

DOCUMENT RESUME

ED 296 712

IR 013 395

AUTHOR White, Barbara Y.; Frederiksen, John R.
TITLE Causal Model Progressions as a Foundation for Intelligent Learning Environments.
INSTITUTION BBN Labs, Inc., Cambridge, MA.
SPONS AGENCY Army Research Inst. for the Behavioral and Social Sciences, Alexandria, Va.; Office of Naval Research, Washington, D.C. Personnel and Training Branch.
REPORT NO BBN-Lab-R-6686
PUB DATE Nov 87
CONTRACT MDA-903-87-C-0545; N00014-82-C-0580
NOTE 68p.
PUB TYPE Reports - Research/Technical (143)

EDRS PRICE MF01/PC03 Plus Postage.
DESCRIPTORS *Artificial Intelligence; Cognitive Processes; Computer Assisted Instruction; Computer Simulation; *Expert Systems; *Instructional Design; Instructional Development; *Learning Strategies; Models
IDENTIFIERS *Causal Models; *Learning Environment

ABSTRACT

This paper describes the theoretical underpinnings and architecture of a new type of learning environment that incorporates features of microworlds and of intelligent tutoring systems. The environment is based on a progression of increasingly sophisticated causal models that simulate domain phenomena, generate explanations, and serve as student models. Constraints on model evolution are discussed in terms of causal consistency and learnability, and a taxonomy of models useful for instruction is outlined. The design principles underlying the creation of one type of causal model (zero-order models for electrical circuit behavior) are given, and possible progressions with respect to model elaboration, order, and perspective are described in the context of presenting a theory of model evolution. Finally, the architecture that enables the pedagogical tools of the intelligent learning environment is described, with an emphasis on the range of instructional interactions and learning strategies that can be supported. Eight figures are included. (51 bibliographic notes/references) (EW)

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

ED296712

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.

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

Report No. 6686

Causal Model Progressions as a Foundation for Intelligent Learning Environments

Barbara Y. White and John R. Frederiksen

November 1987

Prepared for:

The Army Research Institute

IR013395

REPORT DOCUMENTATION PAGE

Fcirm Approved
OMB No. 0704-0188

1a. REPORT SECURITY CLASSIFICATION Unclassified		1b. RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY		3. DISTRIBUTION/AVAILABILITY OF REPORT approved for public release; distribution unlimited.	
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE		Reproduction in whole or in part is permitted for any purpose of the U.S. Government	
4. PERFORMING ORGANIZATION REPORT NUMBER(S) BBN Laboratories Report #6686		5. MONITORING ORGANIZATION REPORT NUMBER(S)	
6a. NAME OF PERFORMING ORGANIZATION BBN Laboratories	6b. OFFICE SYMBOL (if applicable)	7a. NAME OF MONITORING ORGANIZATION Army Research Institute	
6c. ADDRESS (City, State, and ZIP Code) 10 Moulton Street Cambridge, MA 02238		7b. ADDRESS (City, State, and ZIP Code) 5001 Eisenhower Avenue Alexandria, VA 22333-5600	
8a. NAME OF FUNDING/SPONSORING ORGANIZATION Army Research Institute	8b. OFFICE SYMBOL (if applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER MDA-903-87-C-0545	
8c. ADDRESS (City, State, and ZIP Code)		10. SOURCE OF FUNDING NUMBERS	
		PROGRAM ELEMENT NO.	PROJECT NO.
		TASK NO.	WORK UNIT ACCESSION NO.
11. TITLE (Include Security Classification) Causal Model Progressions 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 Nov. 86 TO Nov. 87	14. DATE OF REPORT (Year, Month, Day) November, 1987	15. PAGE COUNT 60
16. SUPPLEMENTARY NOTATION			
17. COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP	
		AI and Education	
		Intelligent Tutoring Systems	
		Cognitive Modelling	
		Qualitative Models	
19. ABSTRACT (Continue on reverse if necessary and identify by block number)			
<p>This paper describes the theoretical underpinnings and architecture of a new type of learning environment that incorporates features of microworlds and of intelligent tutoring systems. The environment is based on a progression of increasingly sophisticated, causal models that simulate domain phenomena, generate explanations, and serve as student models. Constraints on model evolution, in terms of causal consistency and learnability, are discussed and a taxonomy of models, useful for instruction, is outlined. The design principles underlying the creation of one type of causal model are then given (for zero-order models of electrical circuit behavior); and possible progressions with respect to model elaboration, order, and perspective are described in the context of presenting a theory of model evolution. Finally, the simple architecture that enables the powerful pedagogical tools of the intelligent learning environment is described, with an emphasis on the range of instructional interactions and learning strategies that are supportable.</p>			
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 Judith Orasanu		22b. TELEPHONE (Include Area Code)	22c. OFFICE SYMBOL

Causal Model Progressions
as a Foundation for
Intelligent Learning Environments

Barbara Y. White and John R. Frederiksen

BBN Laboratories

Table of Contents

1. Introduction	1
1.1. Mental Models	3
1.1.1. The Importance of Qualitative Reasoning	4
1.1.2. Causal Consistency	5
1.1.3. An Overview of the Qualitative Model	6
1.2. Learning as a Process of Model Transformation	7
1.2.1. The Perspective of a Model	7
1.2.2. The Order of a Model	8
1.2.3. The Degree of Elaboration of a Model	9
1.3. An Overview of the Learning Environment	9
1.4. An Overview of the Paper	11
2. The Models	12
2.1. The Instructional Need for Zero-Order Models	12
2.2. The Zero-Order Models	13
2.2.1. Circuit Topology	14
2.2.2. Device Models	14
2.2.3. Circuit Principles Employed	17
2.2.4. The Device-Centered Propagation	18
2.2.5. The Process-Centered Propagation	21
2.3. Model Limitations	22
2.4. Model Strengths	24
2.5. The Troubleshooting Expert	25
2.5.1. Facilitating Troubleshooting	26
2.6. Problem Solving Capabilities	27
3. Model Evolutions	29
3.1. The Problem of Modifiability	30
3.2. Types of Micro Model Evolutions	31
3.3. Evolutions in Model Order and Degree of Elaboration	33
3.3.1. The Initial Degree of Elaboration: Reasoning about Voltages	33
3.3.2. The Second Degree of Elaboration: Reasoning about Current Flow	36
3.4. Developing Alternative Perspectives	39
3.4.1. Functional Models	39
3.4.2. Reductionistic, Physical Models	41
3.5. Coordinating Multiple Models	42
4. The Learning Environment	43
4.1. Interacting with the Learning Environment	45
4.2. Instructional Effectiveness	47
5. Summary	49

To appear in Artificial Intelligence

Causal Model Progressions as a Foundation for Intelligent Learning Environments

Barbara Y. White and John R. Frederiksen

BBN Laboratories, 10 Moulton Street, Cambridge, MA 02238, U.S.A.

ABSTRACT

AI research in qualitative modelling makes possible new approaches to teaching people about science and technology. We are exploring the implications of this work for the design of intelligent learning environments. The domain of application is electrical circuits, but the approach can be generalized to other subjects. Our prototype instructional system is based upon a progression of qualitative models of electrical circuit behavior. These models enable the system to simulate circuit behavior and to generate causal explanations. They also serve as target mental models for the learner. The model progression is used to create problem sets that motivate successive refinements to the students' mental models. Acquisition of these models allows students, at all stages of learning, to solve interesting problems, such as circuit design and troubleshooting problems. The system enables students to employ different learning strategies and to manage their own learning. For instance, they can create and experiment with circuits, can attempt problems posed by the system, and can ask for feedback and coaching from the models. In pilot trials, the learning environment successfully taught novices to troubleshoot and to mentally simulate circuit behavior.

1. Introduction

Our objective has been to create a powerful intelligent learning environment that enables students to utilize a wide range of learning strategies as they seek to develop expertise within a domain. We wanted to develop a system that would support not only learning from explanations, but also learning by discovery. The domain of instruction we selected is electrical circuits and the laws that govern their behavior. However, the architecture of the learning environment could readily be generalized to other domains.

To create the intelligent learning environment, we designed a progression of qualitative, causal

models of electrical circuit behavior that represent a transition from naivete to expertise. These qualitative models serve to drive circuit simulations, and to generate causal explanations of circuit behavior. They are implemented within the instructional system as a set of "articulate models" that can visually simulate and verbally explain what happens in a circuit when, for example, a switch is closed. When students interact with the computer, they are playing a "guess my model" game in that they are attempting to induce the computer's current model. After they have assimilated a given model, students move on to the next, more sophisticated, model in the sequence. Each such model progression is used to define a problem set that fosters that particular refinement in the student's mental model¹. The central feature of the approach is that, at each stage of learning, the model driving the computer simulation represents the mental model that the student is to acquire.

This new type of instructional environment synthesizes features of intelligent tutoring systems with features of microworlds. As a microworld, it provides for open-ended exploration within an interactive simulation, enables one to discover its laws. As a tutoring environment, it provides facilities for generating explanations, for modelling how to solve problems, for sequencing problems based upon the student's evolving mental model, and for providing intelligent feedback.

While the prototype system described in this article is restricted to helping students understand the behavior of electrical circuits, the approach can be applied to a large number of domains. For instance, any domain whose phenomena can be represented by laws affecting the behavior of objects can potentially be tutored using causal model progressions. This includes physical systems (e.g., Newtonian mechanics -- White & Horwitz, 1987), biological systems (e.g., the human heart -- Feltovich et al., 1987), or even economic systems. Causal modelling can also be applied to the learning of mathematical systems -- as in the work of Feurzeig and White (1983) that used an articulate causal model to help elementary school students understand place notation and its relationship to the standard algorithms for addition and subtraction. Further, this approach could be applied to learning about computers themselves: articulate causal models could help students understand, design, and troubleshoot computer hardware and software (such as when attempting to debug a computer program or to understand how a compiler works). Thus, we argue that causal model progressions have broad applicability as a foundation for intelligent learning environments.

