it. There is a device in parallel with the bulb, the switch. Two devices in parallel have the same voltage drop across them. Voltage drop is directly proportional to resistance: If there is no resistance, there can be no voltage drop. Since the switch has no resistance, there is no voltage drop across the switch. Thus, there is no voltage drop across the light, so the light will be off.

One could be given even "deeper" accounts of the physics underlying circuit causality. For instance, the system could present physical models that attempt to explain why current flow and voltage drop are affected by resistance in terms of electrical fields and their propagation. However, for our present purposes, the system presents a causal account to the depth illustrated by the preceding model.

In explaining the behavior of the light in the preceding example, one could utilize either the voltage drop explanation or the current flow explanation, or both. Our view is that giving students both types of explanations, at least in the initial stages of learning about circuits, would be unnecessary and confusing. It would require students to construct two models for circuit behavior, and this would create a potential for them to become confused about circuit causality. However, later on students may learn to reason in either way about circuit behavior.

We therefore selected only one of the causal models. We chose the voltage drop explanation because current flows as a result of an electromotive force being applied to a circuit; because troubleshooting tasks typically are based upon reasoning about voltages and testing for them; and because research has shown that this is an important way of conceptualizing circuit behavior that even sophisticated students lack, as illustrated by the following quotation from Cohen, Eylon, and Ganiel (1983):

"Current is the primary concept used by students, whereas potential difference is regarded as a consequence of current flow, and not as its cause. Consequently students often use $V=IR$ incorrectly. A battery is regarded as a source of constant current. The concepts of emf and internal resistance are not well understood. Students have difficulties in analyzing the effect which a change in one component has on the rest of the circuit."

In addition, reasoning about how circuits divide voltage is a major component of our first order models. These models reason about changes in resistances and voltages within a circuit, using a qualitative form of Kirchhoff's voltage law. Thus

getting students to reason in terms of voltages is compatible with the type of reasoning that will be required later on in the evolution of the students' models.

Topological search. The rules that embody circuit principles, such as the zero order voltage rule, utilize topological search processes that are needed, for example, to determine whether a device has a conductive path to a source of voltage. The search processes utilize the information maintained by the device data tables concerning the devices' circuit connections, polarity, internal conductivity, and whether or not they serve as voltage sources. The topological search processes can locate conductive paths within the circuit. For example, they can find all conductive paths from one port of a device to another port of the same device, or to a port of another device. They can also check to see if the paths are resistive or non-resistive. The students execute analogous search processes when tracing from one device to another, using the circuit diagram, in order to locate, for instance, a feed path for a device.

Establishing device polarities. The topological search processes are guided by polarities assigned to the ports of each device in the circuit. For example, when the light bulb in the circuit shown in Figure 4 is attempting to evaluate its state, it calls upon the voltage rule which invokes a search for, amongst other things, a conductive path to the positive side of the battery. This search immediately reaches a potential branching point: it could pursue the path starting with resistor R_3 and/or it could pursue the path starting with resistor R_2 . However, the search is reduced to following only the path starting with resistor R_2 , because the polarities of the connecting ports for the light bulb and resistor R_3 are both positive, and therefore, this path through resistor R_3 cannot lead to the positive side of the voltage source. The device polarities can thus be used to prune the topological searches.

Device polarities are established by a general, qualitative circuit orientation algorithm that reorients the circuit whenever a topological change in the circuit occurs or whenever a device alters its status as a source of voltage. The algorithm begins by identifying all electrical nodes⁷ in the circuit, and labelling them. Then, it recursively recognizes and removes all series and parallel subcircuits. Two

⁷Electrical nodes are points of connection between two or more resistive devices. Any non-resistive devices present are collapsed into a single electrical node.




Figure 4.

components that are connected together at both ends are recognized as a parallel subcircuit and are treated as a unit. Two components that are connected only to each other at one end are recognized as a series subcircuit and are also treated as a unit. The algorithm first brackets all parallel subcircuits as units and then, working with what are currently the highest level bracketed units, all series subcircuits. This process of alternately removing parallel and series subcircuits continues until there are no such subcircuits remaining. The algorithm constructs the innermost groupings first and proceeds in this way until the final grouping is reached, which in the case of series-parallel circuits is one that encompasses the entire circuit. The result is a hierarchical parsing of the circuit. The units are then assigned polarities in relation to the voltage source, starting at the outermost grouping and moving inwards. The side of a unit connected to the positive terminal of the battery is assigned a plus, and the other side a minus. Units contained within larger units are assigned the same polarities as those of the larger units which contain them.⁸

This circuit orientation algorithm was designed to be easy for students to learn and execute. However, in the initial progression of models, the complexity of circuits that students are exposed to is not sufficient to require teaching the algorithm. Determining the orientation of devices within the circuit is straightforward. Thus, we assume that students can identify device orientations within the initial progression of circuits, and therefore, the algorithm does not articulate its behavior and is not explicitly taught.

Control structure. The simulation of circuit operation is driven by changes in the states of the devices in the circuit. These changes are produced by (1) changes in states of other devices, such as a battery becoming discharged causing a light to go out; (2) external interventions, such as a person closing a switch, or a fault being introduced into the circuit; and (3) increments in time, such as a capacitor becoming discharged. Whenever a device changes state, its status as a voltage source is redetermined by the device model, along with its internal conductivity/resistance. Whenever any device's internal conductivity or status as a voltage source changes,

⁸This algorithm can identify indeterminacies in the assignment of polarities to a unit. For instance, if a unit has both feed and return paths from each of its ports then its orientation may not be determined. If all of these paths lead to the same voltage source, it is a bridge element in the circuit. If the paths lead to different voltage sources having different polarities, the orientation of the unit is also indeterminate.

then time stops incrementing within the simulation and all of the other devices in the circuit reevaluate their states. This allows any changes in conductivity or presence of voltage sources within the circuit to propagate their effects to the states of other devices. The circuit information used for this reevaluation is the set of device data tables existing at the initiation of the reevaluation (not those that are being created in the current reevaluation cycle). This is to avoid unwanted sequential dependencies in determining device states. If in the course of this reevaluation some additional devices change state, then the reevaluation process is repeated. This series of propagation cycles continues until the behavior of the circuit stabilizes and no further changes in device states have occurred. Time is then allowed to increment and the simulation continues. When any further changes in device internal conductivity or status as a voltage source occur, due either to the passage of time or to external intervention, time is again frozen and the propagation of state changes is allowed to commence once again.

A Sample Zero Order Circuit Simulation. As an illustration of how a zero order model reasons, consider a simulation of the behavior of the circuit illustrated in Figure 5:

Initially suppose that both switches are open, the light bulb is off, and the capacitor is discharged. Then, suppose that someone closes switch #1. This change in the internal conductivity of a device causes the other devices in the circuit to reevaluate their states. The capacitor remains discharged because switch #2 being open prevents it from having a good return path. The light bulb has good feed and return paths, so its state becomes on. Since, in the course of this reevaluation no device changed its conductivity, the reevaluation process terminates. Note that even though the light bulb changed state, its internal conductivity is always the same, so its change of state can have no effect on circuit behavior and thus does not trigger the reevaluation process.

Now, imagine that someone closes switch #2. This change in state produces a change in the conductivity of the switch and triggers the reevaluation process. The light bulb attempts to reevaluate its state and finds that its feed path is shorted out by the capacitor (which is purely-conductive because it is in the discharged state) and switch #2 (which is also purely-conductive because its state is closed), so its state becomes off. The capacitor attempts to reevaluate its state and finds that it has a good feed and return path, so its state becomes charged. This change in state causes it to reevaluate its internal conductivity, and to reevaluate whether it is a source of voltage. As a result of the capacitor becoming charged, it becomes non-conductive, and a source of voltage. This change in the internal conductivity of the capacitor causes the reevaluation process to trigger again. The light bulb reevaluates its state and finds that it has a good feed




Figure 5.

and return path (it is no longer shorted out by the capacitor because the capacitor is now charged and therefore non-conductive) and its state becomes on. This change in the light bulb's state has no effect on the light bulb's internal conductivity so the reevaluation process terminates.

Suppose that someone then opens switch #1. This changes the switches internal conductivity and therefore causes all other devices to reevaluate their states. The light bulb no longer has a good return path with respect to the battery. However, it has a good feed and return path to another source of voltage within the circuit, the capacitor (which is charged and therefore a source of voltage). The state of the light bulb will thus be on. The capacitor no longer has a good return path to a source of voltage and it has a conductive path across it, so its state becomes discharged and it becomes purely-conductive and is not a source of voltage. This change in the capacitors internal conductivity causes the light bulb to reevaluate its state. Since the capacitor is no longer a source of voltage, and since switch #1 is open thereby preventing a good return path to the battery, the light bulb concludes that its state is off. This change in state has no effect on the light bulb's internal conductivity so the reevaluation process terminates.

Notice that this relatively unsophisticated qualitative simulation has been able to simulate and explain some important aspects of this circuit's behavior. It demonstrates how when switch #2 is closed, it initially shorts out the bulb, and then, when the capacitor charges, it no longer shorts out the bulb. Further, it explains how when switch #1 is opened, the capacitor causes the light bulb to light initially, and then, when the capacitor becomes discharged, the light bulb goes out.

The evolution of the control structure. By control structure we mean the determination of what goal to pursue next when reasoning about the behavior of a circuit (what Anderson (1984) terms the "problem solving structure"). An example of control knowledge within the qualitative model is, "when one device changes its conductivity, all other devices in the circuit must reevaluate their states". The system makes such control knowledge clear to the student by simply reasoning out loud. For instance, it might state, "I am trying to determine whether this device has a voltage applied to it (i.e., it states a goal). In order to do that, I must search for a conductive path from one port of the device to a voltage source. Then, ... (i.e., it states the means for achieving its goal)." Thus the system articulates its goals and subgoals, as well as its means for achieving those goals. By so doing, the control structure of the simulation becomes apparent to the student.

One of the most impressive features of the type of qualitative, causal model described in this paper is its utility in helping to solve a wide range of circuit

problems. For example, the student can be asked to predict the state of a single device after a switch is closed, or to describe the behavior of the entire circuit as various switches are opened and closed, or to determine what faults are possible given the behavior of the circuit. Further, students can be asked to locate a faulty switch within a circuit, or to design a circuit such that when the switch is closed, the light in the circuit will be off. Performing this type of mental simulation of circuit behavior is instrumental in solving all of these types of problems.

For instance, even when the student is attempting to predict the behavior of a single device within a circuit such as a test light, it is often necessary to know the states of other devices within the circuit. If there are devices such as capacitors and transistors whose internal conductivity is state dependent, then their state must be determined in order to determine the state of the light bulb. Thus even for this simple type of problem, a mental simulation of the entire circuit is often necessary.




Figure 6.

The complexity of the control structure required for simulating circuit behavior varies with the type of circuit. For circuits that contain only devices like resistors and bulbs whose internal conductivity does not change when their states change, parallel reevaluation is not necessary. For example, consider the circuit shown in Figure 6. Suppose that someone closes the switch in the circuit. Whether light bulb #1 and light bulb #2 reevaluate their states in parallel or in a particular order makes no difference to the behavior of the circuit. The state of one light bulb has no effect on the state of the other light bulb since they both remain conductive-resistive no matter what their state (unless they are faulted open). By initially restricting attention to such circuits, one can begin the instructional sequence with models that

reason serially and then introduce the more complex parallel architecture only when students have already been introduced to basic circuit concepts such as conductivity, resistance, and voltage drop.

In fact, as long as there is only one device in the circuit whose internal conductivity changes when its state changes, serial reasoning can yield accurate predictions concerning circuit behavior as long as a prescribed order of device reevaluations is followed. To illustrate, when switch #2 was closed in the simulation previously presented for the circuit shown in Figure 5, either having the light bulb and the capacitor reevaluate their states in parallel, or having the light bulb reevaluate its state before the capacitor leads to a correct simulation of behavior for this circuit. However, suppose instead, that the capacitor had reevaluated its state first. It would have determined that it had a good feed and ground path and it would have become charged, non-conductive, and a source of voltage. The light bulb would have then have reevaluated its state and found itself to be on, whereas, it should have initially been off. One of the light bulb's state changes would therefore have been missed. Thus either parallel reevaluation or serial evaluation, with the device whose internal conductivity changes when its state changes being reevaluated last, can work for this type of circuit.

However, for circuits containing multiple devices, such as capacitors and transistors, whose internal conductivity changes when their state changes, parallel reevaluation of device states is crucial for accurately simulating the behavior of the circuit. One approach is for students to learn to simulate parallelism the way the computer model does. This is done by introducing a notational scheme that facilitates the parallel reevaluation of device states using device data tables. To elaborate, the circuit diagram provides information about device connections. Students then record device polarity information around each device as is done in Figure 4. Above each device the student must record the device's "data": its state, its internal conductivity, and whether it is presently a source of voltage. For serial evaluation of device states, updating this information is all that is required. However, parallel reevaluation requires keeping two sets of device data for each device. One set records the present data for the device and the other set records the reevaluated data. Students then learn that the device whose change precipitated the reevaluation does not get reevaluated, so that its data remain the same while other devices undergo reevaluation. The remaining devices use the present data of other devices in the

circuit, not the reevaluated data, when reevaluating their own state. If one wants to record the behavior of the circuit as sequences of state changes that occur within the circuit, one simply makes a table of device data. Data for each device are recorded in the table after each reevaluation or time increment. By circling the devices that change state in each column of the table, the sequence of state changes for the circuit can become clear as illustrated in Figures 7 and 8.

This process would become lengthy for large circuits. A second approach may prove to be more efficient, and more direct in terms of the causality of circuit behavior. That is to use a zero order form of Kirchhoff's voltage law to immediately propagate the effects of a change in conductivity of a device on voltages applied across other devices in the circuit. Then, when other devices' states are reevaluated, it will already have been established whether or not there is a voltage being applied to each of them. The qualitative version of Kirchhoff's voltage law states that, in any loop containing a voltage source, there will be voltages applied across any devices in the loop provided there are no shorts across the device within the loop. Whenever a device changes its internal conductivity or its status as a voltage source, the voltages applied to other devices in the circuit are reassessed using the voltage law applied to those devices that are in direct loops with that particular device. Thus, feed and return paths do not have to be (redundantly) determined for each device in the loop. In addition, since changes in voltages applied to other devices within the circuit can be inferred, only devices with a change in voltage applied to them need reevaluate their states. If in the course of reevaluation the internal conductivity or status as a voltage source of any device changes, then the voltage law is triggered again, and so on. In this sequence of reevaluations, the model is similar in control structure to that of its more inefficient predecessor.

Time dependent behaviors - one limitation of qualitative models. A major limitation on time dependent behaviors for qualitative models is that the sequencing of events happens in ordinal, not interval, time. That is, subject to the limitations mentioned below, the state changes happen in the correct order, but the length of time between events is indeterminant. For instance, in the preceding example of the circuit illustrated in Figure 5, when switch #2 was closed, how long did the light stay off before coming on again? Was it an instant or a relatively long time? The model has no way of knowing. Further, the simulation implied that first the capacitor charged and then the light came on. This is not quite accurate since, although the

DEVICE	VARIABLES	INITIAL CONDITIONS	TIME →								
			1	2	3	4	5	6	7	8	
SWITCH 1	STATE CONDUCTIVITY VOLTAGE SOURCE?	OPEN N-C NO	CLOSED P-C* NO	CLOSED P-C NO	CLOSED P-C NO	CLOSED P-C NO	CLOSED P-C NO	CLOSED P-C NO	OPEN N-C* NO	OPEN N-C NO	OPEN N-C NO
SWITCH 2	STATE CONDUCTIVITY VOLTAGE SOURCE?	OPEN N-C NO	OPEN N-C NO	OPEN N-C NO	CLOSED P-C* NO	CLOSED P-C NO	CLOSED P-C NO	CLOSED P-C NO	CLOSED P-C NO	CLOSED P-C NO	CLOSED P-C NO
CAPACITOR	STATE CONDUCTIVITY VOLTAGE SOURCE?	DISCHARGED P-C NO	DISCHARGED P-C NO	DISCHARGED P-C NO	DISCHARGED P-C NO	CHARGED N-C* YES*	CHARGED N-C YES	CHARGED N-C YES	CHARGED N-C YES	DISCHARGED P-C* NO*	DISCHARGED P-C NO
LIGHT BULB	STATE CONDUCTIVITY VOLTAGE SOURCE?	OFF C-R NO	OFF C-R NO	ON C-R NO	ON C-R NO	OFF C-R NO	ON C-R NO	ON C-R NO	ON C-R NO	ON C-R NO	OFF C-R NO
RESISTOR	STATE CONDUCTIVITY VOLTAGE SOURCE?	COLD C-R NO	COLD C-R NO	HOT C-R NO	HOT C-R NO	HOT C-R NO	HOT C-R NO	HOT C-R NO	HOT C-R NO	COLD C-R NO	COLD C-R NO
BATTERY	STATE CONDUCTIVITY VOLTAGE SOURCE?	CHARGED P-C YES	CHARGED P-C YES	CHARGED P-C YES	CHARGED P-C YES	CHARGED P-C YES	CHARGED P-C YES	CHARGED P-C YES	CHARGED P-C YES	CHARGED P-C YES	CHARGED P-C YES

LEGEND:
 Conductivity = Non-conductive (N-C), purely-conductive (P-C), or conductive-resistive (C-R)

- = Device changed state
- * = Device's internal conductivity or status as a voltage source changed which triggers a reevaluation cycle
- ||| = Circuit behavior stabilizes, i.e., end of a sequence of reevaluation cycles

Figure 7. Table of device data tables.







Figure 8. Emergent causalities from table of device data tables.

capacitor would start charging before the bulb would light, it would most likely not be charged to battery voltage before the bulb lit. The limitation has arisen from the attempt to model only steady states within the circuit. This view works for some circuits but not others. In order to accurately simulate the behavior of a larger class of circuits, for example, the capacitor model could be refined to reason about non-steady states as well as steady states. The refined capacitor model would then reason about charging and discharging, as well as the charged and discharged states. It would incorporate rules of the form: if the capacitor is charged and if there is not a voltage applied to it and if it has a conductive path across it, then its state becomes discharging.

However, there are still limitations to such a model's ability to simulate these time dependent behaviors. For instance, even though the simulation can now determine when the capacitor starts discharging, it has no precise means of determining when the capacitor will be discharged. For some circuits, this limitation is crucial. In such cases all the qualitative model can do is to articulate the range of possible behaviors for the circuit. So that, for instance, if the capacitor becomes discharged at a certain point with respect to the behavior of the other devices, the circuit will exhibit one behavior, whereas, if it becomes discharged at another point, the circuit will exhibit a different behavior. The student, or system, must then use knowledge about the purpose of the circuit or quantitative models to determine what is the likely behavior for this particular circuit.

No function in structure. We sought models that would be robust in permitting faults to be introduced or circuits to be modified without requiring a new model for each perturbation in the circuit. By utilizing context free models for devices along with circuit principles for evaluating voltages, we have been able to construct qualitative circuit models that simulate the behavior of a large class of circuits in both faulted and unfaulted states.

The device models are prototypical and behave appropriately (within the limits discussed) no matter what circuit they are placed into. The only circuit-specific information that is required is the set of device interconnections, that is, information about the structure of the particular circuit. Similarly, the circuit principles embody general laws of circuit behavior that work (again within the limits discussed) for all circuits. Thus we are in keeping with deKleer and Brown's (1985) no function in structure principle.

Creating knowledge structures with this property is important in enabling the system's qualitative model to simulate and generate explanations for the behavior of any circuit that the students choose to construct (within the limits discussed). It is also an important property for the students' mental models in that their knowledge will then be in a general form that enables them to understand and predict the behavior of any circuit.

Locality. However, unlike deKleer (1985), our device models do not reason locally. Rather, they typically determine the integrity of feed and return paths (that is, carry out a loop analysis) in order to determine their states. This is a consequence of the causal analysis that we are trying to teach. This feature also enables our models, unlike deKleer's (1985), to avoid making assumptions about the integrity of the circuit and, therefore, to avoid running into contradictions in their reasoning processes.

Causality. However, a potentially serious difficulty introduced by violating deKleer's locality principle is that it requires the introduction of parallelism into the more sophisticated models. Thus more than one device can change state on a given cycle, which could obscure the causal relationships between changes in device states. For instance, suppose that two devices, A and B, change state on a cycle. Then, on the next cycle, another device, C, changes state. From merely observing the state changes, one could not infer whether A or B or both A and B caused C to change state.

The causality could be recovered by imagining that only A had changed state or only B, and then determining whether C still changed state. However, for large circuits, this would require a lot of unnecessary processing. A simpler method for recovering the causality is made possible by the type of reasoning that the zero order models employ. The topological search that is used to determine if a device has a voltage applied to it, can compare the trace of the circuit on one cycle, with the trace of the circuit on the next cycle, in order to determine, for instance, why a given device now has a good feed path, whereas, on the previous cycle it did not. This type of comparison is easy for a student to make, particularly for the small circuits that are utilized to teach basic circuit concepts.

3. Troubleshooting

The problem of troubleshooting a circuit requires students to reason "on their feet" about circuit behavior, and is potentially a very powerful instructional task. Conventionally, however, troubleshooting is preceded by instruction on circuit theory, rather than used as a vehicle for teaching models of circuit behavior. By decomposing troubleshooting strategies along lines that are parallel to those used in the construction of zero order qualitative models, troubleshooting problems can be incorporated within the general instructional sequence.

3.1. The troubleshooting algorithms

The progression of troubleshooting algorithms is based upon a qualitative approach taken by an expert whom we have studied. This expert not only utilizes this approach in actual diagnostic work, but also teaches the technique to students in a technical high school. The method he uses is based upon the fundamental idea of a circuit, and is similar to that of the zero-order models (which was motivated in part by the approach of this expert): For a device to "operate" (e.g., for a test light to light or a capacitor to charge), it must have voltage applied to it. When such an electrical potential exists, a current will flow through the device (provided it is conductive), causing it in some cases to change its state. In order for there to be an electrical potential, there must be a source of voltage. Further, there must be conductive paths leading from each port of the device to, respectively, the positive and negative sides of a voltage source. In a series circuit, one source of faults is the occurrence of opens within either of these paths, which will prevent current from flowing with a resulting effect on the device's state. Another source of faults is the presence of shorts to ground, which introduce non-resistive parallel paths into the circuit. If these shorts occur between the device and the ungrounded side of the voltage source, they will prevent current from flowing through the device. Opens and shorts to ground are types of faults that the troubleshooting algorithm is designed to diagnose.

The goal of the troubleshooting algorithms is to divide the circuit into two parts and then to infer which portion of the circuit contains the fault. The troubleshooting logic is then recursively applied to the faulty segment until the fault has been localized. This is accomplished using the following strategy: First, the circuit is logically divided into two parts by inserting a test light into the circuit between a test point near the center of the circuit and the grounded (negative) side of the voltage

source. Second, the circuit simulation is run to determine the correct state of the test light in a circuit that is not faulted. Third, that state is compared with the actual test light behavior, and inferences are made about possible faults that are consistent with the findings. The logic used depends upon whether or not the test light is supposed to be on, given an unfaulted circuit, and upon the actual behavior of the light in the presence of the fault. For instance, if the test light is supposed to be on and is not on, the fault could be in the part of the circuit in series with the test light, or it could be a short to ground in the part of the circuit in parallel with the test light (at a point before any resistance is encountered). Additional troubleshooting operations are then carried out to isolate the fault to either the portion of the circuit in series with the test light, or the one in parallel with the test light. To accomplish this, the expert detaches the latter portion of the circuit from the test point, and observes the effect. If the test light comes on, the fault can be isolated to the portion of the circuit in parallel with the test light. Namely, it was providing a non-resistive path from the feed path of the light to the ground. If the test light remains off, the problem must be an open or a short to ground in the portion of the circuit in series with the test light. When the fault has been isolated to within a portion of the circuit, the expert moves the test light to a new point within the faulty segment of the circuit and reapplies the troubleshooting logic. This process is repeated until the fault is located.

The troubleshooting logic as described here is restricted to series circuits. However, additional principles allow it to be extended to parallel circuits and to series-parallel circuits. In instruction, the troubleshooting algorithm presented to students increases progressively in complexity. The sequence of troubleshooting algorithms is coordinated with the progression of behavioral circuit models that the students acquire.

3.2. Facilitating troubleshooting

The faults that can be introduced into a circuit, in the current version of the instructional system, are shorts to ground, and opens. The device model has rules for determining how each fault will alter its data. For instance, shorts to ground change the circuit connections for that device whereas opens may change the conductivity of the device. Both types of fault can change the state of the device. When a particular fault is removed from the circuit, the device data are returned to their unfaulted values and the circuit simulation proceeds on that basis. The particular faults that

are introduced at any stage in instruction are chosen to be consistent with the partial model of circuit behavior currently implemented in the simulation. Thus, for instance, shorts to ground are not introduced until students have learned about non-resistive parallel paths.

To facilitate troubleshooting, a test light can be introduced into a circuit. In addition, ports of any device can be disconnected (for example, one can choose to disconnect the portion of a circuit in parallel with a test light). These troubleshooting operations alter the circuit connections and the model simulates the behavior accordingly. The availability of these facilities enables students to troubleshoot for themselves. If at any time they want assistance, they can call upon the system "expert" to demonstrate its techniques on the circuit they are working on and explain its logic. In fact, if they choose, they can plant a fault into a circuit themselves and have the expert demonstrate how it would proceed to isolate it.

4. Model Evolutions

Most of the work on qualitative modelling within the AI community has been concerned with developing relatively sophisticated models for simulating the behavior of physical phenomena (e.g., see Bobrow (Ed.), 1985). The work deals with qualitative derivatives (Brown and deKleer, 1985) and qualitative calculi (Forbus, 1985). This is understandable since these researchers are interested in creating intelligent, not naive, machines. However, our interest is in instruction and in possible transitions from novice to expert behavior. We have developed, therefore, simpler zero order qualitative models for the novice that are easy to learn, that capture important circuit concepts and laws, and that are extendible to more sophisticated ways of reasoning about circuit behavior. Moreover, for purposes of instruction, the zero order models themselves have been decomposed into a succession of models of increasing complexity, each extending the range of electrical circuit problems that can be understood. In tutoring, more advanced qualitative models can be introduced when the students have mastered the concepts and principles contained in the earlier models.

The learning theory on which we base our tutoring system assumes that, in a learning environment in which students are continually solving problems, students will develop mental models on which they can base their problem solving. Initially, these are naive models that have been developed informally as a result of prior experience with electrical systems. The tutoring system seeks to provide a means for students to evolve their models into progressively more sophisticated representations of electrical circuit behavior, and it seeks to do this by presenting problems and offering explanations that motivate particular transformations in the students' models.

In this section of the paper, we outline the types of model transformations that are possible at any given stage of learning. We go on to articulate the factors that must be taken into account when attempting to determine an appropriate path for a particular student to take through the space of possible model progressions. Finally, we describe one curriculum that we implemented in order to teach basic electricity and troubleshooting to high school students.

4.1. Types of Model Evolution

If one takes the view that students "learn how to learn", then students may have numerous learning processes and strategies which they have evolved for themselves, and since these processes are learned, there will be individual differences in the set of learning processes and strategies that a given learner will possess. This view has several important implications. Firstly, it suggests that it would be inordinately difficult to model the particular learning processes whereby a particular learner will transform one model into another since the set of such processes that a learner may possess will be large and will vary from learner to learner. Secondly, as a consequence of the existence of different learning strategies, one needs to allow for different learners to pursue different paths through the space of possible model progressions.

While there are individual differences in the processes by which models may be transformed, it is useful to characterize some of the products of model transformation -- the ways in which models can evolve. These can be broadly classified into modifications of a model's knowledge (declarative and procedural knowledge within the model), and modifications of its structure (the form of knowledge representation within the model).

Evolution of knowledge. In learning, a model's knowledge may be augmented by refining, generalizing, or differentiating an existing concept, by adding a new concept, or by integrating several existing concepts within some larger conceptual framework. Each of these transformations represents a type of knowledge evolution and a possible pedagogical goal for the student to pursue. The following are examples of each of these ways in which a student could choose to progress.

1. Knowledge acquisition -- The student acquires a new concept or law or problem solving skill. For example, many novices, as we have discussed, do not have the basic concept of a circuit.
2. Knowledge refinement -- The student refines an existing concept. For example, students may want to refine their understanding of voltage drop, by noting, for instance, that in parallel circuits, a device only needs for one of its feed and return paths to be "good" in order to operate.
3. Knowledge generalization -- The student learns how an existing concept applies in a wide range of contexts. For example, students could learn that their concept of resistance associated with a resistor can also be applied to a light bulb.
4. Knowledge differentiation -- The student learns about the differences

between two concepts. For instance, students may want to learn how their concept of voltage drop differs from the voltage measured from any point in the circuit to ground.

5. Knowledge integration -- The student integrates two concepts. For example, students may need to synthesize their understanding of non-resistive paths with their conception of voltage drop.

Students may differ in the type of evolution they prefer at different stages of learning. One student may prefer, for example, to generalize first and differentiate later, whereas another may prefer to differentiate first and generalize later.

Evolution of structure. Possible transformations for a mental model are not limited to the preceding changes in the model's knowledge base. A mental model can also change in its form. These structural transformations alter the way in which knowledge is represented and applied. For example, one can choose to include within each specific device model rules for altering device states and variables when the device is faulty. Alternatively, one could choose to keep the rules for making such modifications separate from the device models - as general procedures that operate on device models and infer the effects of a fault on the device's state. Another example involves changes in the control structure of the model. One such transformation was given in the previous section, where we described how propagations of changes in voltages could be evaluated: (1) on a device by device basis, by reasoning backward whenever a device's state is reevaluated; or (2) by propagating forward the changes in voltages that occur whenever any device changes its state. Another transformation in control structure is the shift from serial reevaluations of device states to parallel reevaluations. Structural changes such as these in a mental model may pose particular difficulties for the learner.

4.2. The Problem of Modifiability

If one's theory of learning involves a concept of model transformations and the view that at each stage in learning the student must develop a runnable model on which to base problem solving, then a primary consideration in designing such evolutionary families of models must be their modifiability. Models must be developed with a view towards facilitating their progressive upgrading in response to new problem demands. In this regard, a worthwhile analogy can be made with the programmer's problem of developing code that is maintainable and modifiable. Concepts such as modularity, inheritance, goal decomposition, and the like have

evolved within computer science to serve these needs, and they all have their application to the development of progressions of mental models that can be easily learned. For example, to facilitate learning, all devices of a given type should have a common model and that model should be independent of the circuit context in which the device occurs (modularity), and all device models should have a common form (inheritance). Thus, when the concept of a fault state of a device is introduced, it can be easily generalized to other devices.

In considering the learnability of a particular model progression one must consider not only the concepts and reasoning skills that must be acquired, but also the types of "programming" changes that the new reasoning skill would require to the student's mental model. These changes might involve refinements, rewrites, deletions, or additions of device models or general circuit principles, as well as changes in the model's control structure. Each of these types of change poses its own particular problems for the learner who is attempting to modify his or her current model in an appropriate fashion.

Refining Knowledge. The simplest kind of change that a model transformation could produce is the refinement of a rule or a procedure. For example, as the students' understanding of voltage increases in sophistication, the rule for determining whether or not there is a voltage applied to a device gets refined in a gradual progression. The basic rule remains, qualifiers just get added to it.

Rewriting Knowledge. Another type of change occurs when students are introduced to a new way of conceptualizing some aspect of circuit behavior, as is the case, for instance, when one goes from a zeroth order transistor model to a first order transistor model. The transistor model remains, but some of the rules get rewritten as opposed to simply refined.

Deleting Knowledge. This type of change requires students to completely erase, or at least to no longer access, some aspect of their mental model. An illustration is when students utilize their zero order mental model of circuit behavior to acquire a first order model. Certain rules of the zero order model no longer apply and should not be incorporated into a first order model.

Among these transformations, complete rewrites of aspects of the model are likely to be more difficult for the student to achieve than refinements or deletions. On the

other hand, complete rewrites may sometimes be necessary in order to introduce material in an easily learnable form. For example, the zero order models enable basic circuit concepts to be acquired more easily than if one started with first order models. However, the limitations of a zero order model require the addition of a first order model, which builds upon the knowledge and structure of the zero order model but requires rewrites of many of the zero order rules.