¹By 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., mentally simulating the behavior of a circuit).

1.1. Mental Models

The theoretical framework we adopt is that expertise (in this case, electrical expertise) can be captured by a small set of mental models that embody alternative, but coordinated, conceptualizations of system 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; 1985). The models that we present to students are also influenced by studies of students' reasoning about circuit problems (e.g., Cohen, Ganiel, & Eylon, 1983; Collins, 1985), and also draw upon theoretical AI work on qualitative modelling (Brown & deKleer, 1985; Davis, 1983; deKleer, 1985; Forbus, 1985; Kuipers, 1985; Weld, 1983; Williams, 1985), particularly that of DeKleer & Brown (1983).

We chose mental models as the knowledge structures that we would try to impart to students for several reasons. First, it is our belief that students learn most effectively when they are actively applying their knowledge to solving problems. To support such learning-by-problem-solving, students must, at all times, maintain a knowledge structure that is complete enough to support problem solving. Thus, the purpose of the intelligent learning environment is to develop in students a progression of increasingly sophisticated mental models for reasoning about circuit behavior, each capable of supporting problem solving for a subset of circuit problems.

Second, 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 for making predictions about the behavior of different circuits, for troubleshooting circuits, and for designing circuits. This is in contrast with, for example, troubleshooting knowledge represented as symptom-fault associations, which are context specific and not explicitly causal, and are, therefore, of limited use in helping students to understand how circuits work.

In addition to being efficient and powerful knowledge structures for students to possess, mental models 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. This simulation is capable, by simply reasoning out loud (via a speech synthesis device), of generating causal accounts for the behavior of a circuit that the student is creating and observing. For instance, the student can close a switch and see a light turn on and, at the same time, hear an explanation of how changes in the conductivity of the switch caused a change in the voltage applied to the light, which in turn caused the light to turn on. Thus, we argue that mental models can, if appropriately designed, enable both the instructional system and the student to reason from general principles and to generate causal accounts of circuit behavior.

We will argue that these models should initially be qualitative and able to generate causal accounts of the changes in devices' states that occur during the operation of a circuit. Thus, they should be capable of reasoning about device states and about the processes that lead to changes in those states. 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, with reductionistic, physical models of electricity, 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.

1.1.1. 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. 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 (diSessa, 1983), and (2) to enable them to acquire this important problem solving skill that evidence has shown they lack (Chi et al., 1981; Larkin et al., 1980). Quantitative reasoning should only be introduced after students have been given a qualitative, causal 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 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, causal reasoning for their solution. For instance, in the domain of electrical

circuits, problems involving the prediction of circuit behavior, circuit design, and troubleshooting all have this property. Problems of this type are thus useful in motivating the development of qualitative reasoning skills.

1.1.2. 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 in their form (e.g., $I=V/R$, $V=IR$, $R=V/I$). Instruction then focuses on how to apply these equations in their various forms in order to analyze problems involving circuits of varying degrees of complexity. In these problems, the resulting constraints on voltages and currents in a circuit are used to develop quantitative solutions for the unknown quantities in the problem. Note that by using such constraint-based reasoning, no explicit causal model is presented, 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. In addition, the order in which the constraint equations are applied when solving problems is governed by algebraic considerations rather than by a causal analysis. Thus the problem-solving process that students are taught does not necessarily facilitate an understanding of the physical system under study.

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 what he terms the KCL heuristic). At other times, an increase in voltage across a component causes the current through the component to increase (using an qualitative version of 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 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 the input voltage and 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). Furthermore, this view is consistent with students' models of mechanics, in which forces applied in a particular direction accelerate a moving body. Finally, Cohen, Eylon, & Ganiel (1983) argue that this is an important way of conceptualizing circuit behavior that even sophisticated students lack, as illustrated by the following quotation:

"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."

Qualitative reasoning should therefore be based initially upon a subset of the constraints available in quantitative circuit theory, chosen for their causal consistency. With electrical force 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 within the circuit. Changes in the conductivities of circuit components are thus taken as the independent variables.

1.1.3. An Overview of the Qualitative Model

Within the qualitative model embedded in our tutoring system, any change in conductivity of a circuit component invokes a voltage redistribution process. This process determines the voltages that occur throughout the circuit as a result of such a change in conductivity. Qualitative rules that govern this voltage redistribution process are related to well-known laws of quantitative circuit theory (Kirchhoff's Voltage Law and Ohm's Law). They can also be seen as emergent properties of a dynamic model of the movement of charge carriers within a circuit as it settles into a steady state (Haertel, 1987; Frederiksen & White, 1988). Once the distribution of voltages is determined, the system then considers the effects of changes in voltage on the states of components within the circuit. For example, in a series circuit containing a switch, a light bulb, and a battery, if one closes the switch, the voltage across the switch becomes zero and that across the light bulb becomes nonzero. Then, since the light bulb has a voltage applied to it, it changes its state from "off" to "on". One can envision cases where the change in state of a component may also have an effect on its conductivity (such as when a transistor becomes saturated). This conductivity change in turn leads to another application of the voltage redistribution process, and so on. For example, if a switch is suddenly closed, it may cause a transistor to become saturated, which in turn could cause a light to go on. The sequence of propagation cycles continues until the circuit behavior stabilizes.

Simulating circuit behavior in this way, through the use of 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. This allows the student to see how the behavior of devices is causally determined by changes in other devices' states, through the effects of those changes on the voltage distribution within the circuit. Understanding the causality of circuit behavior thus motivates the need to understand basic circuit concepts such as conductivity and voltage, and also basic circuit principles such as Kirchhoff's Voltage Law. These are non-trivial concepts and laws to master, so we have taken 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.

1.2. 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.

Transformations may involve the addition, modification, differentiation, or generalization of model features, or even the construction of alternative 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. Other model transformations involve changes in a model's control structure. For instance, initially we focus on the behavior of a single device in a circuit, such as a light bulb, 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 within the circuit.

The representation of the learner's current knowledge state is a description of the set of models he or she currently has evolved (Frederiksen & White, in press). This representation, in turn, characterizes the types of problems that the learner can currently solve. Our work has focused on creating upward progressions of increasingly sophisticated models for reasoning about the behavior of electrical circuits. We have also worked on lateral progressions that focus on alternative means for understanding circuit operation, such as models of circuit teleology and reductionistic, physical models of circuit behavior. All of these models furnish learning objectives for different stages in instruction.

For purposes of creating model progressions, we find it useful to define three dimensions on which models may vary: their perspective, their order, and their degree of elaboration. Lateral progressions, that represent alternative means of understanding the domain, involve changes in model perspective. Upward progressions to more sophisticated models involve changes in model order and degree of elaboration.

1.2.1. The Perspective of a Model

The "perspective" of a model refers to the nature of the model's reasoning in explaining a circuit's operation. For instance, it could be reasoning about:

1. The overall functionality of the circuit -- e.g., its purpose and how subsystems within the circuit work together to achieve that purpose -- i.e., functional models;
2. The behavior of the circuit at the level of circuit components -- e.g., how changes in the

state of one device can cause changes in the states of other devices -- i.e., behavioral models; or

3. The behavior of the circuit at a more micro level -- e.g., how electrical charges are redistributed across a resistor when the voltage drop across the resistor changes -- i.e., reductionistic, physical models.

In the first type of model, we represent devices as functional units that transmit and process information. The propagation mechanism involves information flow and takes into account functional dependencies among devices. In the second model, devices are represented as potential charge carriers (conductors) and sources of voltage, and the propagation mechanism utilizes qualitative, causal versions of basic circuit laws, such as Ohm's Law and Kirchhoff's Voltage Law, to reason about the electrical interactions among devices. In the third type of model, devices are represented at the molecular level in terms of their effects on the behavior of charge carriers, and the propagation mechanism is in terms of electrical forces and Coulomb's Law. Each of these models is qualitative and causal, yet each models a device from a different perspective, and propagates the effects of changes in device states by utilizing different laws (i.e., functional dependence versus Kirchhoff's Voltage Law versus Coulomb's Law). They are thus focusing on modeling different aspects of the domain phenomena.

1.2.2. The Order of a Model

In the context of the behavioral models, a further subdivision in this typology of models can be made. We distinguish models that reason on the basis of the mere presence or absence of resistance, voltage, or current, which we call "zero-order models", from those that reason on the basis of incremental 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 there a voltage drop across the light in this circuit and, consequently, is the light off or on?" First-order models on the other hand reason on the basis of qualitative (first-order) derivatives and can answer questions such as, "Is there an increase in voltage across a component when we decrease the resistance of some other component?" Each of these is distinguished from quantitative models that can answer questions of the form, "What is the voltage across two points in a circuit?" All of these orders of model are thus useful for answering questions about circuit behavior of a particular sort².

For example, from the standpoint of the zero-order model, a simple one-transistor amplifier behaves like a switch, since the transistor has only two states: unsaturated (non-conductive) and saturated (conductive). Viewed within a first-order model, however, changes in the input voltage to the transistor cause incremental changes to the output voltage of the amplifier. This enables one to understand in qualitative terms the analogue behavior of the amplifier. From the viewpoint of a

²Zero-order models, for example, are sometimes taught as a basis for learning to troubleshoot electrical circuits (White and Frederiksen, 1984; 1985).

quantitative model, one can ascertain the range of input voltages for which this incremental behavior of the transistor remains true (that is, over what range of voltages the amplifier is linear). One could also synthesize zero- and first-order models by introducing the notion of boundary conditions on the input voltage to the transistor, in order to distinguish regions where the transistor follows the qualitative incremental model and those where it is fixed as unsaturated or saturated. The complexity of models which incorporate boundary conditions suggests that they should be reserved for problems in which such aspects of circuit behavior are important. Over the course of instruction, one seeks to develop an understanding of zero- and first-order models, and then such principles as boundary conditions which allow their integration. At the same time, students should develop a knowledge of the classes of problems for which each type of model is best suited.