Adding Knowledge. Consider next the problem of adding knowledge, as when the student learns something entirely new. An example is when transistors are introduced as devices for the first time. In this case, the concept of a device model existed before, but the particular prototype for a transistor did not. Adding knowledge is a potentially complex model transformation because one has to decide where to place the knowledge. If the instructional approach involved teaching independent condition-action rules, this would not be an issue. However, in the case of mental models it can be a crucial issue. For instance, does one place a new rule or concept in the prototypical device model so that all other device models inherit the knowledge, or does it belong in the device model for, say, capacitors? Even further, possibly the rule is a general principle of circuit behavior and does not belong in a device model at all. Considerations of where a particular piece of knowledge should be embedded in the students' mental model are an important in determining the learnability and useability of the model.

Revising Control Knowledge. A final example of a model transformation that may cause difficulty in learning is the alteration of the control knowledge that students employ to manage their reasoning about circuit behavior. For example, at the beginning of instruction, students may be asked to reason about the behavior of one device within a circuit, such as, a light bulb. For such problems, the student's model needs only to activate one device model plus the basic circuit principles that are needed to determine the behavior of the device within the circuit. However, later in the progression, students are asked to reason about multiple devices within a circuit. Initially this can be done serially, but as soon as devices such as capacitors and transistors are introduced, it must be done "in parallel". Thus, the form of the student's model gets more complex in that control procedures that were initially unnecessary, or at least were very simple, must now increase in complexity. Similar kinds of control complexities are introduced when students go from troubleshooting just opens, or just shorts to ground, to attempting to locate either type of fault

within a circuit. Moreover, for purposes of economy in reasoning, students may wish to retain multiple control structures so that they can reason using the simpler, serial control structure when a problem allows it (such as, for example, when determining the expected state of a test light if a circuit were unfaulted). There is thus the added complexity of learning the contexts in which a particular control structure is applicable.

The problems of modifiability can be particularly complex when one is trying to impart knowledge in the form of a mental model rather than as, for example, a collection of independent condition-action rules, such as a set of symptom-fault-fix associations that many experts use in troubleshooting. For instance, the complexity of control knowledge does not become an issue if the knowledge is in the form of independent condition-action rules such as symptom-fix associations.

Finally, the type of model transformation can affect the ease or difficulty a student has in using the model to reason about circuits. For instance, changes that increase the complexity of the model's control structure could make the model not only more difficult to learn but more difficult to use as well. Creating learnable model progressions must take into account not only their modifiability, but also how easily they can be put into practice in solving problems.

4.3. The Path of Model Evolution

The selection of appropriate model transformation goals during learning involves a consideration of not only students' learning styles and the difficulty of the transformation, but also the purposes for which they are learning about circuit behavior. If, for example, students are learning for the purposes of acquiring skill in troubleshooting circuits, the path through the model progression space that is most appropriate may be different from that for students who have the goal of designing circuits.

At any point in learning, different types of model transformations are possible that increase the sophistication of the model's reasoning in different ways. A particular path through the space of possible model progressions embodies a possible transition from novice to expert status. Our ultimate goal is to create a space of possible model progressions and to add facilities to the learning environment that will help students to select a path through this space based upon their own pedagogical styles and goals.

Within the present project, we have focused on learning to troubleshoot and have constrained the network of possible model evolutions to a linear progression, i.e., a curriculum. We have created and tried out one particular curriculum which had the objective of teaching troubleshooting for opens and shorts to ground in series-parallel circuits.

4.3.1. A Zero Order Curriculum

The progression of zero order models that we selected, in conjunction with the progression of troubleshooting algorithms, captures a possible transition from novice to expert status. The progression thus defines a curriculum for a student. The path through the space of possible model progressions was constrained by (1) teaching circuit concepts and laws needed to enable troubleshooting, and (2) teaching them in an order that would permit students to engage in troubleshooting as early in the progression as possible while still making the principles and causality of circuit behavior clear. By starting with simple zero order qualitative models, the curriculum introduces the fundamental idea of a circuit and of a voltage drop. It then progresses to ideas about resistive and nonresistive paths in parallel circuits. Finally it teaches the troubleshooting of opens and shorts to ground within series and parallel circuits. Within this progression, more than one type of transition is typically incorporated in a step. For example, students may be acquiring a concept of resistance at the same time as they are revising their understanding of when a device has a voltage applied to it. The two changes are integrated in that the need to understand voltage motivates the need to understand resistance.

Voltage, conductivity, and the fundamental idea of a circuit. The zero order curriculum we have implemented, starts by teaching the fundamental idea of electrical potential and its ability to alter a device's state. In order to understand how an electrical potential can be developed across a device in a circuit, the idea of conductive and non-conductive paths to a voltage source are introduced. Series circuits, such as the one shown in Figure 9, containing only a battery, light bulb, wires, and switches are utilized. The fault of open is introduced as a means for creating a non-conductive path. The control structure required of the students' model is kept simple by asking them to make predictions about the behavior of a single light bulb when a switch is opened and closed, or when a wire is faulted open.

Reasoning about more than one device changing state. In this model transition, students learn to generalize the concepts related to electrical potential and device




Figure 9.




Figure 10.

states to cases where they must reason about more than one device changing state in the circuit. They learn that when there is a circuit, voltage is applied to all devices within the circuit. Reasoning about the behavior of more than one device increases the complexity of the control structure of the students' model. Since the circuits being presented at this stage contain only devices such as light bulbs and resistors whose internal conductivity never changes, a serial evaluation of device states is all that is necessary. Thus this model transformation entails only a slight increase in complexity over the control structure of the previous model.

The concept of a common ground. This transition generalizes the concept of a circuit to incorporate a common ground -- a purely conductive path which, when devices are connected to it, serves as a connection between the devices. Circuit problems of the type that the students have already learned to reason about are presented. The only difference is that this time, devices are connected to a common ground instead of directly to the battery, as shown in Figure 10.

Alternative feed and ground paths. In this transition, students are asked to reason about the behavior of light bulbs in circuits, such as that illustrated in Figure 1, which potentially supply multiple feed and ground paths. This type of reasoning is necessary for troubleshooting because when a test light is inserted into a circuit, it could have multiple feed paths. In this transition, the students' concept of voltage must be refined to incorporate the fact that in a circuit with parallel paths, only one good feed and ground path are necessary for a device to have a voltage applied to it.

Shorts across a device. In this model transition, students are exposed to circuits where shorts immediately across a device can exist, and they must expand their circuit principles to account for the effects of such shorts. For example, in the circuit shown in Figure 2, there is a short across the light bulb when the switch is closed. Understanding this type of short is needed when troubleshooting since if there is a purely conductive path in parallel with a test light, the light will be off even if there is no fault in the circuit. In this model progression, students must differentiate their concept of a conductive path into conductive-resistive and purely conductive paths. Thus their concept of conductivity must be refined and this refinement must be integrated into their voltage rule -- the rule must now incorporate the fact that if there is a purely conductive path immediately across a device, then no voltage is applied to that device.

Purely conductive paths in parallel circuits. This model transition generalizes the concept of a purely conductive parallel path (a short) from being immediately across a device to being anywhere on the device's feed path to any point on a return path or even immediately to ground (a short to ground). The circuit principles used to infer when voltages are applied must be refined to incorporate the more sophisticated rule presented earlier in our discussion of circuit principles.

Troubleshooting an open in a series circuit. Students now possess an understanding of circuit behavior that is sufficient to support troubleshooting a series circuit containing a battery, wires, light bulbs, switches, and resistors. The simplest troubleshooting algorithm is thus introduced at this point. This subset of troubleshooting logic allows students to learn the basic troubleshooting heuristics of dividing the search space and making inferences about entire portions of the circuit. By limiting the fault to an open, both the conceptual and procedural aspects of troubleshooting are kept simple.

Troubleshooting shorts to ground in a series circuit. Students now have a basic knowledge of troubleshooting heuristics and an understanding of circuit behavior sufficient to support locating opens and shorts to ground. Learning how to locate shorts to ground is made easier by considering a short to ground as the only possible fault at this point in the learning sequence. In this transition, students thus generalize the troubleshooting heuristics of dividing the search space and making inferences about entire portions of the circuit to situations in which they must locate shorts to ground in a series circuit.

Locating opens or shorts to ground in a series circuit. Finally, students are given problems to motivate an integration of their troubleshooting model for finding opens with that for finding shorts to ground, since in real troubleshooting situations, they will not know which fault is present in the circuit.

Additional model evolutions will include increasing the domain of circuits that the student can troubleshoot to include series-parallel circuits, and increasing the repertoire of device models to include capacitors and devices such as diodes and transistors that have polarities associated with them.

4.4. Further Model Evolutions

We envision further model evolutions within the tutoring environment aimed at developing alternate conceptualizations of circuit behavior. These included: (1) first order models that allow one to reason about changes in resistance and voltage and how they propagate within a circuit; (2) increasing the degree of elaboration of models, such as through extensions of the underlying framework of the analysis to include forward reasoning about the effects of voltage and resistance on current, and backward reasoning about how changes in current have been precipitated by changes in voltage; and (3) quantitative circuit analysis based upon the qualitative constraints on voltage, resistance, and current that have previously been presented in their qualitative forms. Certain of these alternative models will be discussed in more depth in a later section of the paper.

The model evolutions discussed in this section have been with respect to changes in the students' zero order model of circuit behavior. The same principles apply when considering more dramatic evolutions in the students' understanding of how circuits work. For instance, just because students are adept at looking at circuit diagrams and predicting the behavior of circuits does not mean that they have a "deep" understanding of electrical circuits. They may be completely unable to describe the functionality of circuits - the purpose of a circuit as a whole and the role that subsets of devices play in achieving that purpose. Also, they may understand nothing about the physics of device and circuit functioning. Further, they may only be able to reason at a qualitative level and thus be unable to formalize their understanding by constructing a quantitative model of circuit behavior. Thus we claim that in order to attain a "deep understanding" of how a circuit works, students must evolve such alternative conceptualizations of circuit phenomena that exist in conjunction with their zero order model of circuit behavior.

5. The Learning Environment

The learning environment consists of an interactive simulation driven by a qualitative model, and a troubleshooting expert. The system is capable of generating runnable, qualitative, causal models for any circuit that the student or instructional designer might create. Thus students can, for example, use a circuit editor to create circuits and experiment with them by changing the states of devices, inserting faults, and adding or deleting components. They can also ask the system to illustrate and explain the behavior of the circuit, or to demonstrate how to locate a fault within the circuit. In addition, there is a curriculum organized around a progression of models which serves to define classes of problems and facilitate the generation of explanations. Students can thus attempt to acquire an understanding of how circuits work in a more structured way by solving problems designed to induce particular transformations in their understanding and by hearing explanations for how to solve those problems. They can also use the circuit editor to modify and experiment with these circuits presented to them by the system.

This section of the paper describes problem types and learning strategies that are enabled by the learning environment. It then goes on to discuss the findings of instructional trials of the system in terms of the learning strategies actually employed by students, and the effects of the learning environment on students' ability to reason about circuits. Implications of these findings for future revisions of the system are discussed.

5.1. Problem Types

One of the most interesting features of an intelligent learning environment based upon qualitative models is the range of problem types supportable by this architecture.

Predicting device behavior. The student is presented with a circuit and is asked to predict the behavior of a device or devices in the circuit. Similarly, for certain model transformations, the student or computer can insert test lights into various points in the circuit and the student is asked to predict the behavior of the test light. In addition, the student or the computer may change the state of some device (e.g. open or close a switch) or fault a device and the student is again asked to predict the behavior of the light. The system gives the student feedback concerning whether his or her prediction was correct or incorrect. Also, the student is given the

option of having the system give its explanation as to what the state of the device or devices is and why.

Enumerating all possible faults consistent with circuit behavior. The student is presented with a circuit containing a fault unknown to the student and a test light inserted into the circuit between a particular point and ground. The student is then asked to enumerate all possible faults that are consistent with the behavior of the test light. When the student has finished selecting all faults that he or she believes would produce the given test light behavior, the student is given feedback concerning the correctness of her or his selections as well as any omissions he or she has made. At any point in the problem solving process, the student can request to have an unfaulted circuit to work with, complete with the test light, and can experiment with introducing faults into the circuit and observing the behavior of the test light. As in the prediction problems, the student can also request that the system give an explanation of why the test light is in that state. In addition, the student can request to hear the system solve the problem which it can do by hypothesizing all possible faults and running the qualitative simulation to see what test light behavior results. In doing so, it considers five possible fault types and locations, (1) an open in the part of the circuit in series with the test light, (2) a short to ground in the part of the circuit in series with the test light, (3) an open in the part of the circuit in parallel with the test light, (4) a short to ground in the part of the circuit in parallel with the test light before a point where resistance is encountered, and (5) a short to ground in the part of the circuit in parallel with the test light after a point where resistance is encountered. If the test light behavior for any of these fault possibilities is consistent with the given behavior of the test light, then that fault is included in the set of possible faults that are consistent with that test light behavior.

Troubleshooting problems. The computer selects a fault for a given circuit and the student is asked to determine the location and type of fault. The student can insert a test light between any point in the circuit and ground. The student can also disconnect devices from one another. After each such operation that the student performs, the student is asked two questions: (1) given the current behavior of the test light, which portion of the circuit, that in parallel or that in series with the test light (or both), could contain (i) an open or (ii) a short to ground, and (2) can you determine the specific location of the fault yet, and if so, where is it? When the student has located the fault, the computer gives the student feedback as to whether

the choice is right or wrong. At any point in the troubleshooting process, the student can request to hear how the computer would troubleshoot the circuit.

Circuit design and modification problems. The student is asked to, using the circuit construction kit, create a circuit that achieves a particular purpose. For example, when learning about non-resistive parallel paths, the student could be asked to create a circuit such that when the switch in the circuit is closed, the light bulb goes from on to off. A simpler form of problem is a circuit modification problem where students are asked to alter a circuit so that its behavior changes. For instance, they could be asked to insert a switch into the circuit so that when the switch is closed, the light will go off. At any point in the circuit construction process, the student can request to see and hear an explanation for the behavior of the circuit that they have created. They must then decide, based upon the behavior of the circuit, whether their design is correct or incorrect.

Problems in model design, modification, and debugging. In addition to creating and troubleshooting circuits, the learning environment could allow the student to create and debug qualitative models for circuit behavior (the system currently does not have this facility). All of the types of problems that apply to circuit behavior (troubleshooting, prediction, etc.), apply to mental model behavior as well. Thus students could be asked, for example, to locate the buggy device model, or an erroneous circuit principle, or faulty control knowledge contained in a given model (e.g., Brown and Burton, 1978; Brown & Van Lehn, 1980; Richer & Clancey, 1985). In order to determine this, students could present the model with circuits and observe how it simulates them. Further, they could inspect the model by looking at, for instance, the rules within its device models.

5.2. Problem Selection

With respect to the different types of problems, predictive problems were chosen as the initial method of inducing a model evolution because they require only the running of the students' mental model for their solution. Enumerating possible faults consistent with circuit behavior is the next type of problem presented. Solving this type of problem requires running a mental model, for each of the five possible fault types and locations mentioned previously, to see what circuit behavior results.

The troubleshooting and circuit design problem types require knowledge that goes beyond a mental model of circuit behavior. For instance, troubleshooting

problems require in addition a knowledge of troubleshooting heuristics. This type of problem was presented to students after they had a model of circuit behavior of sufficient complexity to support troubleshooting. Circuit design and modification problems require a knowledge of circuit functionality as well as circuit behavior. We are currently working on extending the learning environment to incorporate functional models of circuits (a model of the organization and operation of the circuit derived from (1) its overall purpose, and (2) an analysis of the operations required to achieve that purpose, and the necessary relationships between those operations). Because this class of model is not currently implemented, we did not include this type of problem in our curriculum.

Problems involving qualitative model design and troubleshooting are potentially a most interesting method for facilitating model evolution. The current implementation of the system does not have facilities for allowing students to create and debug mental models so we were unable to utilize this problem type.