1.2.3. The Degree of Elaboration of a Model

During learning, qualitative models developed can 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 states of circuit components on the behavior of other components. The qualitative rules employed are drawn from the repertoire of constraints of quantitative circuit theory, and initially focus on principles relating resistances to voltages within a circuit. 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 (particularly, Kirchhoff's Voltage Law). In subsequent models, further qualitative rules are introduced (based on Ohm's Law and Kirchhoff's Current Law) to allow inferences about current flow within a circuit (as described in section 3.3).

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, students must learn to apply flexibly the full set of constraints available in circuit theory, and to reason "algebraically" in finding and applying multiple constraints.

1.3. An Overview of the Learning Environment

Our QUEST (short for: Qualitative Understanding of Electrical System Troubleshooting) intelligent learning environment is a descendant of the SOPHIE (Brown et al., 1982) and STEAMER (Forbus & Stevens, 1981) intelligent tutoring systems: It makes use of interactive simulations, qualitative explanations, and a troubleshooting expert. However, unlike SOPHIE and STEAMER, QUEST is based upon a progression of qualitative models that correspond to the desired evolution of the learner's mental model. Furthermore, the interactive simulation of circuit behavior is driven by these qualitative models, and not by a quantitative model. Also, the system is capable of generating a runnable, qualitative, causal

model for any circuit that the student or instructional designer might create (within the limits discussed in the next section). Thus students can, for example, use a circuit editor to create circuits and experiment with them. 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 the progression of models which serves to define classes of problems and facilitates the generation of explanations. Making use of this model progression, students can 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.

The instructional system thus 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). Students can also use the circuit editor to experiment with circuits by changing the states of devices, inserting faults, and adding or deleting components.

The objective is 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 the voltage redistribution processes operate, they can leave a visible trace of the path they are currently pursuing so that, for example, when they determine 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 allowing the student 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.

1.4. An Overview of the Paper

In Section 1, we presented a theory of expertise and its evolution: learning can be regarded as a process of acquiring a set of coordinated mental models. We argued that in the initial stages of instruction, the models should be qualitative and causal. Constraints on model evolution, in terms of causal consistency and learnability, were discussed and a taxonomy of models useful for instruction was outlined. We then went on to provide an overview of our prototype intelligent learning environment based upon a progression of qualitative, causal mental models.

Section 2 describes the intelligence underlying our intelligent learning environment: the zero-order models of electrical circuit behavior and the troubleshooting expert. In presenting these models, we illustrate the instructional principles that governed their design. We conclude by describing the range of problem types that a tutoring system, or a human, possessing such expertise can solve. The reader who is not interested in the details of the qualitative models may choose to skip Section 2. Section 3 illustrates possible progressions, with respect to model elaboration, order, and perspective, in the context of presenting a theory of model evolution. Section 4 describes the simple architecture that enables the powerful pedagogical tools of the intelligent learning environment. The emphasis is on the range of instructional interactions and learning strategies that are supportable. Instructional trials of the system are then briefly discussed. Section 5 presents a summary of the paper.

2. The Models

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 can introduce basic circuit concepts and principles, 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's (1985) EQUAL reasons in terms of qualitative derivatives obtained from qualitative versions of the constraint equations ("confluences") used in quantitative circuit analysis. These, together with several heuristics, enable it to analyze the effects of incremental change on circuit behavior. The difficulty with using such a model, at least at the initial stage of instruction, is that novices typically do not have well developed concepts of voltage or resistance, let alone of changes in voltages or resistance (Collins, 1985; Cohen, Eylon, and Ganiel, 1983). For example, as part of a trial of our instructional system, we interviewed seven high school students who had not taken a 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 has no resistance when closed did not matter. In fact, five of the students stated that they did not know what was meant by the term "resistance". They thus did not have the electrical concept of resistance and of how a non-resistive path could affect circuit behavior.

 insert Figures 1 & 2 about here

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: the presence or absence of voltages in a circuit and their effect on the states of devices. 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.

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.

 insert Figure 3 about here

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 model is no longer simply reasoning about whether or not there is a voltage applied to a device; rather, it reasons about whether the voltage will increase or decrease. 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 incremental changes. 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.

2.2. The Zero-Order Models

A progression of zero-order, behavioral models is implemented within our prototype intelligent learning environment. These models incorporate knowledge of (1) the topological structure of the circuit, (2) the behavior of the devices within the circuit (device models), and (3) 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, such as the effects of opening or closing a switch, on the voltages applied to other components of the circuit. The instructional system also includes a progression of general troubleshooting strategies for localizing faults within a circuit. These strategies

utilize the behavioral models as part of their problem solving process. Both the behavioral models and troubleshooting strategies can articulate their thinking, both visually and verbally, when simulating the behavior of a given circuit or when carrying out troubleshooting.

2.2.1. 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.

Circuit parsing and orientation. The zero-order models utilize an algorithm that can parse any circuit, based upon its topology, into series and parallel subcircuits. This algorithm recursively recognizes and groups all series and parallel subcircuits by first bracketing 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 no such subcircuits remain. In the case of bridge circuits, only parallel or series subcircuits within the circuit are grouped. In the case of series-parallel circuits, the final, top-level circuit is a series circuit. As the algorithm progresses, it also assigns conductivities to each of the subcircuits in the parse. The result is a hierarchical parsing of the circuit.

Establishing polarities of devices. Devices within the circuit are "oriented" by establishing polarities for each device and subcircuit. Polarities are assigned in relation to the voltage source, starting at the outermost grouping and moving inwards. Units contained within larger units are assigned the same polarities as those of the larger units which contain them.³ The main purpose of the circuit orientation is to facilitate the tracing of circuit paths and circuit loops. Since the circuits used in the initial stages of instruction are not very complex, and since humans have spatial skills that make circuit tracing generally quite effortless, circuit orientation is not explicitly taught until necessary.

2.2.2. Device Models

The zero-order models contain models for devices typically found in circuits. The devices modelled are batteries, switches, resistors, bulbs, diodes, fuses, capacitors, transistors, test lights, and wires. 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. The conditions for the rules that determine device states

³The algorithm can identify indeterminacies in the assignment of polarities to a unit. For instance, if a unit has both feed paths to the positive side of a voltage source and return paths to the negative side of a voltage source 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 indeterminant (see Ritter (1987)).

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. When a device's state changes, the device model activates additional rules which reevaluate a set of state variables associated with the device. These variables include (1) the conductivity of the device (is it purely conductive⁴, 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 models⁵ follow:

Battery

States: 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-Conductive

Voltage 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 permanently nil.

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

⁴The term "purely conductive" is introduced to represent the conductivity of devices such as wires, whose resistance is very much lower than that of other circuit components.

⁵The devices are modelled as ideal components. Thus, for example, the battery is modelled as purely conductive because an ideal battery as a pure voltage source has no resistance, even though, real world batteries are resistive.

increasingly sophisticated and adequate models for such devices⁶. The initial capacitor model is illustrated below. It should be noted that when a device state change occurs during a given causal propagation cycle, it takes some indeterminate amount of time. Thus when the capacitor goes from the charged state to the discharged state, it can serve as a voltage source during that propagation cycle as illustrated further in section 2.2.4.

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 its state is charged and if it has a conductive path across it, then its state becomes discharged; otherwise it remains charged.

Internal Conductivity:

If it is charged then it is non-conductive in the polarity with which it is charged, but it is purely conductive in the opposite polarity.

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 nil.

When a particular device, such as a capacitor, 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,

⁶The initial capacitor model only incorporates two discrete states: charged and discharged. One limitation of such a zero-order capacitor model is that it does not explicitly introduce 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 higher voltage battery, 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.

resistive, or nonconductive), (4) the device polarity, and (5) its fault status. When the student is performing a mental simulation of a particular circuit, the student must also keep track of this information. The information must be recorded, either above the device in the circuit diagram, in a device data table (as illustrated in Figure 5), or in memory. Most of the circuits we use to introduce basic circuit concepts and laws are simple enough that the student can remember this information -- especially since the student is typically focusing on the behavior of one device, such as a light bulb.

A mental model for a device, of the form illustrated for batteries and capacitors, enables the student to determine the state of the device for any circuit environment in which it is placed⁷. 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 or 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.

2.2.3. Circuit Principles Employed

Device models reevaluate their states on the basis of the voltages that are currently being applied to them in the simulation. Whenever a change occurs in a device's state, a change may also occur in the device's conductivity or its status as a voltage source. These changes in turn effect the distribution of voltages within the circuit and, consequently, the voltages that are being applied to other circuit devices. The qualitative simulation employs general circuit principles for propagating the effects of these changes in a device's state on the voltages applied to other devices. We have developed two alternative forms for carrying out this propagation: (1) a device-centered causal propagation in which devices effectively ask themselves, "What effect did device x's change in state have on me?", and (2) a process-centered causal propagation where any change in device state triggers an explicit voltage redistribution process. Forbus (1985) has argued for the utility of process-centered models in making the causality of physical systems more explicit. We argue that both propagation mechanisms have their utility in instruction, in problem solving, and in unpacking the causality of circuit behavior. For instance, the device-centered mechanism enables one to focus on the behavior of a single device, and by so doing to introduce basic concepts such as that of a circuit. In addition, when solving circuit problems such as in troubleshooting, this approach enables one to focus on the behavior of a single device such as a test light and to envision what possible faults would do to its behavior. On the other hand, the process-centered mechanism

⁷It 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.

enables one to see more explicitly the redistribution of voltages that can occur when a device changes state, and is more efficient when envisioning the behavior of large circuits.