Defining Problem Sets

With respect to the predictive problems, the current linear progression of partial models defines sets of problems that are deemed appropriate for the students at different stages in learning. A first pass at defining problems sets came from grouping problems that can be solved by a given model but cannot be solved by the previous model in the sequence. The theory is that by giving students problems in this group, i.e., problems that are just beyond their level of competence, that students would be motivated to revise their model. This model revision would be facilitated because it would require only a small change to their model in an environment where feedback and explanations are available to help them to understand the model transformation. Students should thus be motivated and able to transform their model into the next model in the sequence.

In addition to problems requiring the transformed model, some problems were interspersed from the earlier set. In some cases these problems provided negative exemplars of a concept. If students were learning, for example, that a short from a point on a device's feed path to a point on its ground path prevented the device from having a voltage drop across it, and if all the problems were cases of this sort (i.e., where there was always a short from feed to ground), then students would never see negative instances (i.e., cases where there was no short). As Bruner, Goodnow, and

Austin (1956) have argued and illustrated, negative instances of a case are very important to learning. Providing some problems from the previous set often served this function.

Presenting problems that the transformed model can handle but the current model cannot turned out not to be a sufficiently refined method of selecting problems in that not all problems in that category are of equal instructional value. There are additional subcategories of problems that can be classified by their pedagogical effects:

Illustrate a Prototypical Case. Certain problems have the property of making the model difference as clear as possible and, if presented first to the students, have a high likelihood of causing the correct model transformation. They illustrate the difference in the simplest possible instance with no distracting other possible causes for differences in circuit behavior.

Illustrate an Extreme Case. If, however, students have difficulty inferring a model transformation from a prototypical case, it is often useful to present a problem which embodies an extreme case. For instance, introducing a short immediately around a device, instead from any point on its feed path to any point on its ground path, often helps make the concept of a short easier to understand.

Produce Incorrect Model Transformations. Certain problems in the category could, if presented first in the problem solving sequence, induce wrong transformations to the student's model. For instance, the circuit illustrated in Figure 11 causes some students to infer that the light lights because one of the feed paths had no resistance in it. This is incorrect and once inferred needs to be corrected.

Fix Incorrect Transformations. If, however, the student has made an incorrect inference, certain problems in the category are particularly good at undoing the incorrect inference. They address the particular erroneous inference by focusing students attention on what is wrong with that inference. For example, students who have erroneously inferred that the light in the circuit shown in Figure 11 lights because the one of the feed paths has no resistance, can be given the same circuit problem, only this time with all of the feed paths resistive.

The initial sequence of problems in a set is crucial to facilitating a correct

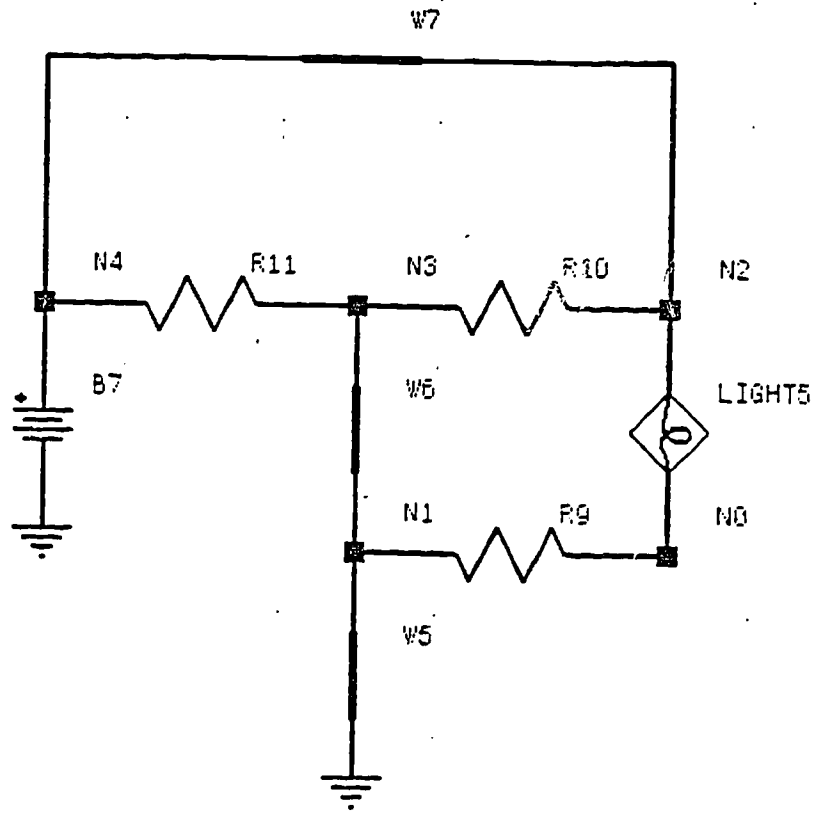


Figure 11.

model transformation. After the correct model transformation has been induced, the remaining problems serve the function of giving students practice in utilizing their new mental model for circuit behavior. Our approach to the design of problem sets, as elaborated in the next subsection, was to focus on creating, for the initial problems, prototypical and extreme case problems and to avoid the creation of "bug inducing" problems. We thereby avoided having to create "bug fixing" problems. The remaining problems in the set were simply derived from the more general category of problems that can be solved by the new model but not by its predecessor, interspersed with problems from the previous set. This latter section of the problem set thus included problems that, if presented in the initial stages of model transformation, could have caused the induction of buggy models.

5.3. Design Philosophy

We have focused on creating a progression of models that makes a gradual transition from naivity to expertise. To facilitate this transition we:

1. motivated learning via problem solving and appropriate problem selections;
2. emphasized qualitative, causal analysis that builds upon novices existing intuitive knowledge; and
3. generated explanations that make the causality of circuit behavior, as derived from basic concepts and principles, as clear as possible.

The assumption is that (i) by giving the students problems that (i) present a manageable cognitive challenge to the student, that is, problems that they could solve with a small revision to their mental model, and (ii) are inherently interesting, such as troubleshooting or circuit prediction problems, and (2) by presenting students with examples of model reasoning via verbal and visual descriptions of circuit behavior, that the student's model, at any stage in the learning process, will be transformed to match that of the system's. If students make incorrect model transformations, we assume that the fault is in the model progression (which affects problems selection and explanation generation) and revise the model progression. That is, we do not assume that wrong inferences are a necessary consequence of the learning process and, therefore, we do not attempt to diagnose and treat wrong model transformations. The hypothesis is that if the model progression and problem sets are designed appropriately, one does not get incorrect model transformations.

Selecting an appropriate progression for a given student is non-trivial. It

requires a decomposition of the domain knowledge and reasoning skills that builds gradually on the learners prior knowledge. It also requires knowing the learning strategies a given student can utilize to transform his or her model. In addition, it requires understanding the purpose for which the student is learning about the domain -- a progression that may be relevant and of interest to one student's purposes, such as learning to troubleshoot, may be inappropriate for another.

Given that students may manage their own learning ineffectively and select inappropriate model progressions, diagnosing and treating wrong models may become necessary. Such a diagnosis could be achieved by constructing buggy models and, based upon students answers to problems, identifying their buggy models, and adjusting the selection of problems and explanations accordingly (as do, for example, Anderson et al., 1984; Brown & Burton, 1978; Goldstein, 1982; Johnson & Soloway, 1984; Soloway et al., 1983; Reiser et al., 1985). However, our initial focus is on developing good model progressions, problem sets, and explanations.

5.4. Learning Strategies

Basing the system on a progression of qualitative models makes it possible for students to have considerable freedom in determining the way they interact with the learning environment. Students can choose whether to advance to new levels in the progression or to review earlier problems. They can attempt to solve problems on their own or can request the tutor to give demonstrations and explanations. They can use a circuit editor to alter existing problems or create new circuits, and can add or remove faults from a circuit they have been given or one they have created. The system supports this wide range of activities by being able to simulate the behavior of a circuit that is constructed and by providing explanations of its operation. Finally, the concept of a progression of models allows the student to understand what electrical knowledge has been mastered and what remains to be learned.

This architecture for an intelligent learning environment permits great flexibility in the students' choice of an instructional strategy. Particular strategies that can be followed include the following:

Open-ended exploration. Students can construct circuits, explore their behavior (by changing the states of devices, inserting faults, and adding or deleting components), and request explanations for the observed behaviors. Students can thus create their own problems and experiment with circuits. The system thereby permits an open-ended exploratory learning strategy.

Problem-driven learning. In addition, the progression of models enables the system to present students with a sequence of problem solving situations that motivate the need for developing particular transformations of their models of circuit behavior. In solving new problems, the students attempt to transform their models of circuit behavior in concordance with the evolution of the system's models. The focus is on having students solve problems on their own, without providing them first with explanations for how to solve them. Only when they run into difficulty, do they request explanations of circuit behavior.

Example-driven learning. Alternatively, students can be presented with tutorial demonstrations for solving example problems by simply asking the system to reason out loud about a given circuit using its present, qualitative, causal model. Students can thus hear explanations of how to solve each type of problem in the series, followed by opportunities to solve similar problems. Since the focus is on presenting examples together with explanations prior to practice in problem solving, we term this learning strategy "example-driven".

Student directed learning. The classification of problems created by the progression of models provides facilities students can use in pursuing instructional goals of their own choosing. Problems can be classified on the basis of the concepts and laws required for their solution, and on the instructional purpose served by the problem.⁹ This enables students to pursue goals such as acquiring a new concept or differentiating two concepts. The students can thus make their own decisions about what problems to solve and even about what learning strategy to employ.

5.5. Instructional Effectiveness

The learning environment was tried out on seven high school students who had had no formal instruction in circuit theory. The students were initially shown a demonstration of how to use the various facilities of the system and then given the opportunity to use those facilities to control the functions of the system while learning. Thus, the students could browse through the topics in the curriculum (as embedded in the progression of qualitative models), select problem sets to try, decide for themselves when to go on to a new topic (i.e., a more sophisticated model), and

⁹In the current implementation, the classification of problems is in terms of the linear progression of models.

could use the circuit editor to alter a given circuit. In addition, whenever they so desired, they could ask the associated circuit model to simulate the circuit's behavior and to articulate its reasoning. They could also point to any device in the circuit and ask for an explanation as to why the device was in a particular state. Similarly, they could ask the troubleshooting algorithm resident at that state in the progression to demonstrate and explain how it would locate a fault in the circuit.

The students were given, as a pretest, a set of circuit problems and asked to explain the behavior of each circuit as the states of devices within it were manipulated. As described earlier in the paper, initially the students exhibited serious misconceptions about circuit behavior and lacked key electrical concepts. Further, none of them had had any experience with troubleshooting. The students then spent from five to six days, an hour a day, working with the system. The students were then given the same eight circuit problems they had attempted in the pretest and asked to explain the behavior of the circuit or to troubleshoot.

All of the students were remarkably conservative in the use of the system. Typically, they did a large proportion of the problems in a given set, even though after the first few problems, they were getting them all correct. The reason they often gave was that they were afraid of missing a "tricky" problem near the end of the set -- "something I don't understand might be lurking in there". They rarely skipped a topic and went through them in the linear order of the curriculum. They only occasionally experimented with a circuit by, for instance, flipping switches or disconnecting parts. Instead they primarily employed the learning strategy of going to a new topic (as embedded in the next qualitative model in the progression), trying a problem, getting it wrong, asking for an explanation, and then solving the rest of the problems (usually correctly). Occasionally, when the new topic was particularly novel (e.g., troubleshooting for the first time), they would request a demonstration/explanation before attempting a problem.

There are numerous possible explanations for why the students employed this "conservative" learning strategy. The fact that the system presented a curriculum to the students in a sense implied to them that its designers thought it was a good idea to progress through the models in this linear order. If, instead, they had been presented with a network of increasingly sophisticated models, they would have been forced to decide on their own path through the model space and problem sets, and their behavior may have been quite different. Further, the students' conception of

how one learns, as derived from their school environment, is primarily that of following a curriculum by hearing explanations then doing problems. So the fact that they employed this learning strategy when using the instructional system may simply be an instantiation of their school model of learning. The implication is that when we extend the learning environment, we should explicitly teach alternative learning strategies. A final possible explanation for the students conservative behavior is that, when interacting with the system, they were always being observed. This may have inhibited their exploratory behavior and increased their desire to "do the right thing" by focusing on getting correct answers rather than discovering things for themselves.

After five hours of working within the learning environment on an individual basis, all seven of the students were able to make accurate predictions about circuit behavior and could troubleshoot for opens and shorts to ground in series circuits. They went from getting all of the pretest questions incorrect to getting all eight correct on the posttest (with the exception of one student who got two of the questions on the pretest correct since he already had the basic concept of a circuit).

The most impressive results were reflected in the students' troubleshooting behaviors. Several of the students modified the troubleshooting algorithm that the system demonstrated to make it more efficient. In other words, they understood circuit behavior and the troubleshooting heuristics (such as divide the search space) well enough to make modifications. Another noteworthy aspect of the students' troubleshooting performances was that, when they made erroneous inferences, they were usually able to recover. For instance, they would reach a contradiction and recognize that one of the inferences they had made earlier was premature. Finally, on the posttest, students were given a troubleshooting problem of a type they had not seen before, and all of the students were able to get the correct answer (although a few of them got the correct answer even though they did not accurately generate all possible fault locations at each step).

Despite the apparent success of the learning environment, several deficiencies became apparent as the student worked with the system. For example, the sequencing of problems within a given set was crucial (which is not surprising). As we discussed earlier, all problems that can be solved by the present model but not by the previous model, are not of equal instructional value. In particular, the initial problems in a set should be selected so that they can be solved by the transformed model but not by some other erroneous model transformation.

With respect to the explanations generated by the models, there need to be more levels of explanation. For instance, often a student simply made a slip when making a prediction about circuit behavior. In asking for an explanation, they merely wanted to locate their slip and did not need extensive explanations of the circuit's behavior. Rather, they needed a summary trace of the model's reasoning. In contrast, there were other times when students wanted "deeper" explanations than the two levels of explanation currently available. For example, many of the students wanted to know why there is a voltage drop across a resistor but not across a wire. Simply being told that, *voltage drop is directly proportional to resistance -- if there is no resistance, there can be no voltage drop -- a wire has no resistance so there is no voltage drop across it*, was not sufficient to completely satisfy them. The learning environment needs to incorporate deeper causal models of circuit behavior such as a "pressure-flow" model: *it takes pressure (voltage) to make the current flow through a resistor, so the more resistive the resistor, the bigger the pressure drop (voltage drop) across it.*

Were the students' mental models in the form of the qualitative, causal models driving the learning environment? There is some evidence that they were. The instructional strategy was to tell students that whenever they go to a new topic, the computer will have a slightly more sophisticated model for predicting and explaining circuit behavior. The students were thus playing a "guess my model" game which, aside from any interest the students may have had in learning about electricity, was motivating in its own right. When the students were reasoning out loud on the posttest circuit prediction problems, their reasoning was usually identical to that of last qualitative, causal model embedded in the curriculum.

However, when it came to the troubleshooting problems, there were, as was alluded to earlier, some interesting differences between the students' strategies and that of the computer. Occasionally students made premature inferences about the location of the fault. This was due to a deficiency in the model progression and in the availability of problem types. There were not enough problems of the form: *identify all possible faults that are consistent with a given behavior of the test light inserted into a given circuit.* On the other hand, some of the students' troubleshooting strategies were different from the computers in a more positive sense. They had the goal of locating the fault as quickly as possible and thus recursively used the split-half technique. The computer troubleshooting "expert" uses the split-half technique initially and then does a serial search in the section of the circuit

known to contain the fault. This strategy was selected to avoid errors that occur when students attempt to remember the bounds on the part of the circuit that they have determined contains the fault. Indeed, this error of forgetting the bounds on the faulty portion of the circuit did occur in students who derived the recursive split-half strategy.

We think that these discrepancies between the computer's strategy and those derived by the students were not due to the inevitability of such differences, but rather, were due to the form of troubleshooting knowledge embedded in the learning environment. For example, the rule concerning where to insert the test light into the circuit should have been more flexible -- anywhere in the suspected faulty portion of the circuit is reasonable. This change would enable the students or the system to choose a point based upon considerations of efficiency, ease of insertion, or knowledge of likely fault locations. It would enable the students or the instructional system to generate a set of reasonable next test light locations that could be chosen at a given point. This decision in the troubleshooting process could then be based upon general principles, such as consider likely fault locations, as opposed to simply following a predetermined rule. This alternative, more general form of encoding this particular aspect of troubleshooting knowledge would enable the students to be more flexible and principled in their behavior, and would enable the system to provide the students with better feedback and explanations when they are in the process of troubleshooting.