2.2.4. The Device-Centered Propagation

In this control logic, whenever any device's internal conductivity or status as a voltage source changes, 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 devices⁸. 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. Whenever 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, the propagation of state changes commences once again.

Circuit principles. In order for a device model to reevaluate its state, it must establish whether or not it currently has a voltage applied to it. To do this, the device models call upon rules to determine, based upon the circuit topology and the states of devices, whether or not the device has a voltage applied to it. 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 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.⁹

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.

Topological search. The rules that embody circuit principles, such as the Zero-Order Voltage Rule

⁸The 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.

⁹By "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.

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.

A Sample Zero-Order Circuit Simulation. As an illustration of how this zero-order model reasons, consider a simulation of the behavior of the circuit illustrated in Figure 4:

 insert Figure 4 about here

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 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. This behavior is summarized in Figure 6.

 insert Figures 5 & 6 about here

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 et al. (1984) term 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 must ... (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 causal propagation becomes apparent to the student.

When the student is attempting to predict the behavior of a single device within a circuit such as a test light, it is sometimes 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 can be necessary.

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, a serial reevaluation of device states is all that is necessary. 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 simulate parallelism the way the computer model does by using device data tables as shown in Figure 5. Students 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. By circling the devices that change state on each propagation cycle (i.e., each column of the table), the sequence of state changes for the circuit can become clear as illustrated in Figures 5 & 6.

The device-centered propagation becomes lengthy for large circuits. A second approach that is

more efficient, and more direct in terms of the causality of circuit behavior, is the process-centered approach described below.

2.2.5. The Process-Centered Propagation

In this approach, the circuit principles are embedded in a process called voltage redistribution. This uses rules based upon Kirchhoff's Voltage Law to immediately propagate "forward" the effects of a change in conductivity of a device on voltages applied to other devices in the circuit. Then, only those other devices whose applied voltages have changed will undergo reevaluation. Furthermore, since it will already have been established whether or not there is a voltage being applied to each device, redundant searches for paths to a voltage source are avoided. This approach of immediately propagating the effects of changes in a device's conductivity on the distribution of voltages is more closely related to the causality of our reductionistic, physical model, which focuses on the voltage redistribution process.

Rules for circuit loops. The propagation logic for redistributing voltages within a circuit begins by applying rules for assigning voltages within circuit loops containing a voltage source. These loops are those in the top-level of the circuit parse, and are typically made up of a combination of devices and bracketed subcircuits (together referred to as components of the loop). If there are no nonconductive components in the loop (i.e., the loop is conductive), voltage drops are assigned to all resistive components. While no voltage drops occur across purely conductive components¹⁰, they nonetheless have voltages applied to them. These rules apply to conductive circuit loops. If, on the other hand, the circuit loop contains one or more nonconducting components, a voltage drop is assigned only across the nonconductive segment of the circuit.

Rules for subcircuits. Next, if any components of the loop to which voltages have been assigned consist of subcircuits, then voltages remain to be redistributed within those subcircuits. Rules for doing this depend upon the conductivity of the subcircuit. For example, if the subcircuit contains a purely conductive path, then no components of the subcircuit will have voltage drops and only components of the purely conductive path within the subcircuit will have voltages applied to them. On the other hand, if the subcircuit is resistive, then voltage drops are assigned to all resistive components within its conductive paths. For any nonconductive paths, the voltage drop will occur across the nonconductive segment of the path. Similar rules are applied if the subcircuit is nonconductive.

Control structure. Whenever a device changes its conductivity, one first updates the conductivities of any subcircuits containing that device. Then the rules for assigning voltages within circuit loops are

¹⁰Since devices having very low resistance are regarded as purely conductive within the zero-order model, they will have no voltage drop across them even when they are in a circuit loop in which they have a voltage applied to them. For purely resistive devices, it is therefore necessary to keep track of applied voltages. For devices that are resistive or nonconductive, there will always be an applied voltage whenever there is a voltage drop across them.

applied for all loops containing components whose conductivity (or status as a voltage source) has changed. Next, voltages are redistributed within subcircuits of those loops for which the voltage has changed. If the result of the above voltage redistribution process causes a change in voltage across a component of some intersecting circuit loop, then a further voltage redistribution is carried out within that loop, and so on. Finally, once the voltages have been redistributed, those devices for which there have been changes in applied voltages are prompted to reevaluate their states. If any of these change their conductivity or status as a voltage source, another propagation cycle begins. Propagation ceases when the circuit behavior stabilizes.

An example of circuit simulation using voltage redistribution. As an illustration of how this zero-order model reasons, we will apply it to the same circuit analyzed earlier (see Figure 4).

We will begin with the case where switch #2 is open and switch #1 is closed. First, voltages are distributed within the series circuit loop made up of the battery, the resistor, the parallel subcircuit (containing the light bulb, connected across the capacitor and switch #2), and switch #1. Since all of these three circuit components are conductive, voltage drops are assigned across the two resistive components (the resistor and the subcircuit). Switch #1 has no voltage drop across it, since it is regarded as purely conductive, although it does have a voltage applied to it. Second, voltages are assigned within the parallel subcircuit. Since the subcircuit has a voltage across it, resistive components within its conductive path (here, the light bulb) also have voltages across them. However, within the nonconductive path, the voltage drop appears only across the nonconductive segment (here switch #2) and no voltage is applied to the capacitor. Since the light bulb has a voltage applied to it, it enters the "on" state.

Next, we will suppose that switch #2 is closed. Since there has been a change in its conductivity, the conductivity of the subcircuit containing it has to be revised. Since the capacitor (which is initially discharged) and switch #2 are both purely conductive, the subcircuit contains a purely conductive path and is therefore also purely conductive. Within the circuit loop containing the battery, the subcircuit now has no voltage drop across it (since it is purely conductive), although it has a voltage applied to it. This implies that, within the subcircuit, any circuit paths connected across the purely conductive path cannot have voltages applied to them, and only components within the purely conductive path can have voltages applied to them. Changes in voltages thus occur for two devices: the light bulb and the capacitor. These devices therefore are prompted to reevaluate their states. The light bulb no longer has a voltage applied to it and goes from the "on" to the "off" state (but does not change in conductivity), while the capacitor now has a voltage applied to it and goes from the discharged to the charged state. This causes it to change in conductivity from purely conductive to nonconductive (it also becomes a source of voltage).

These changes in state of the capacitor cause another redistribution of voltages within the circuit. First, since the subcircuit containing the capacitor changes from being purely conductive to resistive, and since it remains in a conductive loop with the battery, it now has a voltage across it. Consequently, within the subcircuit there are now voltage drops across the capacitor and light bulb, neither of which is purely conductive. Second, since the capacitor is now a voltage source, and since it is in a conductive loop with the light bulb and switch #2, it will also apply a voltage to the light bulb. Since the polarity of this voltage drop is the same as that previously assigned, there is no ambiguity and the voltage can be assigned to the light bulb. Finally, since the light bulb now has a change in its voltage, it reevaluates its state, this time going from off to on. As this state change does not alter its conductivity, no further changes in the distribution of voltages occur and the circuit behavior stabilizes.

2.3. Model Limitations

Two limitations of these qualitative models emerge when multiple events occur on the same propagation cycle. For example, imagine that on a given cycle a transistor becomes saturated and a capacitor becomes charged. The first difficulty relates to determining the order of the events and the second relates to determining the cause of subsequent events.

Time. A major limitation of qualitative models is that the sequencing of events happens in ordinal, not interval, time. That is, the state changes happen in a given order, but the length of time such events take is indeterminant. For instance, if closing a switch causes a transistor to now have a voltage drop across its base-emitter path, how long does it take for the transistor to become saturated? Is it an instant or a relatively long time? The simulation has no way of knowing. Further, imagine that closing that same switch causes a voltage drop not only across a transistor, but also across a capacitor. According to the model, the transistor becomes saturated and the capacitor becomes charged. Both events are caused by closing the switch and happen on the same propagation cycle. However, the two events do not necessarily take the same amount of time. Thus, it could be the case that the transistor becomes saturated before the capacitor becomes charged, or vice versa, or both new states could be reached at the same time. To further complicate matters, for certain circuits, the subsequent behavior will vary depending upon which of these possibilities occurs. Thus, accurately simulating their behavior will depend crucially on distinguishing between these possibilities. In such cases, all zero-order, qualitative models can do is to articulate the range of possible behaviors for the circuit. If the capacitor becomes charged at a certain point with respect to the behavior of the transistor, the circuit will exhibit one behavior, whereas, if it becomes charged at another point, the circuit will exhibit a different behavior. The student, or system, must then use knowledge about the purpose of the circuit, knowledge about the relative times for events to occur, or quantitative models to determine what is the likely behavior for such a circuit.

Cause. This multiplicity of events occurring within a given propagation cycle raises yet another model limitation. Imagine, for the sake of argument, that the transistor and capacitor in the above example reach their new states at the same time, or, alternatively, that in another circuit, someone closes two switches simultaneously. Further, imagine that on the next propagation cycle for either circuit, another device, such as a light bulb, changes state. In each case, was it the change in state of device #1, or device #2, or both that caused the light bulb to change state? The device-centered propagation could unpack the causality by keeping histories of feed and return paths for each device. The model could then investigate these histories to determine whether the state change for device #1 and/or device #2 completed or destroyed a feed and/or return path for the light bulb. Alternatively, one could hypothesize that only device #1 or device #2 changed state and then envision whether the test light would have changed state under either of these conditions.

These limitations and complexities of zero-order, qualitative models in producing the correct sequence of state changes and unpacking circuit causality are serious from a simulation perspective. Yet, from an instructional perspective they are not. It is the limits of particular models that help to define pedagogically appropriate classes of problems. For instance, when designing sequences of circuit problems to tutor the concepts of resistance, voltage, and the voltage redistribution process, it is not

necessary to introduce circuits so complex as to broach these complexities. In fact, to do so would be poor instructional design. This again argues for the utility of qualitative models, not only for introducing fundamental circuit concepts and principles, but also for being able to define the problem context in which they should be introduced.

2.4. Model Strengths

We sought models that would be powerful enough to enable the system and the students to envision and explain circuit behavior. We also 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. Finally, unlike deKleer (1985), we wanted to avoid making assumptions about the integrity of the circuit and, therefore, to avoid running into contradictions during the envisioning process.

No-function-in-structure. By utilizing context-free models for devices along with circuit principles for determining the distribution of 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) in whatever circuit they are placed. 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 our models are in keeping with deKleer and Brown's (1983) no-function-in-structure principle.

Creating knowledge structures that follow this principle 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. It is also an important property for the students' mental models in that it ensures that their knowledge will then be in a form that enables them to understand and predict the behavior of any circuit.

In addition to obeying the no-function-in-structure principle, the model's control structure, circuit principles, and the format for device models are general and compatible with higher order models. Thus, when students are introduced to a new type of circuit (e.g., logic circuits) incorporating new types of devices (e.g., logic gates), they should be able to generalize their knowledge to this new situation. Moreover, they should be able to transfer their knowledge of model form and process to the acquisition of new models of circuit behavior (such as first order models that involve incremental qualitative reasoning). Generality thus enhances the learnability of new models that are extensions of the earlier models.

Finally, one of the most impressive features of the type of qualitative, causal models 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, to describe the behavior of the entire circuit, or to determine what faults are possible given the behavior of the circuit. Further, students can be asked to locate a faulty component within a circuit, or to design a circuit. The ability to perform this type of model-based simulation of circuit behavior is instrumental in solving all of these types of problems.

2.5. The Troubleshooting Expert

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.

The progression of troubleshooting strategies 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 methods he uses are based upon the fundamental idea of a circuit, and employ circuit principles that are similar to those of the zero-order model (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. In order for there to be voltage applied to a device, there must be a source of voltage, and conductive paths leading from each port of the device to, respectively, the positive and negative sides of the voltage source. In a series circuit, one source of faults is the occurrence of opens within either of these paths. 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 (usually positive) side of the voltage source, they will prevent voltages from being applied to the device. These two types of faults, opens and shorts to ground, are the ones that the troubleshooting expert is designed to diagnose.

The strategy used by the expert is a "space splitting" strategy. Its goal is to divide the circuit into two parts and then to infer which portion of the circuit contains the fault. Once the fault has been isolated to one or the other segment, then the troubleshooting logic is recursively applied to the faulty segment. The process continues until the fault has been localized.

This top level strategy is accomplished using the following approach: First, for purposes of measurement and inference, the circuit is logically divided into two parts by inserting a test light near the center of the circuit (i.e., between the test point and the grounded (negative) side of the voltage source). As can be seen in Figure 3, this places the test light in series with one segment of the circuit (that between the positive side of the voltage source and the test point) and in parallel with the other segment

(that between the test point and the grounded side of the voltage source). Next, the circuit simulation is run to determine the correct state of the test light if the circuit had no fault. That state is then compared with the actual test light behavior, and inferences are made about possible faults that are consistent with the findings¹¹. Additional troubleshooting operations are then carried out to isolate the fault possibilities to either the portion of the circuit in series with the test light, or the one in parallel with it¹². When the fault possibilities have been isolated to within a particular segment of the circuit, the expert moves the test light to a new test point within the faulty segment of the circuit and reapplies the troubleshooting logic. In selecting the next test point, the troubleshooting expert may again apply the strategy of dividing the circuit in half, or it may instead employ a serial strategy of simply moving the test light, one component at a time, away from the earlier test point. 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 strategy presented to students increases progressively in complexity, in coordination with the progression of behavioral circuit models that the students acquire (see White & Frederiksen, 1986a).

2.5.1. 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. Each 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 capabilities of the student's (and the system's) present model of circuit behavior. 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 by the student into a circuit. In addition, any part of a 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 of the circuit accordingly. The availability of these facilities enables

¹¹These inferences depend 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 but is not on, the fault could be either an open or a short to ground 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).

¹²In the example shown in Figure 3, the expert detaches switch SW2 from N4, and observes the effect. If the test light comes on, the fault has been isolated to the portion of the circuit in parallel with the test light. More particularly, it was providing a non-resistive path from the feed path of the test light to the ground, and thus shorting out the test light. 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.

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 themselves plant a fault into a circuit and have the expert demonstrate how it would proceed to isolate it.

2.6. Problem Solving Capabilities

One of the most impressive features of an intelligent learning environment based upon causal models and problem solving experts is the range of problem types for which the system's (and the student's) expertise is applicable. The following describes the types of problems supportable by an environment based upon the causal models and troubleshooting expert described in the previous sections:

Predicting the effects of a change in state: envisioning circuit 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 mode' 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.

Determining the possible origins of a given behavior: 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 (as in Figure 3). 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.

Isolating the cause of an unexpected behavior: troubleshooting problems. The system 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, he or she is asked two questions (Feurzeig & Ritter, in press): (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.

Restructuring in order to achieve a particular behavior: circuit design and modification problems. The student is asked to create, using the circuit construction kit, 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 students 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.

Creating models of system behavior: 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.

3. 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). 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 this section we will characterize the ways in which models can evolve. The major evolutions can be broadly classified using the taxonomy developed in the introduction: (1) an increase in degree of elaboration, (2) order, and (3) a change in perspective. In addition, there are finer grained evolutions that can occur within a model of a particular perspective, order, and degree of elaboration -- such as conceptual refinement, generalization, or differentiation. This taxonomy of model evolutions will allow us to create a space of causal model progressions, and will ultimately allow for different learners to pursue alternative paths through the space, depending upon their learning styles and pedagogical goals.

Some of the above model evolutions can involve a change in model form as well as content, that is, a transformation in the way knowledge is represented and applied. For example, when introducing the concept of a fault, one could simply add to each device model rules for altering that device's state variables when the device is faulty. Alternatively, one could introduce 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 form of propagation mechanism utilized by 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 tracing backward to the voltage source; or (2) by propagating forward, making voltage redistribution into an explicit process which evaluates the changes in voltages that occur whenever any device changes its state. Another more dramatic transformation in form involves splitting a model into multiple models, such as we have done in creating zero-order models as distinct from first-order models. The fact that causal model progressions can involve transformations in knowledge form as well as content increases the richness of the theory of knowledge evolution -- in contrast with, for example, trying to impart knowledge in the form of a collection of independent condition-action rules, such as the set of symptom-fault-fix associations that many experts use in troubleshooting.

3.1. 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 to the student's mental model that the new reasoning would require. These changes can occur to the device models and general circuit principles, as well as to 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.

Complete rewrites of aspects of the model are likely to be more difficult for the student to achieve than model refinements. However, 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.

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 model 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, a transistor? Even further, it is possible that 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 important in determining the learnability and useability of the model.

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 only 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" or, alternatively, one must introduce the voltage redistribution process. 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 control structure when a problem allows it. There is thus the added complexity of learning the contexts in which a particular control structure is applicable.

Finally, it should be noted that the type of model transformation can also 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 model progressions must take into account not only models' modifiability, but also how easily they can be put into practice in solving problems.

3.2. Types of Micro Model Evolutions

This section will discuss the types of "micro" evolutions that can occur within a model of a particular perspective, order, and degree of elaboration, such as evolutions that could occur within the zero-order, behavioral models. In learning, a model's knowledge may be augmented by refining, generalizing, or differentiating an existing knowledge component, by adding or substituting a new component, or by integrating several existing components 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, presented in the order of a possible knowledge evolution:

1. Knowledge acquisition -- The student acquires a new concept or law or problem solving skill. For example, many novices, as we have discussed, need to acquire the basic concepts of conductivity, circuit, and voltage drop.

2. Knowledge differentiation -- The student learns how an existing concept can be differentiated into two concepts. For instance, once students have acquired the concept of conductivity, they learn that it can be differentiated into conductive-resistive (i.e., a resistance) and purely conductive.
3. Knowledge integration -- The student integrates two concepts. For example, students need to synthesize their new concept of purely conductive paths with their concept of voltage drop -- purely conductive devices do not have voltage drops across them and any device connected directly across such a purely conductive device cannot have a voltage drop across it.
4. Knowledge generalization -- The student learns how an existing concept applies in a wide range of contexts. For example, students could generalize this new concept of a purely conductive parallel path (i.e., a short) from being immediately across a device to being anywhere on the device's feed path to any point on its return path or even immediately to ground (i.e., a short to ground). Also, students could learn that their new concept of resistance acquired in the context of a resistor can also be applied to other devices, such a light bulb or the collector-emitter path of a transistor.
5. Knowledge refinement -- The student refines an existing concept. For example, students can refine their understanding of voltage drop by observing that in parallel circuits, a device only needs for one of its feed paths to be "unshorted" in order to have a voltage applied to it.
6. Knowledge substitution -- The student replaces an earlier concept or skill with another. For example, students may go on to substitute causal propagation based upon voltage redistributions for the device-by-device tracing of feed and ground paths. Note that substitutions can be regarded as reversible (as in this example) or nonreversible.

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. 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 one 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 model space based upon their own pedagogical styles and goals. Within the present project, we have focused on tutoring troubleshooting, and have constrained the network of possible model evolutions to a linear progression, i.e., a curriculum. This curriculum has the objective of teaching troubleshooting for opens and shorts to ground in series-parallel circuits using the zero-order circuit model and the approach to troubleshooting described in the previous section. We have evaluated our implementation of this curriculum within the QUEST system using a small group of high

school students, and have found the tutoring approach to be effective. This work is described in more depth in White & Frederiksen (1986a; 1987), as well as in section 4.2.