To summarize, we argue that any difference between the students' mental models and those we were trying to teach were not due to the inevitability of bugs or misconceptions, but rather, were due to limitations of the learning environment. In other words, the cognitive theory underlying the learning environment needs to undergo further evolution. The derivation of erroneous mental models was due to a non-optimality in either the form of the knowledge we were trying to impart, or the progression of models, or the type of problem selected to induce a particular model transformation. Thus our future research will focus on developing further the theory underlying model forms, model transformations, problem types, and instructional strategies.

6. Multiple Alternative Conceptualizations

We will begin this section by analyzing a digital logic circuit using the zero order model. It will be seen that, while such an analysis can describe the behavior of such circuits at a level that is sufficient for understanding the causal sequence of device state changes, there are other fundamental questions concerning circuit operation that are not explained. These include understanding the purpose of components in the circuit that have no apparent function under the first order model, and accounting for the behavior of a circuit when there are quantitative changes in its input signal. First order qualitative models will be introduced for reasoning about the behavior of a circuit when quantitative changes in voltage must be explained, as, for example, when feedback is employed. These first order models reason about the first order derivatives of voltage, and resistance, rather than about their presence or absence. Finally, some of the limitations of qualitative models will be discussed, as well as the role quantitative models may play in supplementing an analysis based upon the zero and first order qualitative models.

6.1. Zero Order Qualitative Models

An application of zero order logic to a simple logic circuit is illustrated in Figure 12 (Horowitz and Hill, 1980, p. 86). The circuit is used in an automobile to control a buzzer. Whenever a person is seated in the drivers seat (causing switch S_3 to close) and either front door is ajar (either switch S_1 or S_2 is closed), the buzzer sounds. Otherwise, the buzzer is silent.

The behavior of the circuit can be derived by applying the zero order model. This produces (1) a sequence of state changes that occur in devices, and (2) explanations for the state changes in terms of the causal dependencies among devices in the operating circuit. A summary of the explanations produced by the model follows. Remember that the model evaluates states of devices in parallel. When any device changes its internal conductivity or status as a voltage source, all other devices reevaluate their states. If in generating explanations we have only those devices that change state during their reevaluations explain their behavior, a causal sequence of device state changes will be generated. In addition, when each device changes state, it provides an explanation of the cause of its change in state. In this example we will assume that initially the left door is open (S_1 is closed) and that the seat is occupied (S_3 is closed).

1. Transistor Q_1 attempts to evaluate its state. It finds a path from its

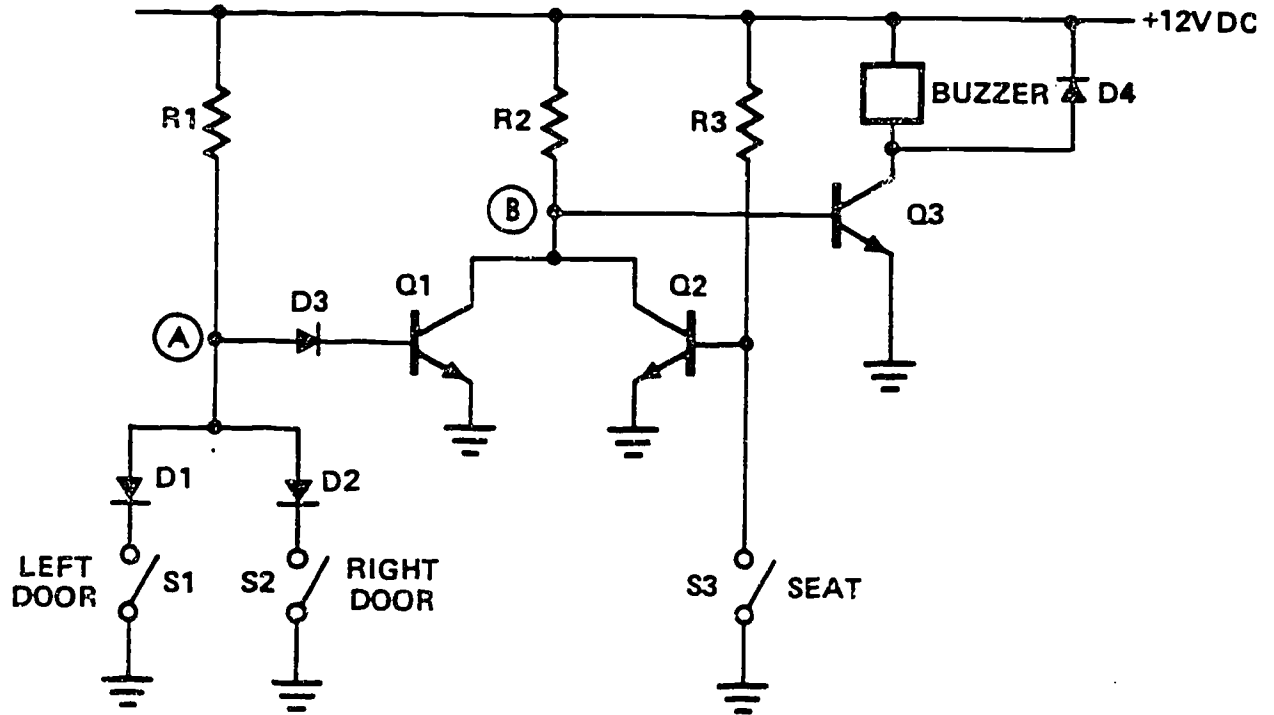


Figure 12.

base (B) to battery + via diode D_3 and resistor R_1 . It finds a path from its emitter (E) to ground (battery -). However, it finds a short from a point on the positive, feed path to ground via diode D_1 and switch S_1 and concludes that, since there is no voltage drop across the base and emitter, the collector-emitter (C-E) circuit of the transistor is in the non-conductive state. (Note that the diode is simply regarded as providing a conductive, non-resistive path to ground. There is no concept of a diode drop in this model.)

2. Q_2 attempts to evaluate its state. It finds a path from its base to battery + through R_3 , and a path from its emitter to ground. However, it finds a short from the feed path to ground through S_3 , and concludes that, since there is no voltage drop across B-E, the C-E circuit is non-conductive.

3. Q_3 attempts to evaluate its state. It finds a path from B to battery + through R_2 and a path from E to ground. Furthermore, it finds that, since the C-E circuit of Q_1 and that of Q_2 are both non-conductive, there is a voltage drop across its base and emitter and, consequently, the C-E circuit of the transistor is conductive.

4. Finally, the buzzer attempts to evaluate its state. It finds a path from one port to battery + and a path from its other port to ground via the C-E circuit of transistor Q_3 which is in the conductive state. Finally, it finds that there is no short from its feed to its ground path, since diode D_4 is non-conductive in that direction. It concludes therefore that its state is ON.

6.1.1. Models of Functional Interactions Among Devices

It is apparent from further application of the zero order logic that a change in any switch position will initiate a particular sequence of changes in device states, constituting the behavioral analysis of the circuit. It is also apparent that there are general dependencies among the devices in the circuit that are due to the effects of changes in conductivity of certain devices on the voltages drops across other devices, which in turn determine their states. These dependencies can be summarized verbally in a series of statements such as:

1. The state of the buzzer depends upon the state of Q_3 .
2. The state of Q_3 depends upon the states of Q_1 and Q_2 .
3. The state of Q_2 depends on that of S_3 .
4. The state of Q_1 depends on those of S_1 and S_2 .

These facts can be expressed in the dependency graph shown in Figure 13.




Figure 13.

Thus, by considering only the changes in device states that occur during the simulation of circuit operation, a sequence of state changes emerges together with reasons for the state changes in terms of the changes in states of other devices. An understanding of these dependencies among devices within a circuit is important for a student in bridging between behavioral and functional accounts of a circuit's operation. Developing an alternative conceptualization of a circuit in terms of functional interactions among devices can provide an important alternative way for reasoning about the operation of a circuit in its unfaulted state. For example, in the auto buzzer circuit, switches S_1 and S_2 together with diodes D_1 and D_2 constitute an OR gate which causes the voltage at A (that on the base of transistor Q_1) to go from positive (high) to zero (low) when either door is opened. Transistors Q_1 and Q_2 serve as an AND gate which causes the voltage at B to be high whenever both of its input voltages (the voltages on the bases of the transistors Q_1 and Q_2) are low. The first of these inputs is the output of the OR gate, which is low when either car door is ajar, and the second is determined by the switch in the drivers seat, and is low when the switch is closed. Finally, transistor Q_3 turns on the buzzer whenever the output of the AND gate is high. Expressed as a diagram, this functional model of the circuit is shown in Figure 14.




Figure 14.

Reasoning using functional models. Reasoning about the functioning of a circuit at this level of device interaction can be important in troubleshooting. Troubleshooting based upon a functional model allows one to reason about circuit behavior at a level at which the "parts" of the circuit are functional units representing subcircuits of the original circuit, rather than individual circuit components. The critical test points in the circuit are the input and output lines to each subcircuit. Propagation of effects of parts (subcircuits) changing state is based upon the functional interactions among parts as discussed above. When a circuit contains a fault, the troubleshooter can use a strategy such as dependency-directed backtracking to identify what functional part of the circuit is faulty, and to determine what tests can be performed to determine the particular part that is at fault. The functional model in this way allows one to determine what the outputs of the various functional parts of the circuit should be for various input conditions, and to reason about what functional parts of the circuit could be at fault given a discrepancy between the behavior of the circuit and its expected behavior in the unfaulted condition.

To illustrate, suppose the fault in our auto buzzer circuit is a bad transistor Q_1 (its C-E circuit is open). The symptoms are that the buzzer sounds whenever someone is in the driver's seat, and is independent of whether or not a door is open. Reasoning from the functional model of Figure 14, since the buzzer sounds regardless of states of S_1 and S_2 , if the fault is a single fault it is most likely associated with either the AND gate or the OR gate, on whose functioning the AND gate depends. Placing a test light at A (the output of the OR gate), the input to the AND gate is found to be good. Knowledge of the correspondence between the AND gate and the structural model of the circuit allows us to immediately localize the fault to that

portion of the circuit corresponding to the branch of the AND gate connected to the OR gate, namely, transistor Q_1 , or its connections to other circuit components. This can be substantiated by moving the test light to B, the output of the AND gate, and opening S_1 , S_2 , and S_3 . Since Q_1 is supposed to be conductive when there is a positive voltage on its base, the fault is confirmed.

Expert troubleshooters learn to reason in this way using functional models as well as using behavioral models for a circuit, and they can coordinate inferences made while reasoning with one model with inferences made while reasoning with another model. There are cases, however, where troubleshooting is confined to the functional level, as when the replaceable units are subcircuits corresponding to functional units rather than individual components.

Some limitations of zero order models. While zero order models allow one to predict the behavior of the circuit in either unfaulted or faulted states and help one derive a functional account of circuit behavior, there are a number of features of the circuit they do not account for. For example, they do not explain the purpose of the resistors, which among other things serve to limit current through the transistors. Nor do they explain the function of the diodes in the circuit. One needs to know, for example, that diodes have a constant voltage drop across them when a current flows through them (the diode drop). Given this fact, diodes D_1 and D_2 provide a switching in and out of a constant voltage (the diode drop), which is the same whether either one or both switches are closed. Diode D_3 compensates for this diode drop to set the voltage on the base of Q_1 at zero when either of the switches is closed. The purpose of diode D_4 is apparent only when one has a model for the buzzer as an inductive device which can generate large voltages when the circuit is broken by the clapper. Diode D_4 shorts out these inductive surges and thus protects transistor Q_3 . These explanations depend upon quantitative properties of diodes and transistors, and are explainable in those terms.

There are other circuits, for example amplifiers, that require yet another level of reasoning to understand their behavior. These devices, unlike the digital circuit we have been analyzing, change their output voltage in proportion to changes in input voltage, and may include feedback pathways whereby a portion of the output is mixed with the input signal. To understand such circuits, models for circuit behavior must be constructed that allow one to reason about how circuits respond when there is a change in the input voltage or a change in resistance of a component, that is, models that reason about derivatives of voltage or resistance.

6.2. First Order Qualitative Models

The zero order models we have discussed allow one to reason about circuits where the outcomes of changes in device states are discrete -- a light is either on or off, or a transistor is on (conductive) or off (non-conductive). Such models do not allow one to understand the operation of analogue circuits, that is, circuits in which changes in resistance or voltage produce incremental effects on other components -- a light becomes brighter or dimmer, or a transistor becomes more or less conductive. The first order models represent attempts to understand, using qualitative, causal logic, how such circuits operate. In first order models, the qualitative logic developed for the zero order models is extended to permit reasoning about changes in the magnitude of resistances and voltages and how they propagate within a circuit to cause other changes in voltages or resistance. For instance, first order models can predict how increasing the resistance of a device will alter voltages within the circuit.

The zero order models we have described have been implemented and used to create the instructional system described in this article. The first order models have not yet been implemented. The purpose of describing them in this article is to illustrate how the progression of zero order models can be extended to model and teach more sophisticated reasoning about circuit behaviors.

Device States

Within the zero order models we have discussed, devices are modelled as having multiple states, each of which may be thought of as a discrete level of some underlying variable describing an attribute of the device. For example, a transistor may have a collector-emitter circuit that is either purely conductive or non-conductive, depending on whether it is in the saturated or unsaturated state, and a capacitor may be either non-conductive or purely conductive, depending on whether it is in the charge or discharged state. In each case, the underlying variable referred to is the conductivity or resistance of the device. In the zero order model, these changes in conductivity influence the states of other devices in the circuit by their effects on voltage drops across those other devices. In the first order models, reasoning about the behavior of circuit components is based upon the occurrence of changes in voltages across components. These changes cause incremental changes in device variables, rather than absolute changes in those variables. Thus, the response of a transistor when there is an increase in its controlling (base-emitter) voltage is to decrease its collector-emitter resistance. Note that a series of qualitative changes

in a device variable has a cumulative effect on that variable, in this case, resistance. The general point is that, in a first order model, the existence of qualitative derivatives of circuit variables (voltage, resistance) implies that the integral of these qualitative derivatives is a scale of attribute value that is quantitative, at least at an ordinal level.¹⁰ This can provide a bridge to models which reason quantitatively about circuit variables.

Principles for reasoning about voltage

The first order qualitative models differ from the zero order models in that they reason about changes in voltage and resistance rather than about simply their presence or absence. Within the first order models, the behavior of devices within a circuit is determined by considering how changes in the conductivity (or resistance) of circuit components cause changes in the voltages across those and other components. These changes in voltages in turn cause other devices to change their states, that is, to increment or decrement some variable associated with them. Just as in the zero order models, the sequence of these device state changes that results constitutes a prediction about the overall behavior of the circuit, here one based upon a model that considers first order derivatives of voltage and resistance in reasoning about circuit operation.

As in the initial zero order model, the propagation of changes in voltages within the initial first order model is based upon reasoning using the $R \rightarrow V$ rule and upon a qualitative version of Kirchhoff's voltage law. Each time a device reevaluation leads to a change in resistance of the device, the $R \rightarrow V$ rule is employed to infer what change in voltage occurs across the device. When there is a change in the voltage across the device, Kirchhoff's voltage law (in a qualitative form) is then employed to determine the effects of that change on voltages across other circuit components. These changes in voltages, in turn, cause other components to reevaluate their states. This cycle of state changes and propagation of effects of those changes on voltage distributions within the circuit continues until the circuit stabilizes and there are no more changes in device states.

¹⁰There is thus an inconsistency within models that assume a common qualitative scale type for variables and their derivatives (cf., DeKleer, 1985). Segregating reasoning about zero order and first order derivatives into separate models (1) avoids this inconsistency, and (2) allows us to focus explicitly on the cumulative effects of multiple increments on a quantitative circuit variable.

The $R \rightarrow V$ rule and Kirchhoff's voltage law are applied in reasoning under a first order model in the following way: Whenever a device within a circuit (such as, for example, a variable resistor or the collector-emitter circuit of a transistor) changes in resistance, there is an immediate propagation of this change to a change in the voltage drop across the device using the $R \rightarrow V$ rule¹¹. A decrease/increase in resistance of a device causes a decrease/increase in the voltage across that device (except in the case where the device is connected directly to a voltage source; i.e., the voltage is fixed). The effect of this decrease/increase in voltage across a device is to alter the voltages across other devices in loops with the device, following a qualitative version of Kirchhoff's voltage law. This law states that the voltage drops across the components within such loops must sum to zero so that the loop maintains its equilibrium. This principle allows one to deduce what changes will occur in voltages across each of the components within the loop. In adding voltages for any circuit loop, the polarities of the components must be known. Propagations are possible whenever a set of like-signed voltages can be equated to a voltage whose direction of change is known.

For whatever direction one uses in traversing the loop, positive voltages are assigned whenever the polarity of a component is in the order plus-minus, and negative voltages are assigned if the polarity is minus-plus. The polarities of devices in the circuit are determined by applying the circuit orientation procedure discussed earlier. Battery polarities are given. In cases where the polarity of a component cannot be inferred from the circuit structure (for example, the bridging element in a bridge circuit), the polarity takes a preassigned value determined by the circuit designer, which is the one that would be determined if we were to take into account the particular quantitative values of resistors in the circuit. Note that if opens or shorts to ground in such a circuit create a new circuit in which a previously uninferable device polarity could now be inferred, the inferred value would override the preassigned value.

As an example of a propagation based upon these circuit principles, consider the simple series circuit of Figure 2a (in which the switch is assumed to remain open). In this circuit, increasing the resistance of R_1 causes an increase in the voltage across

¹¹The $R \rightarrow V$ rule can be shown to be true, for example, for devices connected to any circuit having a Thevenin equivalent with a non-zero resistance, or to one having a Norton equivalent (i.e., that contains a current source).

that resistor (the $R \rightarrow V$ rule). Since the voltages across R_1 and the light bulb add to the battery voltage which is unchanged, the increase in voltage across R_1 causes a decrease in voltage across the light bulb (Kirchhoff's current law).

Device models

In addition to principles for reasoning about changes in the distribution of electrical forces within a circuit, the first order qualitative models contain device models which state how a device increments or decrements some attribute (such as its resistance) in response to changes in the voltages that are applied to it. As an example of a device model, consider that for an NPN transistor¹² (This model will assume that the transistor is forward biased, that is, that the collector is always more positive than the base; cf. Horowitz and Hill, 1980, p. 51):

States: Increasing saturation, decreasing saturation.

If there is an an increase in the base-emitter voltage, then the transistor becomes more saturated.

If there is a decrease in the base-emitter voltage, then the transistor becomes less saturated.

Internal Conductivity:

If the transistor increases in saturation, then the collector-emitter path within the transistor becomes less resistive.

If the transistor becomes less saturated, the collector-emitter path becomes more resistive.

The transistor is conductive from base (+) to emitter (-).

It is non-conductive from emitter (-) to base (+).

It is non-conductive from emitter (-) to collector (+).

It is non-conductive in either polarity from base to collector.

Voltage Source: The transistor is not a voltage source.

The internal conductivity rules for the transistor are similar to those used in

¹²For the PNP transistor, reverse all polarities.

the zero order qualitative models. The state rules link qualitative changes in the resistance of the collector-emitter circuit of the transistor to qualitative changes in the controlling voltage, the voltage applied to the base and emitter.

Control structure

1. When a device has changed state. In analyzing the behavior of a circuit under the first order model, when a device within a circuit changes its state, the $R \rightarrow V$ rule is first applied to determine changes in voltages across that component resulting from the change in resistance of the device. Then, the qualitative version of Kirchhoff's voltage law is used to propagate the effects of that change in voltage on all other voltage drops within any loops in which the component is a part. Whenever there have been changes in voltages within the circuit as a result of a device changing its state, all other devices in the circuit are prompted to reevaluate their states.

2. When a device reevaluates its state. To establish whether or not there have been any changes in voltages within the circuit that may influence their states, each device undergoing reevaluation (a) looks to see if a change in the voltage across its controlling ports has occurred, or (b) employs a circuit tracing procedure similar to that of the zero order model to find out if any changes in voltage across components in loops with the controlling ports of the device have occurred. In the latter event, Kirchhoff's voltage law (in its qualitative form) is employed to ascertain the change in voltage across the device undergoing reevaluation. If a change in voltage across the controlling ports of a device is inferred, that device then changes state following the rules stated in its device model.

3. Reevaluations are parallel. As in the zero order model, to avoid unwanted order effects in evaluating the effects of devices changing state on other devices, when an initial state changes occurs, the voltage distributions within the circuit that result are frozen until all other devices have ascertained their states within those changed conditions. Then, on the next cycle, those devices that change state propagate their effects on voltage distributions within the circuit, and those conditions are frozen while the rest of the devices in the circuit reevaluate their states. This process continues until no further changes in device states occur.

An example of reasoning using a first order qualitative model

With this background, it is now possible to give an example of how a first order model reasons about the behavior of a circuit. Figure 15 shows a simple Schmitt trigger circuit. The Schmitt trigger is a positive feedback circuit that reacts to an increase in input voltage by changing its output voltage from some initial low value (determined by the resistances in the voltage divider formed by R_2 , the collector-emitter circuit of Q_2 , and R_4) to a high value, namely the battery voltage. The circuit serves the function of "monitoring" an input signal; when that signal reaches a critical or threshold value, the circuit switches its output voltage from the low to the high or battery voltage level. For purposes of the example, suppose that initially the input is zero.

1. The initial event is an increment in the input voltage, the voltage between the base of transistor Q_1 and ground.

2. Transistor Q_1 attempts to evaluate its state. Applying the circuit tracing procedure, it finds that its controlling parts (the collector and emitter) are in a loop with a component (the input impedance) whose voltage has changed. Applying the Kirchhoff voltage law, the voltages across the base and emitter of transistor Q_1 and across resistor R_4 can be inferred to increase, since they sum to that of the input.

3. Under the transistor model, since there is an increase in its base-emitter voltage, the resistance of the collector-emitter circuit of the transistor decreases.

4. Applying the $R \rightarrow V$ rule, this causes the voltage across the collector and emitter to decrease.

5. This change is then propagated within the two loops which contain the collector-emitter of transistor Q_1 : (a) within the voltage divider made up of R_1 , the collector-emitter of Q_1 , and R_4 , the voltage across each of the resistors increases since, given the polarities of the components within the loop, the three voltages must sum to that of the battery, and if one of those voltages decreases the others must show a compensating increase; and (b) within the loop consisting of the collector-emitter of Q_1 , resistor R_3 , and the base-emitter of transistor Q_2 , the voltages across the resistor and transistor Q_2 decrease since, given the polarities of the components, they must sum to the voltage across transistor Q_1 , which decreased.

6. Transistor Q_2 attempts to evaluate its state. It checks for a change in the voltage across its base and emitter. Since there has been a decrease in the voltage across these terminals of the transistor, the transistor model causes the resistance of the collector-emitter circuit of Q_2 to increase.




Figure 15.

7. Applying the $R \rightarrow V$ rule, this causes the voltage drop across the collector and emitter of the transistor also to increase.

8. Propagating this change in voltage within the voltage divider formed by R_2 , the collector-emitter circuit of Q_2 , and R_4 , this increase in voltage across the transistor causes a decrease in the voltage across each of the resistors.

9. Transistor Q_1 again attempts to evaluate its state. Applying the circuit tracing procedure, the base and emitter of this transistor along with resistor R_4 are found to form a voltage divider connected across the source voltage, which is assumed to be unchanged. Since the voltage across resistor R_4 has decreased, the voltage across the base and emitter of Q_1 must have increased (since the two voltage drops within the voltage divider sum to the source voltage).

10. Therefore, applying the transistor model, this causes the resistance of the collector-emitter circuit of Q_1 to show a further decrease.

11. This in turn causes the voltage across the collector and emitter of the transistor to decrease still further, and again this change in voltage is propagated as in steps 4 and 5.

12. Transistor Q_2 , prompted by the change in state of the first transistor, again reevaluates its state as in step 6, and so on. As the positive feedback cycle is repeated, transistor Q_1 becomes more and more conductive, while transistor Q_2 becomes less and less so.

Limitations of First Order Models

Recognizing terminal states. The first order model reasons only about changes in resistance and voltage, and thus cannot terminate this positive feedback loop. In reality, at a certain point Q_2 would be "turned off", the collector-emitter circuit would become effectively an open, and the output voltage would become the battery voltage. Moreover, the first order model cannot "know" about threshold input or output voltages or even saturated and unsaturated states of a transistor. Lacking a quantitative representation of the input voltage and a quantitative model of the transistor, the circuit model will "trigger" on the first increase in input voltage and continue endlessly in a positive feedback loop. Even more seriously, if a decrease in input voltage occurs, the model will encounter an ambiguity as to the voltage change across the base and emitter of transistor Q_1 in step 9. Here the positive feedback will cause an increase in the voltage drop across Q_1 , while the input change causes a decrease in the same voltage. Lacking a quantitative representation of voltage, the model cannot weigh these alternative influences. If the model could recognize its

terminal state (for example, when Q_1 becomes totally saturated and Q_2 becomes unsaturated), a mechanism would exist for ending the positive feedback, at which point, a decrease in input voltage would then initiate a second feedback period which would lead back to its untriggered state.

In order to overcome this problem, the device models for devices such as transistors could contain rules of the form: if the transistor is becoming more saturated/less saturated, then after a certain number of model cycles, the device will change its state to purely saturated/unsaturated. This would permit the model to recognize when a transistor becomes saturated, which would in turn influence the behavior of the model. In the above example, when transistor Q_2 becomes saturated, it will no longer have the effect of decreasing the voltage across transistor Q_1 , and the effect of a decrease in input voltage will be unambiguous: The trigger will be able to respond by decreasing its output voltage to its low level, that is, return to its "untriggered" state. However, the model will still be unable to recognize thresholds, or to resolve ambiguities that occur when the input voltage decreases while it is still in the positive feedback cycle leading toward the "triggered" state.

An interesting point is that qualitative models in a sense "know" the limits of their own reasoning processes. For instance, they know that they cannot determine when the transistor will become saturated. By simply articulating their reasoning, they can then communicate this important knowledge to the student. They can also recognize when they encounter ambiguities, and can report those to the student. The student's mental model will thus also know its own limits. This lack of determinacy will also motivate students to want to acquire quantitative models for circuit behavior.

6.3. Quantitative Extensions

The limitation of first order models to reasoning qualitatively about increments and decrements in resistance and voltage precludes an understanding, on that basis alone, of one of the fundamental functions of the trigger circuit: hysteresis. The trigger circuit is designed to have a high threshold before it has triggered, and to have a low threshold after it has triggered. Thus, the voltage required to trigger the circuit is initially higher than that required to return it to the untriggered state. This prevents the circuit from wildly triggering on and off when there is noise in the input signal and when the input voltage is near the triggering value. This feature of the circuit is basic to its design, and illustrates the kind of problem that will require

quantitative extensions of the first order model. Such extensions of the circuit theory are needed to reason about the notion of a threshold, let alone the means by which the circuit achieves differential thresholds for its untriggered and triggered states.

Quantitative extensions of the first order model include (1) using proportional reasoning about voltages within a resistive network in place of qualitative reasoning about increments and decrements in voltages, and (2) using heuristics for analyzing a circuit, such as examining extreme cases. Proportional reasoning about voltages, for example, states that within a loop, the proportion of the total voltage across an individual resistive component is given by the ratio of the resistance of that component to the total resistance in the loop. An example of reasoning from extreme cases is an examination of the trigger circuit in its untriggered and triggered states. These two techniques enable one to explain a variety of quantitative circuit phenomena, still without resorting to algebra and/or calculation. They offer a bridge to purely quantitative analyses of circuits, while at the same time using techniques that are often employed by engineers in reasoning about circuits.

We shall illustrate these techniques by using them to explain how hysteresis is produced in the Schmitt trigger. This is accomplished in the circuit design by making the resistance of R_1 greater than that of R_2 .

1. The voltage across the base and emitter of Q_1 is determined (a) by the input voltage, which is applied to the voltage divider formed by R_1 and R_4 , and (b) the voltage applied to R_1 , which is determined by other other loops within the circuit and depends upon the state of the trigger circuit.
2. To raise the threshold of the trigger, the voltage drop across R_4 should be made high, so that a larger proportion of the input voltage is applied to R_4 than to the base-emitter of Q_1 . To lower the threshold, the voltage drop across R_4 should be made low, so that the proportion of the input voltage applied to Q_1 will be large.
3. Examining the extreme cases, when the circuit is untriggered, Q_1 will be unsaturated and Q_2 saturated. Similarly, when the circuit is triggered, Q_1 will be saturated and Q_2 unsaturated. (These states of transistors Q_1 and Q_2 can be determined by applying the zero order model to the circuit.)
4. When the circuit is in the untriggered state, the voltage across R_4 is determined by the voltage divider formed by R_2 and R_4 , since the collector-emitter of Q_2 is purely conductive and that of Q_1 is nonconductive (open). To keep the threshold of the trigger high, the voltage across R_4 must be made high. Therefore, the resistance of R_2 should be low.

5. When the circuit is in the triggered state, the voltage across R_4 is determined by the voltage divider formed by R_1 and R_4 , since Q_1 is now conductive and Q_2 is nonconductive. To keep the threshold of the trigger low, the voltage across R_4 must be kept low. Therefore, the resistance of R_1 should be high.
6. To create hysteresis, then, the resistance of R_1 should be high and that of R_2 should be low.

This example serves to illustrate how, by the use of proportional reasoning and by examining and comparing the extreme states of a circuit, one can reason about the relative magnitudes of resistances and voltages needed to create a particular circuit behavior. In the example, an important inequality that is fundamental to the design of the Schmitt trigger can be derived. Such techniques could also be employed in reasoning from the design (e.g., the inequality in resistances) to its effects on the behavior of the circuit. Such reasoning about the effects of increasing or decreasing the resistance of a component on the quantitative behavior of a circuit constitutes an important transition step in learning quantitative circuit theory. It can also be more valuable than algebraic reasoning using constraint equations if one is attempting to develop an understanding of the operation of a circuit.

8.4. Model Similarities

We have seen that, in many respects, the first order models are similar to the zero order models. They share important concepts of device models, circuit tracing logic, the notion of devices being oriented within the circuit, the underlying qualitative circuit laws governing the occurrence of voltage drops throughout the circuit, and control structure. However, within the first order models, circuit tracing now seeks devices that have changed their voltage drops, not just sources of voltage. Devices are not modelled in terms of discrete, qualitative states, but are qualitatively incremented or decremented. It is thus implicitly assumed that there is an underlying quantitative attribute for a device whose value represents its state. The logic is now based upon changes in voltages and resistances, not just on their presence or absence. And finally, in the control structure, when devices change state, the effects must now be immediately propagated so that the implications of those changes for distribution of voltages across other devices are derived.

Given these similarities, the zero and first order models are clearly from the same family, and it should be possible for students to learn to reason using either

model to obtain alternative views of the operation of a circuit. Finally, we have seen that for some purposes it may be desirable to coordinate reasoning using these two model types, for example, in order to understand boundary conditions or to reason about quantitative behavior of the circuit.

6.5. Understanding of a Domain

The system we have described attempts to give causal accounts of circuit behavior in terms of voltages and resistances. As we have seen, this is not the only way of conceptualizing how a circuit works. We argue that whether or not a person has an understanding of a domain cannot be assessed with respect to a single conceptualization only. For instance, just because an individual is very adept at looking at circuit diagrams and predicting the behavior of circuits, does not mean that the individual has a "deep" understanding of electrical circuits. Such an individual may be completely unable to describe the functionality of circuits - the purpose of a circuit as a whole and the role that subsets of devices play in achieving that purpose. Also, the person may understand nothing about the physics of device and circuit functioning. Further, he or she may only be able to reason at a qualitative level and thus be unable to formalize his or her understanding by constructing a quantitative model of circuit behavior. Thus we claim that "deep understanding" relates to having such alternative conceptualizations of the phenomena of a domain.