3.3. Evolutions in Model Order and Degree of Elaboration

In addition to the "fine grained" progressions that can occur in the development of a particular model, "larger grained" progressions also occur, characterized by changes in model order and degree of elaboration¹³. In the following, we will illustrate how, in developing behavioral models of circuits, transitions from qualitative to quantitative models might be interleaved with increases in the degree of model elaboration. In this discussion, two levels of model elaboration will be distinguished, one in which the circuit principles included allow the student to reason about voltage distributions within a circuit, and a second in which additional principles are added to incorporate reasoning about current flow within the circuit. To help illustrate the progression, the reasoning of these different models will be presented for two simple transistor circuits (shown in Figures 7 & 8).

3.3.1. The Initial Degree of Elaboration: Reasoning about Voltages

Zero-order model. In the initial, zero-order model (described more fully in the previous sections), students are introduced to the basic idea of a circuit. In this context, they learn that there are two polarities of electrical force, and that for an electrical force to be applied to a device (such as a light bulb), both polarities must be applied, respectively, to two ports of the device. They also learn that devices have properties such as being a source of voltage (e.g., batteries) or not, and being conductive (e.g., bulbs and wires) or non-conductive (e.g., open switches). Finally, they learn that devices can have more than one state (e.g., a switch can be open or closed), and that the state can determine the device's properties (e.g., non-conductive or conductive). Experience with circuit problems leads then to a general idea of a series circuit: it must contain one or more voltage sources, and all elements in it must be conductive for there to be a voltage across all of its components -- if a single device within the circuit is nonconductive, then that device will have a voltage applied to it and no other device will have a voltage drop across it. Thus, students develop knowledge of where voltage drops will be present in a series circuit containing conductive and nonconductive components.

Next, students are presented with a more sophisticated model that knows about series-parallel circuits and can introduce the concept of a short, i.e., a purely conductive path that eliminates voltage drops across components connected in parallel with it. This particular model transformation can be motivated by giving students problems where they have to predict, for instance, the behavior of the light bulb in the circuit shown in Figure 2 as the switch is opened and closed. In this model progression,

¹³The only models that are at present implemented in our prototype learning environment are the zero-order models, at the level of elaboration described in Section 2

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 concept -- the concept 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.

Then, using a light bulb as a fledgling voltmeter, students learn that the distribution of voltages within a series-parallel circuit follows a zero-order version of Kirchhoff's Voltage Law: there is no voltage drop across purely conductive components (which can include parallel subcircuits) within a series circuit, and there is a voltage drop across all resistive elements (provided that the circuit contains no nonconductive elements). Using this zero-order model, students learn how changes in conductivity of components can alter the distribution of voltages in the circuit, and thereby cause other devices to change their state (for example, a transistor may go from the unsaturated to the saturated state). Students first learn to evaluate the effects of changes in conductivity on a device-by-device basis, and then they learn principles for redistributing voltages as an alternative technique for carrying out propagations.

Students learn to apply this knowledge in troubleshooting, in carrying out circuit design problems, and in predicting the behavior of circuits. For example, one could apply the zero-order model to reason about the behavior of a transistor amplifier such as the common-emitter amplifier shown in Figure 7. This requires a device model for the npn transistor which employs the following rules:

Transistor

States: Saturated or Unsaturated

If the base-emitter (B-E) voltage of the transistor is positive, then it is in the saturated state; otherwise, it is in the unsaturated state.

Internal Conductivity

If the state is unsaturated, then the collector-emitter (C-E) is nonconductive.

If the state is saturated, then the C-E is purely conductive.

These rules treat the transistor as a switch that is controlled by the voltage applied to the base-emitter. Applying this model to the common-emitter amplifier, when the input voltage to the amplifier is positive, the C-E of Q_1 will be purely conductive, creating a short across the output, which causes the output voltage to be zero. When the input voltage is zero, the C-E of transistor Q_1 is nonconductive, and there is then a good feed path from the positive side of the output to V_{CC} (the positive terminal of the battery) through R_1 . The voltage across the output is therefore positive. Thus, the zero-order model allows one to discover that the amplifier inverts the input signal (high in, low out, and vice versa). The common-emitter amplifier could actually be used in this way within digital circuits as long as the input voltage is switched to high enough levels to ensure that the transistor becomes fully saturated.

 Insert Figure 7 about here.

The first-order model. Within a first-order model, students generalize the zero-order concepts to reason about the effects of incremental changes. For example, students learn that reducing the resistance of a component within a series circuit causes a drop in the voltage across that component (the R->V Rule¹⁴) and an increase in voltage across the other components in the circuit (Kirchhoff's Voltage Law or KVL). Thus they learn how incremental changes in resistance and voltage influence, in an incremental manner, the distribution of voltages within a circuit, and how these changes in voltage can cause devices within the circuit to incrementally change their state variables in a sequence of causal propagations.

Students could apply these principles widely. For example, in simple analogue circuits, they could develop an understanding of how a simple amplifier works, such as the common-emitter amplifier shown in Figure 7. In first-order models, the transistor is modelled as increasing in conductivity across the collector-emitter (C-E) pathway whenever there is an increase in the base-emitter (B-E) voltage, and as decreasing in C-E conductivity (or, alternatively, increasing in resistance) when the B-E voltage decreases.¹⁵ In the common-emitter amplifier, an increase in the B-E voltage of the transistor (the input signal) causes an increase in the conductivity of the C-E circuit of the transistor (as specified in the first-order model of a transistor). This in turn causes the voltage across C-E to decrease (the R->V Rule), which is the output voltage of the amplifier. Thus, an important characteristic of the common-emitter amplifier can be deduced: changes in the output voltage of the amplifier will be the mirror image of changes in the input voltage (i.e., the output will be 180 degrees out of phase with the input).

Reasoning such as this can also be applied by students in understanding such circuit concepts as feedback (White & Frederiksen, 1986a). For example, consider the common-collector amplifier shown in Figure 8. In this amplifier, an increase in the input voltage causes an increase in the B-E voltage of transistor Q_1 (using KVL applied to the loop consisting of the input voltage source, the B-E of Q_1 , and R_2). Applying the device model for the transistor, the conductivity of the C-E of Q_1 increases, and the voltage across C-E decreases (by the R->V Rule). Since the voltage across R_2 must therefore increase (by KVL for the loop containing the voltage source, C-E of Q_1 , and R_2), the output voltage increases. However, in this amplifier there is negative feedback. Since the voltage across R_2 increased, the voltage across B-E

¹⁴The R -> V Rule states that a decrease in the resistance of a component causes a decrease in the voltage across the component, except if the component is directly connected to a constant voltage source.

¹⁵Increase in conductivity of the C-E of course correspond to increases in the C-E current through the transistor. Modelling the transistor with a direct causal link between B-E voltage and C-E current is reserved for models at the second degree of elaboration (see below).

of Q_1 must decrease (by KVL applied to the loop made up of the input voltage source, now assumed to be fixed, the B-E of Q_1 , and R_2). Thus, the feedback is a voltage in the opposite polarity to the initial increment in the input voltage, and this implies that the feedback is negative.¹⁶ First-order models are essential for detecting and understanding feedback, as DeKleer (1985) has argued.¹⁷

 Insert Figure 8 about here.

The quantitative model. In the quantitative model, Kirchhoff's Voltage Law is generalized to its quantitative form: the sum of voltages in any loop is zero. Students are introduced to the concept of a voltage divider (a set of resistive devices in series, connected across a voltage source), and shown that the total voltage across the divider is given by the sum of the voltage drops across the individual resistive devices, and that the magnitude of each individual voltage drop is proportional to the resistance of that device relative to the total resistance in the circuit. Thus, quantitative laws can be presented as model extensions that are useful when quantitative problems are encountered. They are not presented as a substitute for qualitative models, or even as more desirable or more correct. Indeed, students by this point would already have appreciated the importance of qualitative models in inferring the behavior of a complex circuit such as a digital logic circuit, in troubleshooting, and in circuit design.

3.3.2. The Second Degree of Elaboration: Reasoning about Current Flow

Zero-order model. Within the zero-order models, current is initially defined locally for each device (Ohm's Law). Whenever a voltage is applied to a conductive device, current will flow through the device in the direction of the negative source of voltage¹⁸. When no voltage is applied to a device, no current flows. Furthermore, if a device is nonconductive, no current can flow through it; and if it is purely conductive, it offers no resistance to a current flow. These rules allow students to derive current flow locally after using voltage distribution laws to assign voltages throughout the circuit. For example, suppose that a light bulb were connected across the output of the common-collector amplifier of Figure 8. Applying the zero-order model for redistributing voltages within the amplifier allowed us to deduce that when the input voltage is positive, there is a voltage applied to the light bulb (the output). If we now apply the local rule for determining current (Ohm's Law), we can conclude that a current will flow through the

¹⁶If one continued the propagation sequence, the amplifier would be found to cycle endlessly between increasing the output voltage and decreasing it. Such repeating cycles within a qualitative theory correspond to continuous processes in a quantitative theory (Weld, 1986), and the student needs to learn to recognize them as such.

¹⁷A zero-order analysis of the common-collector amplifier simply shows the transistor switching from the unsaturated state (for a zero input voltage) to the saturated state (for a positive input voltage), with the output going from zero to positive. This is, in fact, an accurate portrayal of the gross behavior of the amplifier, since the feedback can never completely cancel the input. (Rather, it has the function of increasing the input impedance). Within first-order models, reasoning about the relative magnitudes of input and feedback signals involves importing external knowledge that is not contained in the qualitative models themselves.