We define understanding of a domain with respect to a number of dimensions. The first relates to the number of different types of mental models that a person has with respect to the domain - e.g., mechanistic, behavioral, functional, etc. The second relates to the form of the model - e.g., does it utilize qualitative or quantitative reasoning? The third dimension has to do with the level of understanding that a person has with respect to their set of mental models for the domain - e.g., what level model, in terms of their degree of elaboration, does the person possess? The fourth and final dimension relates to the ability to make use of and coordinate these alternative models for reasoning within a domain - e.g., can the person utilize, in coordination, both functional and behavioral models when solving circuit problems?

Finally, by creating causal models that reason about circuit behavior in terms of forces and equilibrium, we hope to create potential links between circuit behavior and the physical laws underlying how circuits work. Thus we claim that a person who has this type of qualitative causal model has a deeper understanding of circuit behavior

than someone who has a model of circuit behavior that enables predictions but whose reasoning is not causally consistent. Further, the person who has this type of model will be better able to link their knowledge of circuit behavior with even deeper accounts of the physics underlying how circuits work.

Summary. The use of progressions of models as the foundation for an intelligent learning environment has served not only a pedagogical function, but has also allowed students to develop multiple models of circuit behavior. Reasoning about a circuit in multiple ways allows for different conceptualizations that in turn serve different purposes. For example, zero order models facilitate reasoning about gross circuit behavior, and can be used in studying the behavior of digital circuits and their functionality. They can also be used in analyzing extreme cases when one is studying the behavior of analogue circuits such as the trigger circuit. First order models are useful in studying analogue circuits, and can explain feedback, or how such circuits respond to changes in input voltages. Furthermore, they can serve as a bridge to reasoning using quantitative models. Quantitative models can explain such features of circuit behavior as thresholds, can provide the reason certain components are present within a circuit, and can of course be used to calculate actual voltages and currents within a circuit. An important problem for future research is the theory selection problem: how do experts invoke appropriate conceptualizations for a particular problem at hand, and how can students be taught how to select and coordinate multiple models in problem solving.

7. The Extendibility of This Approach to Other Domains

One might be tempted to conclude that the design for intelligent learning environments articulated in this paper would have utility for only a small number of domains. For instance, one could infer that it applies only to physical systems such as electrical or mechanical systems. However, we argue that this approach can be applied to any domain that can be taught by problems solving in the context of interactive simulations.

As an example, White (1981) utilized problem solving in the context of a dynamic simulation to teach high school students about the implications of Newton's laws of motion. The computer simulation embodied Newton's first two laws by simulating the motion of an object on a display screen (diSessa, 1979). The student could control the object's motion by applying fixed-sized impulse forces to it in various directions via keyboard commands. The object responded to the application of the impulse forces in accordance with Newton's second law ($F = ma$) by accelerating instantaneously to the appropriate velocity. The resulting motion of the object across the display screen also obeyed Newton's laws since it moved with a constant velocity until another force was applied, or until it crashed into an obstacle. Within this simulation context, students were given game-like problems solving activities where they had to, for instance, navigate the object around a track.

The design of the simulation and problem solving activities was based upon a cognitive analysis which considered the relevant physical theory, the misconceptions and preconceptions that students bring to this domain, and the form of expert knowledge in the domain. The results of the cognitive analysis then constrained the form of computer representations used to portray physical concepts and laws, the nature of the educational activities (problems and examples) embedded in this environment, and the sequencing of these activities. The learning environment proved effective at helping students to learn to reason about force and motion problems (White, 1984).

This work was based on a quantitative model and did not have the explanatory capability of the instructional system described in this paper. Further, it was based upon a single model not upon a progression of models. White and Horwitz are currently extending this research to incorporate a progression models, both qualitative and quantitative, for reasoning in the domain of elementary mechanics. Unlike electricity, where we argued for teaching purely qualitative reasoning for an

extended period, in the domain of mechanics we argue for the need to introduce simple, already-understood quantitative models early in the model progression.

To elaborate, initially students learn to reason qualitatively in one dimension. For instance, they learn that more impulses applied in the direction of motion produces more speed. A qualitative simulation that captures such relationships could be devised. Following the acquisition of such qualitative rules, students go on to learn that the effects of impulse forces can be modelled by scalar arithmetic, which they mastered in second grade. For instance, they learn that the effects of impulses add and subtract, e.g., $3 \rightarrow + \leftarrow 2 = 1 \rightarrow$. Such a quantitative model could also enable the computer to accurately simulate the effects of forces on the motions of objects. When motion in two dimensions is introduced, the focus is again initially on qualitative reasoning. Simple quantitative models are introduced once students have understood the domain in qualitative terms. By (i) focusing on qualitative reasoning and introducing previously acquired, simple quantitative models, by (ii) restricting the application of forces to fixed sized impulses in one of four orthogonal directions, and by (iii) conceiving of motion in terms of its orthogonal velocity components, we have enabled sixth graders to accurately predict the effects of impulse forces on the motion of objects.

The central thesis of this paper is that, at any point in the learning progression, the model driving the computer simulation should be in the form of the desired student mental model. This constraint does not restrict one to purely qualitative models, although, as we have argued, it is of primary importance to teach qualitative understanding. The focus is on producing progressions of models that link to the students' prior knowledge. These model progressions enable the learning environment to (1) aptly represent the domain phenomena, (2) let students interact with that phenomena via experimentation and problem solving, and (3) provide students with feedback and explanations.

The claim is that any domain whose phenomena can be captured by laws affecting the behavior of objects can be tutored via problems and examples in the context of a simulation driven by a progression of causal models. This includes aspects of physics, chemistry, biology, medicine, and even mathematics (e.g., Feurzeig & White, 1983), as well as more applied domains such as automotive troubleshooting or airplane maintenance.

8. Acknowledgements

This research was supported by the Office of Naval Research and the Army Research Institute, under ONR contract N00014-82-C-0580 with the Personnel and Training Research Program. The work benefited from the perspectives and innumerable contributions of Wallace Feurzeig. The current implementation of the system was programmed by Frank Ritter, who helped with the system design, particularly the circuit orientation scheme and the user interface. In addition, we would like to thank Allan Collins and Paul Horwitz for their many thought provoking comments.

9. References

- Anderson, J.R., Farrell, R., & Sauers, R. (1984). Learning to program in LISP. Cognitive Science, 8, 87-129.
- Bobrow, D.G. (Ed.) (1985). Qualitative Reasoning about Physical Systems. Cambridge, MA: MIT Press.
- Bonar, J.G. & Logan, D. (1986). Intelligent tutoring of basic electricity. Paper presented at AERA Annual Meeting, San Francisco, California.
- Bond, L., Eastman, R., Gitomer, D., Glaser, R., & Lesgold, A. (1983). Cognitive task analyses of technical skills: Rationale and approach. LRDC Report, University of Pittsburgh, Pennsylvania.
- Brown, J.S., & Burton, R.R. (1978). Diagnostic models for procedural bugs in basic mathematical skills. Cognitive Science, 2, 155-192.
- Brown, J.S., & Van Lehn, K. (1980). Repair theory: A generative theory of bugs in procedural skills. Cognitive Science, 4, 379-426.
- Brown, J.S., Burton, R.R., & deKleer, J. (1982). Pedagogical, natural language and knowledge engineering techniques in SOPHIE I, II and III. In Sleeman & Brown (Eds.), Intelligent Tutoring Systems. New York: Academic Press.
- Brown, J.S. & deKleer, J. (1985). A qualitative physics based upon confluences. In Bobrow, D.G. (Ed.). Qualitative Reasoning about Physical Systems. Cambridge, MA: MIT Press.
- Bruner, J.S., Goodnow, J.J., & Austin, G.A. (1956). A Study of Thinking. New York: Wiley.

Clancey, W.J. (1982). Tutoring rules for guiding a case method dialogue. In Sleeman & Brown (Eds.), Intelligent Tutoring Systems. New York: Academic Press.

Chi, M. T. H., & Glaser, R. (1980). The measurement of expertise: Analysis of the development of knowledge and skill as a basis for assessing achievement. In E. L. Baker & E. S. Quellmalz (Eds.), Educational Testing Evaluation. Beverly Hills, CA: Sage Publications.

Chi, M., Feltovich, P., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. Cognitive Science, 5, 121-152.

Cohen, R., Eylon, B., & Ganiel, U. (1983). Potential difference and current in simple electric circuits: a study of students' concepts. American Journal of Physics, 51(5), 407-412.

Collins, A. (1985). Component models of physical systems. Proceedings of the Seventh Annual Conference of the Cognitive Science Society, University of California, Irvine.

Davis, R. (1983). Reasoning from first principles in electronic troubleshooting. International Journal of Man-Machine Studies, 19, 403-423.

deKleer, J. (1979). Causal and teleological reasoning in circuit recognition. TR-529, MIT Artificial Intelligence Laboratory, Cambridge, MA.

deKleer, J. (1995). How circuits work. In Bobrow, D.G. (Ed.). Qualitative Reasoning about Physical Systems. Cambridge, MA: MIT Press.

diSessa, A. Dynamics: (1979). Learning physics with a dynaturtle. In Papert, S., Watt, D., diSessa, A., & Weir, S., *Final report of the Brookline Logo project, Part II: Project summary and data analysis* (Memo 545). Pp. 6.1-6.20. Cambridge, MA: M.I.T., A.I. Laboratory.

Feurzeig, W. & White, B.Y. (1983). Development of an articulate instructional system for teaching arithmetic procedures. BBN Report No. 5484, BBN Laboratories, Cambridge, Massachusetts.

Forbus, K., & Stevens, A.S. (1981). Using qualitative simulation to generate explanations. BBN Report No. 4490. Cambridge, MA: BBN Laboratories.

Forbus, K.D. (1985). Qualitative process theory. In Bobrow, D.G. (Ed.). Qualitative Reasoning about Physical Systems. Cambridge, MA: MIT Press.

Frederiksen, J., & Warren, B. (1986). A cognitive framework for developing expertise in reading. In Glaser, R. (Ed.). Advances in Instructional Psychology. NJ: Lawrence Erlbaum Associates.

Gentner, D., & Stevens, A. L. (1983). Mental Models. Hillsdale, NJ: Lawrence Erlbaum Associates.

Gitomer, D. H. (1984). A cognitive analysis of a complex troubleshooting task. Unpublished doctoral dissertation, University of Pittsburgh, Pittsburgh, PA.

Goldstein, I.P. (1982). The genetic graph: A representation for the evolution of procedural knowledge. In Sleeman, D., & Brown, J. S. (Editors). Intelligent Tutoring Systems. London: Academic Press.

Horowitz, P., & Hill, W. (1980). The Art of Electronics. Cambridge, England: Cambridge University Press.

Johnson, L. & Soloway, E. (1984). Intention-based diagnosis of programming errors. In Proceedings of the National Conference on Artificial Intelligence. Austin, Texas: NCAL.

Kuipers, B. (1985). Commonsense reasoning about causality: Deriving behavior from structure. In Bobrow, D.G. (Ed.). Qualitative Reasoning about Physical Systems. Cambridge, MA: MIT Press.

Larkin, J.H., McDermott, J., Simon, D.P., & Simon, H.A. (1980). Expert and novice performance in solving physics problems. Science, 208, 1335-1342.

Lewis, M.W. & Anderson, J.R. (1985). Discrimination of operator schemata in problem solving: Learning from examples. Cognitive Psychology, 17, 26-65.

O'Shea, T. (1982). A self-improving quadratic tutor. In Sleeman, D., & Brown, J. S. (Eds.). Intelligent Tutoring Systems. London: Academic Press.

Rasmussen, J., & Jensen, A. (1974). Mental procedures in real life tasks: a case study of electronic troubleshooting. Ergonomics, 17, 193-307.

Reiser, B.J., Anderson, J.R., & Farrell, R.G. (1985). Dynamic student modelling in an intelligent tutor for LISP programming. In Proceedings of IJCAI-85. Los Angeles, CA: IJCAI, 8-14.

Richer, M.H. & Clancey, W.J. (1985). Guidon-Watch: A graphic interface for viewing a knowledge-based system. Department of Computer Science Report No. STAN-CS-85-1068, Stanford University, California.

Riley, M. S. (1984). Structural understanding in performance and learning. Unpublished doctoral dissertation. University of Pittsburgh, Pittsburgh, PA.

Rouse, W. B., & Morris, N. M. (1985). On looking into the black box: Prospects and limits in the search for mental models. Center for Man-Machine Systems Research Report No. 85-2, Atlanta, GA: Georgia Institute of Technology.

Smith, E. E., & Goodman, L. (1984). Understanding instructions: The role of an explanatory schema. Cognition and Instruction, 1, 359-396.

Sleeman, D. & Hendley, R.J. (1982). ACE: A system which analyses complex explanations. In Sleeman, D., & Brown, J. S. (Eds.). Intelligent Tutoring Systems. London: Academic Press.

Sleeman, D., & Brown, J. S. (Eds.) (1982). Intelligent Tutoring Systems. London: Academic Press.

Soloway, E., Rubin, E., Woolf, B., Bonar, J., & Johnson, W.L. (1983). Meno-II: An AI-based programming tutor. Journal of Computer-Based Instruction, 10, 1.

Steinberg, M.S. (1983). Reinventing electricity. In Proceedings of the International Seminar, Misconceptions in Science and Mathematics, Ithaca, New York.

Weld, D. (1983). Explaining complex engineering devices. BBN Report No. 5489, BBN Laboratories, Cambridge, Massachusetts.

White, B.Y. (1981). Designing computer games to facilitate learning. Technical Report AI-TR-619, Artificial Intelligence Laboratory, M.I.T., Cambridge, Massachusetts.

White, B.Y. (1984). Designing computer activities to help physics students understand Newton's laws of motion. Cognition and Instruction, 1, 69-108.

White, B.Y., & Frederiksen, J.R. (1984). Modeling Expertise in Troubleshooting and Reasoning about Simple Electric Circuits. In the Proceeding of the Annual Meeting of the Cognitive Science Society, Boulder, Colorado.

White, B. Y., & Frederiksen, J. R. (1985). QUEST: Qualitative understanding of electrical system troubleshooting. ACM SIGART Newsletter, 93, 34-37.

White, B.Y. (in preparation). Multiple Muddled Models: Novice Reasoning about Force and Motion.

Williams, B.C. (1985). Qualitative analysis of MOS circuits. In Bobrow, D.G. (Ed.) (Ed.). Qualitative Reasoning about Physical Systems. Cambridge, MA: MIT Press.