¹⁸This is the usual convention, which refers to the flow of positive charge.

light bulb. Similarly, when the input voltage is zero, there will be no voltage applied to the light bulb and therefore no current flow through it. The same reasoning can be applied to derive currents through the other components in the circuit.

After they are acquainted with the local application of Ohm's Law for determining current flow, students would then be introduced to additional principles for reasoning about current flow within a circuit. In particular, they would be presented with a model which incorporates Kirchhoff's Current Law (KCL). This law states that (1) the current flowing into a device will equal that flowing out of a device, and (2) if there is a current flowing out of a port of a device, there will be a current flowing into any device(s) connected to that device port. In addition, they learn to apply Ohm's Law (relating resistance to current flow) to reason about current flow within parallel circuits: if one branch of a parallel circuit contains no resistance, then all of the current flowing into the parallel circuit will flow through the purely conductive branch and no current will flow through the resistive branch.

As well as the introduction of these circuit principles, revisions are also made to device models for devices such as the transistor: In addition to a rule relating the B-E voltage to the conductivity of the C-E, a new rule would be added allowing a direct inference to be made about the C-E current -- when a positive voltage is applied to the B-E, there will be a current through the C-E (otherwise, there is no C-E current). With these new principles, students would now have an alternate way of reasoning about current flow within series and parallel circuits, which provides an alternative means for deriving the current flow through components (in addition to using voltage propagations and the local rule for current flow). For example, we shall apply this alternative model to analyze current flow in the common-collector amplifier discussed above (given in Figure 8). When the input voltage applied to the amplifier is positive, the C-E of Q_1 will have a current through it (by the revised transistor model). Then, applying KCL, if there is a current flowing out of the emitter of Q_1 , there will be a current flowing into both R_2 and the light bulb, which is connected across the output (since both offer paths towards the negative side of the voltage source). Thus, using KCL provides an alternate way to propagate the effect of an input voltage on current flow through the light bulb. In this way, students would have a first experience with a model which contains redundancies and that allows alternative correct solutions to a circuit problem.

First-order model. In this model, increasing the voltage applied to a device having a fixed resistance causes an increase in current through the device (Ohm's Law). Likewise, when the voltage across a device is fixed, decreasing its resistance will cause there to be an increase in current through the device (also Ohm's Law). Finally, for a fixed resistance, increases in the current through a device can be inferred to have been brought about by an increase in the voltage across the device (this is presented as a form of backwards reasoning using Ohm's Law). Since in first-order models variables such as resistance, voltage, and current have qualitative derivatives (rather than being regarded merely as

present or absent as in the zero-order models), Kirchhoff's Current Law (KCL) can be restated: (1) for any component, increasing the current flowing into the component results in an increase in that flowing out of the component, and (2) increasing the current flowing into a node (or point where two or more components are connected) will result in an increase in that flowing out of the node. With these additional principles, students could learn to reason directly about changes in current in both series and parallel circuits.

In the first-order model, a further revision is made to the transistor model to relate increments (decrements) in B-E voltage to increases (decreases) in the C-E current. This revision then allows alternative propagations in modelling the operation of transistor circuits. For example, for the common-emitter amplifier previously discussed (see Figure 7), increasing the input voltage (across B-E of Q_1) causes an increase in the current through the C-E of Q_1 (by the transistor model), which causes an increase in the current through R_1 (by KCL). Since the resistance of R_1 is fixed, we can use Ohm's Law to infer that the voltage across R_1 must have increased. Finally, since the voltage across R_1 increased, that across C-E of Q_1 (which is the output) must have decreased (by applying KVL to the loop made up of R_1 , C-E of Q_1 , and the voltage source)). Reasoning such as this is very much in the spirit of that of DeKleer (1985). We should also note that other propagations for the behavior of such circuits are possible using the set of constraints offered by the qualitative forms of Ohm's Law, Kirchhoff's Voltage and Current Laws, and the alternative rules within the device models. Thus, students will learn that the behavior of a physical system can be envisioned in more than one way, due to redundancies in the set of circuit principles that are used for deducing the system's behavior.

Quantitative circuit theory. The quantitative relations between voltage and current, and between resistance and current, are initially introduced as proportionalities: e.g., when the resistance is fixed, the current is directly proportional to voltage. Kirchhoff's Current Law is restated as: The sum of currents leaving any node of a circuit equals that entering the node (i.e., there is conservation of charge). By introducing a series of problems and their qualitative and quantitative analyses, and by collecting data from circuits used in these problems, students would derive for themselves a number of laws, such as Ohm's Law and the formula for the resistance of resistors connected in parallel. The quantitative theory is thus presented simply as an extension of the qualitative theory, but one which allows algebraic equations for expressing the underlying circuit principles (which by now should be quite familiar to the student). Students would also learn that problems can be solved through the algebraic manipulation of equations, but that each manipulation is based on circuit concepts that can be thought about in qualitative terms. Students should also discover that there are many classes of problems that are not amenable to such quantitative solutions, such as reasoning about the behavior of circuits, troubleshooting, or circuit design. In most cases, the quantitative model serves as a useful adjunct to qualitative reasoning in solving problems rather than as a replacement for it.

Summary. This progression of models increasing in order and degree of elaboration should enable 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 amplifier circuits. 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 (such as current limiting resistors), and can of course be used to calculate actual voltages and currents within a circuit. An important problem for future research is the model 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.

3.4. Developing Alternative Perspectives

The model evolutions discussed so far have been with respect to the students' behavioral models of circuit operation. However, just because students are adept at looking at circuit diagrams and predicting circuit behavior 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 underlying physics of device and circuit functioning. For instance, they may be unable to answer questions of the form: "Why does the collector current of a transistor depend upon the base-emitter voltage?", and "Why are voltage drops within a series circuit proportional to resistances?" Thus we argue that in order to attain a "deep understanding" of how a circuit works, students must evolve such alternative conceptualizations of circuit phenomena and be able to apply them in conjunction with their models of circuit behavior.

3.4.1. Functional Models

We have explored two approaches to developing in students a means for mapping from their qualitative models of circuit behavior to a functional understanding of circuits. Both approaches are based upon a study of the behavior of components of a circuit, and how their behavior depends upon the behavior of other components within the circuit. In each approach, the function of a circuit emerges from a study of its behavior as inputs to the circuit are allowed to vary. The two approaches differ in the nature of the components that are the focus of study: In the first approach, the focus is on devices within the circuit that have multiple states, while in the second, it is on designated subcircuits that constitute functional units in the design of the circuit.

The first approach adopts a device focus and emphasizes the discovery of patterns of interdependence among devices within a circuit. To facilitate this, students are shown how to construct a representation of the dependencies among devices within a circuit. For example, in a simple circuit, the state of one transistor may depend upon the states of two other transistors in the circuit: Suppose further that when the first two transistors are both saturated (e.g., "on"), the dependent transistor is also "on", while if either of them is unsaturated ("off"), the dependent transistor remains "off". Device interdependencies such as this one can be represented in a state dependency graph whose nodes represent devices and whose links represent particular dependencies. In the present example, the third transistor would be linked to the other two. By further parsing this dependency graph, sets of links and nodes performing specific functions (in the example, a logic gate) can be identified, and an analysis of the circuit as a set of interconnected functional units can be performed (White & Frederiksen, 1986a). In the present case, the third transistor is acting as an AND gate and the outputs of the other two transistors are providing inputs to that gate. In this way, circuit functions emerge from a study of the behavior of devices within the circuit. The implementation of these functions within the circuit can be seen by looking back to the specifics of the causal propagations that generated the specific functional dependencies at issue¹⁹.

The second approach to constructing a link between behavioral and functional models focuses on a study of the input-output relations for designated "chunks" of a circuit. In this approach, interesting circuit chunks (or subcircuits) are presented to the student, and the student's task is to discover the function performed by the particular subcircuit. The subcircuits chosen are the functional units circuit designers use in building more complex circuits.²⁰ For such interesting subcircuits (e.g., the common-collector and common-emitter amplifiers in Figure 7 and 8), the student explores the outputs that occur for each of the possible inputs they may receive in a larger circuit context. To illustrate, if the context is that of digital logic circuits, students would employ zero-order models and analyze the behavior of the amplifiers for cases where the input voltages are either high or zero. They would then discover the function performed by each -- the common-emitter amplifier would be seen to act as an inverter, while the common-collector amplifier would simply pass through the input logic value. If on the other hand the context is an analogue circuit, students would need to use first-order (and quantitative) models to understand the functionality of

¹⁹The alternative of basing the functional analysis on the causal trace which results from applying the behavioral model to the circuit, as DeKleer (1985) has advocated, has the disadvantage that the causal derivations of circuit behavior are not unique, and depend as we have seen on the particular qualitative constraints one has chosen to employ in explaining circuit behavior. For example, classifying the function of a resistor as either a "current-to-voltage converter" or a "voltage divider" depends upon the particular circuit laws used in deriving the circuit's behavior. Basing the functional classification on such causal derivations thus leads to ambiguous classifications.

²⁰Circuit designers do not create circuits "bottom up"; rather, they implement a top-level function using a repertoire of subcircuits having known functions, which are realized by the behavioral interaction of their subcomponents. For this reason, one can argue that while the student should understand the way in which the subcircuits implement their functions and how such subcircuits together implement top level functions, they do not need to learn how to discover functionally significant subcircuits within a larger circuit context.

the amplifiers. Thus, the form of qualitative model applied depends upon the context in which the subcircuit occurs. The analysis of circuit function is based on its input-output relations viewed from the perspective of a model of a particular order, and not on the specific steps of the causal propagation employed in deriving its behavior.