Human Resources Lab
 D
 FB, TX 78235

rt Ahlers
 1
 ctors Laboratory
 aining Systems Center
 FL 32813

ken
 onnel R&D Center
 o, CA 92152-8800

am E. Alley
 T
 FB, TX 78235

A. Alluisi
 L (AFSC)
 FB, TX 78235

e Andriole
 ason University
 f Information
 ology & Engineering
 iversity Drive
 VA 22030

Director, ARI
 enhower Avenue
 ia, VA 22333

cia Baggett
 y of Colorado
 at of Psychology

CO 80309

S. Baker
 onnel R&D Center
 o, CA 92152-8800

Black
 Columbia Univ.
 121st Street
 NY 10027

r S. Blaiwes

aining Systems Center
 FL 32813

Dr. Jeff Bonar
 Learning R&D Center
 University of Pittsburgh
 Pittsburgh, PA 15280

Dr. Jomills H. Braddock II
 Center for the Social
 Organization of Schools
 The Johns Hopkins University
 3505 North Charles Street
 Baltimore, MD 21218

Dr. Ann Brown
 Center for the Study of Reading
 University of Illinois
 51 Gerly Drive
 Champaign, IL 61280

Dr. John S. Brown
 XEROX Palo Alto Research
 Center
 3333 Coyote Road
 Palo Alto, CA 94304

Dr. Patricia A. Butler
 OERI
 555 New Jersey Ave., NW
 Washington, DC 20208

Dr. Tom Cafferty
 Dept. of Psychology
 University of South Carolina
 Columbia, SC 29208

Dr. Robert Calfee
 School of Education
 Stanford University
 Stanford, CA 94305

Joanne Capper
 Center for Research into Practice
 1718 Connecticut Ave., N.W.
 Washington, DC 20009

Dr. Susan Carey
 Harvard Graduate School of
 Education
 337 Gutman Library
 Appian Way
 Cambridge, MA 02138

Dr. Pat Carpenter
 Carnegie-Mellon University
 Department of Psychology
 Pittsburgh, PA 15213

LCDR Robert Carter
 Office of the Chief
 of Naval Operations
 OP-01B
 Pentagon
 Washington, DC 20350-2000

Dr. Fred Chang
 Navy Personnel R&D Center
 Code 51
 San Diego, CA 92152-8800

Dr. Alphonse Chapanis
 8415 Bellona Lane
 Suite 210
 Buxton Towers
 Baltimore, MD 21204

Dr. Paul R. Chatelier
 OUSDRE
 Pentagon
 Washington, DC 20350-2000

Dr. Michelene Chi
 Learning R & D Center
 University of Pittsburgh
 3929 O'Hara Street
 Pittsburgh, PA 15213

Mr. Raymond E. Christal
 AFHRL/MOE
 Brooks AFB, TX 78235

Dr. Allan H. Collins
 Bolt Beranek & Newman, Inc.
 50 Moulton Street
 Cambridge, MA 02138

Dr. John J. Collins
 Director, Field Research
 Office, Orlando
 NPRDC Liaison Officer
 NTSC Orlando, FL 32813

Dr. Stanley Collyer
 Office of Naval Technology
 Code 222
 800 N. Quincy Street
 Arlington, VA 22217-5000

LT Judy Crookshanks
 Chief of Naval Operations
 OP-112G5
 Washington, DC 20370-2000

Dr. Cary Czichon
 Intelligent Instructional Systems
 Texas Instruments AI Lab
 P.O. Box 680245
 Dallas, TX 75288

LT John Deaton
 ONR Code 125
 800 N. Quincy Street
 Arlington, VA 22217-5000

Dr. Natalie Dehn
 Department of Computer and
 Information Science
 University of Oregon
 Eugene, OR 97403

Dr. Gerald F. DeJong
 Artificial Intelligence Group
 Coordinated Science Laboratory
 University of Illinois
 Urbana, IL 61801

Goery Delacote
 Directeur de L'informatique
 Scientifique et Technique
 CNRS

15, Quai Anatole France
 75700 Paris FRANCE

Defense Technical
 Information Center
 Cameron Station, Bldg 5
 Alexandria, VA 22314
 Attn: TC
 (12 Copies)

Dr. Richard Elster
Deputy Assistant Secretary
of the Navy (Manpower)
OASN (M&RA)
Department of the Navy
Washington, DC 20350-1000

Dr. Susan Embretson
University of Kansas
Psychology Department
426 Fraser
Lawrence, KS 66045

Dr. Randy Engle
Department of Psychology
University of South Carolina
Columbia, SC 29208

ERIC Facility-Acquisitions
4833 Rugby Avenue
Bethesda, MD 20014

Dr. Edward Esly
Department of Education, OERI
Room 717D
1200 19th St., NW
Washington, DC 20208

Dr. Paul Fellovich
Southern Illinois University
School of Medicine
Medical Education Department
P.O. Box 3928
Springfield, IL 62708

Mr. Wallace Feurzeig
Educational Technology
Bolt Beranek & Newman
10 Moulton St.
Cambridge, MA 02238

Dr. Gerhard Fischer
University of Colorado
Department of Computer Science
Boulder, CO 80309

J. D. Fletcher
9931 Corsica Street
Vienna VA 22180

Dr. Kenneth D. Forbus
University of Illinois
Department of Computer Science
1304 West Springfield Avenue
Urbana, IL 61801

Dr. Barbara A. Fox
University of Colorado
Department of Linguistics
Boulder, CO 80309

Dr. John R. Frederiksen
Bolt Beranek & Newman
50 Moulton Street
Cambridge, MA 02138

Dr. R. Edward Gelselman
Department of Psychology
University of California
Los Angeles, CA 90024

Dr. Michael Genesereth
Stanford University
Computer Science Department
Stanford, CA 94305

Dr. Dedre Gentner
University of Illinois
Department of Psychology
603 E. Daniel St.
Champaign, IL 61820

Dr. Herbert Ginsburg
Teachers College, Columbia Univ.
525 West 121st Street
New York, NY 10027

Dr. Robert Glaser
Learning Research
& Development Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15260

Dr. Marvin D. Glock
13 Stone Hall
Cornell University
Ithaca, NY 14853

Dr. Gene L. Gloye
Office of Naval Research
Detachment
1030 E. Green Street
Pasadena, CA 91106-2485

Dr. Joseph Goguen
Computer Science Laboratory
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025

Dr. Susan Goldman
University of California
Santa Barbara, CA 93106

Dr. Sherrie Gott
AFHRL/MODJ
Brooks AFB, TX 78235

Jordan Grafman, Ph.D.
2021 Lyttonville Road
Silver Spring, MD 20910

Dr. Wayne Gray
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. James G. Greeno
University of California
Berkeley, CA 94720

Dr. Henry M. Hallf
Hallf Resources, Inc.
4918 33rd Road, North
Arlington, VA 22207

Dr. Ronald K. Hambleton
Prof. of Education & Psychology
University of Massachusetts
at Amherst
Hilla House
Amherst, MA 01003

Dr. Bruce W. Hamill
Johns Hopkins University
Applied Physics Laboratory
Johns Hopkins Road
Laurel, MD 20707

Mr. William Hartung
PEAM Product Manager
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Barbara Hayes-Roth
Department of Computer Science
Stanford University
Stanford, CA 95305

Dr. Frederick Hayes-Roth
Teknowledge
525 University Ave.
Palo Alto, CA 94301

Dr. Joan I. Heller
505 Haddon Road
Oakland, CA 94608

Dr. Jim Hollan
Intelligent Systems Group
Institute for
Cognitive Science (C-015)
UCSD
La Jolla, CA 92093

Dr. Melissa Holland
Army Research Institute for the
Behavioral and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Keith Holyoak
University of Michigan
Human Performance Center
330 Packard Road
Ann Arbor, MI 48109

Dr. James Howard
Dept. of Psychology
Human Performance Laboratory
Catholic University of
America
Washington, DC 20064

Dr. Ed Hutchins
Intelligent Systems Group
Institute for
Cognitive Science (C-015)
UCSD
La Jolla, CA 92093

Dr. Dillon Inouye
WICAT Education Institute
Provo, UT 84057

Dr. Zachary Jacobson
Bureau of Management Consulting
365 Laurier Avenue West
Ottawa, Ontario K1A 0S5
CANADA

Dr. Robin Jeffries
Hewlett-Packard Laboratories
P.O. Box 10490
Palo Alto, CA 94303-0971

Margaret Jerome
c/o Dr. Peter Chandler
83. The Drive
Hove
Sussex
UNITED KINGDOM

Dr. Joseph E. Johnson
Assistant Dean for
Graduate Studies
College of Science and Mathematics
University of South Carolina
Columbia, SC 29208

CDR Tom Jones
ONR Code 125
800 N. Quincy Street
Arlington, VA 22217-5000

Dr. Douglas A. Jones
P.O. Box 6640
Lawrenceville
NJ 08648

Dr. Marcel Just
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Ruth Kanfer
University of Minnesota
Department of Psychology
Elliott Hall
75 E. River Road
Minneapolis, MN 55455

Dr. Milton S. Katz
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Frank Keil
Department of Psychology
Cornell University
Ithaca, NY 14850

Dr. David Kieras
University of Michigan
Technical Communication
College of Engineering
1223 E. Engineering Building
Ann Arbor, MI 48109

Dr. Paula Kirk
Oakridge Associated Universities
University Programs Division
P.O. Box 117
Oakridge, TN 37831-0117

Dr. David Klahr
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Kenneth Kotovsky
Department of Psychology
Community College of
Allegheny County
800 Allegheny Avenue
Pittsburgh, PA 15233

Dr. Benjamin Kuipers
University of Texas at Austin
Department of Computer Sciences
T.S. Painter Hall 3.28
Austin, Texas 78712

Dr. Pat Langley
University of California
Department of Information
and Computer Science
Irvine, CA 92717

Dr. Jill Larkin
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

Dr. Paul E. Lehner
PAR Technology Corp.
7626 Jones Branch Drive
Suite 170
McLean, VA 22102

Dr. Alan M. Lesgold
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

Dr. Clayton Lewis
University of Colorado
Department of Computer Science
Campus Box 430
Boulder, CO 80309

Library
Naval War College
Newport, RI 02940

Library
Naval Training Systems Center
Orlando, FL 32813

Dr. Charlotte Linde
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025

Dr. Marcia C. Linn
Lawrence Hall of Science
University of California
Berkeley, CA 94720

Dr. Sandra P. Marshall
Dept. of Psychology
San Diego State University
San Diego, CA 92182

Dr. Richard E. Mayer
Department of Psychology
University of California
Santa Barbara, CA 93106

Dr. Arthur Melmed
U. S. Department of Education
724 Brown
Washington, DC 20208

Dr. James R. Miller
MCC
9430 Research Blvd.
Echelon Building #1, Suite 231
Austin, TX 78759

Dr. Mark Miller
Computer Thought Corporation
1721 West Plano Parkway
Plano, TX 75075

Dr. William Montague
NPRDC Code 13
San Diego, CA 92152-6800

Dr. Donald A. Norman
Institute for Cognitive Science
University of California
La Jolla, CA 92093

Director, Training Laboratory,
NPRDC (Code 05)
San Diego, CA 92152-6800

Director, Manpower and Personnel
Laboratory,
NPRDC (Code 06)
San Diego, CA 92152-6800

Director, Human Factors
& Organizational Systems Lab,
NPRDC (Code 07)
San Diego, CA 92152-6800

Library, NPRDC
Code P201L
San Diego, CA 92152-6800

Dr. Harold F. O'Neil, Jr.
School of Education - WPH 801
Department of Educational
Psychology & Technology
University of Southern California
Los Angeles, CA 90089-0031

Dr. Stellan Ohlsson
Learning R & D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

Office of Naval Research,
Code 1142EP
800 N. Quincy Street
Arlington, VA 22217-5000

Office of Naval Research,
Code 1142PT
800 N. Quincy Street
Arlington, VA 22217-5000
(6 Copies)

Director, Technology Programs,
Office of Naval Research
Code 12
800 North Quincy Street
Arlington, VA 22217-5000

Psychologist
Office of Naval Research
Branch Office, London
Box 39
FPO New York, NY 09510

Special Assistant for Marine
Corps Matters,
ONR Code OOMC
800 N. Quincy St.
Arlington, VA 22217-5000

Psychologist
Office of Naval Research
Liaison Office, Far East
APO San Francisco, CA 96503

Assistant for MPT Research,
Development and Studies
OP 01B7
Washington, DC 20370

Dr. Judith Orasanu
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Prof. Seymour Papert
20C-109
Massachusetts Institute
of Technology
Cambridge, MA 02139

CDR R. T. Parlette
Chief of Naval Operations
OP-112G
Washington, DC 20370-2000

Dr. Douglas Pearce
DCIEM
Box 2000
Downsview, Ontario
CANADA

Dr. Virginia E. Pendergrass
Code 711
Naval Training Systems Center
Orlando, FL 32813-7100

Dr. Nancy Pennington
University of Chicago
Graduate School of Business
1101 E. 58th St.
Chicago, IL 60637

Military Assistant for Training and
Personnel Technology,
OUSD (R & E)
Room 3D129, The Pentagon
Washington, DC 20301-3080

Dr. Ray Perez
ARI (PERI-11)
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. David N. Perkins
Educational Technology Center
337 Gutman Library
Appian Way
Cambridge, MA 02138

Dr. Tjeerd Plomp
Twente University of Technology
Department of Education
P.O. Box 217
7500 AE ENSCHEDE
THE NETHERLANDS

Dr. Martha Polson
Department of Psychology
Campus Box 346
University of Colorado
Boulder, CO 80309

Dr. Peter Polson
University of Colorado
Department of Psychology
Boulder, CO 80309

Dr. Joseph Psotka
ATTN: PERI-1C
Army Research Institute
5001 Eisenhower Ave.
Alexandria, VA 22333

Dr. Lynne Reder
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Fred Reif
Physics Department
University of California
Berkeley, CA 94720

Dr. Lauren Resnick
Learning R & D Center
University of Pittsburgh
3839 O'Hara Street
Pittsburgh, PA 15213

Dr. Gil Ricard
Mail Stop C04-14
Grumman Aerospace Corp.
Bethpage, NY 11714

Dr. Mary S. Riley
Program in Cognitive Science
Center for Human Information
Processing
University of California
La Jolla, CA 92093

Dr. Linda G. Roberts
Science, Education, and
Transportation Program
Office of Technology Assessment
Congress of the United States
Washington, DC 20510

Dr. William B. Rouse
Search Technology, Inc.
25-b Technology Park/Atlanta
Norcross, GA 30092

Dr. Walter Schneider
Learning R&D Center
University of Pittsburgh
3839 O'Hara Street
Pittsburgh, PA 15280

Dr. Alan H. Schoenfeld
University of California
Department of Education
Berkeley, CA 94720

Dr. Janet Schofield
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15280

Dr. Marc Sebrechts
Department of Psychology
Wesleyan University
Middletown, CT 06475

Dr. Judith Segal
OERI
555 New Jersey Ave., NW
Washington, DC 20208

Dr. Sylvia A. S. Shatto
OERI
555 New Jersey Ave., NW
Washington, DC 20208

Dr. Ben Shneiderman
Dept. of Computer Science
University of Maryland
College Park, MD 20742

Dr. Lee Shulman
Stanford University
1040 Cathcart Way
Stanford, CA 94305

Dr. Randall Shumaker
Naval Research Laboratory
Code 7510
4555 Overlook Avenue, S.W.
Washington, DC 20375-5000

Dr. Robert S. Siegler
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Vila M Simulia
Functional Technology
Systems Area

Eisenhower Avenue
Alexandria, VA 22333

J. Wallace Sinsiko
Power Research
and Advisory Services
African Institution
North Pitt Street
Alexandria, VA 22314

Derek Sleeman
Ford University
School of Education
Ford, CA 94305

Edward E. Smith
Bolt Beranek & Newman, Inc.
Moulton Street
Bridge, MA 02138

Richard E. Snow
Department of Psychology
Ford University
Ford, CA 94308

Cathryn T. Spoehr
University
Department of Psychology
Providence, RI 02912

Robert Sternberg
Department of Psychology
University
1A, Yale Station
New Haven, CT 06520

Robert Stevens
Bolt Beranek & Newman, Inc.
Moulton St.
Bridge, MA 02238

Paul J. Slichter
Senior Staff Scientist
Engineering Research Division
0
S. Washington
Alexandria, VA 22314

Dr. Thomas Slichter
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. John Tangney
AFOSR/NL
Bolling AFB, DC 20332

Dr. Robert P. Taylor
Teachers College
Columbia University
New York, NY 10027

Dr. Perry W. Thorndyke
FMC Corporation
Central Engineering Labs
1185 Coleman Avenue, Box 580
Santa Clara, CA 95052

Dr. Douglas Towne
Behavioral Technology Labs
1845 S. Elena Ave.
Redondo Beach, CA 90277

Dr. Paul Twohig
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Headquarters, U. S. Marine Corps
Code MPI-20
Washington, DC 20380

Dr. Beth Warren
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

Dr. Barbara White
Bolt Beranek & Newman, Inc.
10 Moulton Street
Cambridge, MA 02238

LCDR Cory deGroot Whitehead
Chief of Naval Operations
CP-112G1
Washington, DC 20370-2000

Dr. Michael Williams
IntelliCorp
1975 El Camino Real West
Mountain View, CA 94040-2218

A. E. Winterbauer
Research Associate
Electronics Division
Denver Research Institute
University Park
Denver, CO 80208-0454

Dr. Robert A. Wisher
U.S. Army Institute for the
Behavioral and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Martin F. Wiskoff
Navy Personnel R & D Center
San Diego, CA 92152-6800

Dr. Wallace Wulfeck, III
Navy Personnel R&D Center
San Diego, CA 92152-8800

Dr. Joe Yasutake
AFHRL/LRT
Lowry AFB, CO 80230

Dr. Joseph L. Young
Memory & Cognitive
Processes
National Science Foundation
Washington, DC 20550