Once students have understood the functions of subcircuits, they then learn to model the behavior of larger circuits composed of these subcircuits. They are introduced to a new (and simple) way to analyze the interactions among such components, viewing each as a receiver and transmitter of information. Information received by a circuit component (the set of input signals) is processed (for example, the logical operation AND is performed, or an input signal is inverted), and an output signal is then created and sent on to other components. This simple information-processing model for simulating the functional behavior of large circuits is similar to that used by experts, and can be used as a basis for tutoring troubleshooting (Frederiksen, White, Collins, & Egan, in press).

3.4.2. Reductionistic, Physical Models

Explanations of circuit behavior rely on qualitative models of devices and on general circuit principles which incorporate the basic laws of circuit theory relating voltage, resistance, and current. However, neither the device models nor the circuit principles are explicitly derived for the student from more elementary physical principles; rather, they are simply represented as rules that are applicable to circuits in general. To foster a deeper understanding, it is possible to develop coherent models for devices and explanations of the basic steady-state circuit principles in terms of a more elementary physical theory. In this theory, electrical forces within a circuit are derived from, and in turn influence, the distributions of charged particles within the components of a circuit. Such qualitative physical models are quite commonplace in introductory textbooks, where they are used to explain the behavior of devices such as capacitors, diodes, and transistors. What is lacking, however, is a bridge from the physics of electrically charged particles (e.g., Coulomb's Law) and mechanics (e.g., Newton's Laws) to an understanding of the basic circuit laws (e.g., Ohm's and Kirchhoff's Laws) that form the backbone of circuit theory.

We are working on a simple, physical model that focuses on the electrical fields that exist when there are non-uniform distributions of positive and negative charge carriers within a circuit²¹. For example, a battery produces an excess of positive charge carriers on one terminal and an excess of negative charge carriers on the other. An electrical field thus exists around each battery terminal. These

²¹Haertel has developed a similar type of model, although it differs in certain key respects from ours (see, Haertel, 1987, and, Frederiksen & White, 1988).

electrical fields can act upon any mobile charge carriers that are within adjacent regions²², attracting those with a like charge and repelling those with an unlike charge. This movement of charge carriers in turn alters their distribution within adjacent regions of a circuit, which then induces additional migrations of charge carriers, and so forth. This simplified physical model can explain how within a circuit a voltage drop develops across a resistor (a device which, within the model, slows down or impedes the movement of charge carriers), and how a constant current is developed within a circuit containing a voltage source. Moreover, applying the model to simple series and parallel circuits can reveal the mechanisms underlying circuit principles such as Kirchhoff's and Ohm's Laws. These principles become emergent properties of such circuits, representing the behavior of charge carriers when the physical model has reached a steady state. For example, one can illustrate how in a series circuit containing a nonconductive element, the entire voltage drop develops across that element. Furthermore, it is possible to show how, when the nonconductive element (e.g., a switch) becomes conductive, the voltages are redistributed across resistive elements of the circuit. Depending upon the form of measurements that are taken during the operation of the physical model and the circuit problems to which it is applied, the model can be used to emphasize particular circuit principles (involving voltage or current) such as are introduced within the progression of behavioral models, in either their qualitative (zero- or first-order) or quantitative forms. The particular circuit principles emphasized can be chosen to be compatible with the order and degree of elaboration of the behavioral model currently being developed. We hypothesize that linking the behavioral models to students' intuitive notions of "attractive pulls" and "repulsive pushes" via the introduction of this reductionistic, physical model will increase the learnability of the behavioral model, as well as improve the depth of the students' understanding.

A qualitative, physical model such as this, when implemented, will allow the student to request a change in perspective from macroscopic to microscopic models of circuit behavior. The microscopic model will be able to explain and simulate the causality of circuit operation in terms of electrical forces and their effects on the behavior of charge carriers within the circuit. These explanations will be consistent with the voltage redistribution process of the macroscopic models, and the macroscopic circuit principles will emerge from the operation of the microscopic model.

3.5. Coordinating Multiple Models

This paper has addressed issues relating to the learnability of a set of mental models. In discussing learnability, our primary construct was that of causal consistency. This construct applies to the relationship among models of differing orders and degrees of elaboration. For instance, we argued that

²²Coulomb's Law of course states that the force between charged particles decreases with the square of the distance; however, within the qualitative physical model, we limit the influence of an electrical charge to an adjacent "region", beyond which it has no direct influence. Thus, we allow action only at "near" distances.

electric force should be the causal agent within both qualitative and quantitative models of circuit behavior; and we argued further that when these models are elaborated to incorporate principles for reasoning about changes in current, that these changes should be explained in terms of electric force. The construct of causal consistency also applies across models of different perspectives. For instance, we stated that our objective in creating microscopic models of circuit behavior was to enable the principles governing our macroscopic models (such as KVL) to become emergent properties of the microscopic models, and similarly we illustrated (in White & Frederiksen, 1986a) how the principles governing circuit functioning can emerge from the types of causal models of circuit behavior described in this paper. Thus we argue that causal consistency across domain models of differing perspective, order, and degree of elaboration can increase the learnability of the set of mental models that are necessary to "deeply" understand domain phenomena.

We can also apply similar arguments to define the properties that this set of mental models must possess in order to enable them to play a role in helping students to understand the phenomena of other domains. For example, can acquiring this set of electrical circuit models facilitate learning the concepts, laws, and models needed to understand other physical systems? The answer to this question depends upon the causal generality of the concepts and laws embedded in the models. For example, if the models of microscopic circuit behavior are based upon concepts such as force, acceleration, and equilibrium (which we are attempting to do), then the acquisition of such a model can foster and be fostered by an understanding of mechanics. Similarly, if one's models of circuit functionality are based upon general principles of information flow, the acquisition of such a model can foster and be fostered by an understanding of computer science. The thesis is that understanding and learning are facilitated by a consistency of concepts and causal relations across different types of models and across different domains. This paper has attempted to show that such consistency is possible. The view is compatible with the "generality hypothesis" that governs much AI research, as well as research on physical theories (both qualitative and quantitative).

4. The Learning Environment

We have focused on creating a progression of models that makes a gradual transition from naivete to expertise. To facilitate this transition, the instructional system:

1. emphasizes a qualitative, causal analysis that builds upon novices' existing knowledge;
2. motivates learning via problem solving and appropriate problem selections; and
3. generates causal explanations of circuit behavior, and illustrates problem solving strategies.

The assumption is that (1) by giving students problems that present a manageable cognitive challenge (that is, problems they could solve with a small revision to their mental model) and 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, students will, at any stage in the learning process, be able to transform their model to match that of the system. Our view is that incorrect inductions are not a necessary consequence of the learning process. The hypothesis is that in learning environments where the model progression and problem sets are designed appropriately, incorrect model transformations are unlikely. The tutoring system therefore does not actively attempt to diagnose and treat wrong model transformations. It does, however, provide feedback, and it allows the student to compare his or her reasoning with that of the model that is currently driving the system. Juxtaposing the student's and system's problem solving allows the student to debug his or her model when difficulty is encountered.

The progression of models is used to define problem sets that motivate transformations in the student's mental model. At any stage in learning, the instructional goal of the student is to master the model that is currently driving the simulation environment. 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. The theory is that, by giving students problems in this set (i.e., problems that are just beyond their level of competence), students will be motivated to revise their model. This model revision will be facilitated because it requires only a small change to their present model, in an environment where feedback and explanations are available to help them understand the model transformation. Moreover, the models have been designed with the requirement of modifiability in mind. Students should thus be motivated and able to transform their model into the next model in the sequence²³.

The initial sequence of problems in each set is crucial to facilitating a correct model transformation. In particular, the initial problems need to be selected so that they can be solved by the transformed model but not by some other erroneous model transformation. This helps to avoid the induction of "buggy" mental models. After the correct model transformation has been induced, the remaining problems serve the function of giving students practice in using their new mental model of circuit behavior.

Model progressions also determine the focus of explanations. To facilitate model transformations, the system can turn any problem into an example for the student by reasoning out loud while it solves the problem, focussing its explanations on the difference between the transformed and the untransformed

²³In addition to problems requiring the transformed model, some problems need to be interspersed from the previous problem set, in order to provide negative exemplars of a concept. For example, if students were learning that a short from a point on a device's feed path to a point on its ground path prevents 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 demonstrated, negative instances of a concept are very important to learning. Providing some problems from the previous set can serve this function.

model. The difference between models also defines what aspects of reasoning should be represented graphically 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.

In summary, the particular model transformation undertaken at each stage 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.

4.1. Interacting with the Learning Environment

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.

All of these functions of the instructional system are provided, at a given point in the learning progression, by the model that is currently driving the learning environment. Each model can:

1. Simulate circuit behavior. A given 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. Tutor the students. By reasoning out loud, the model can generate qualitative, causal explanations for circuit behavior.
3. 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.

Each model can thus serve as a circuit simulator, a tutor, and a student model.

This architecture for an intelligent learning environment allows students great flexibility in their choice of an instructional strategy. Particular strategies they might use include the following:

Exploration and discovery learning. 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.

Learning via induction and feedback. 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 -- they are just given feedback as to whether their answers are right or wrong and have to induce the model transformation for themselves.

Learning from examples and explanations. 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. This strategy thus focuses on presenting examples together with explanations prior to practice in problem solving.

Mixed strategies. Students can of course also use mixed learning strategies. For example, they can combine inductive learning with the use of examples and explanations. They can also use the circuit editor to modify and experiment with problems presented to them by the system.

The tutoring system also provides students with additional tools to help them direct their own learning. The concept of a progression of models allows students to understand which models have been mastered and which remain to be learned. They are provided with a map of the learning space that is defined by the progression. This provides students with information about the instructional objective of each model. The students can use this map to select topics of study and to access the associated problem sets.

The systems's capabilities of allowing different learning strategies and providing tools for self-directed learning are important from several respects. First, they allow for individual differences in learning style -- not all students learn best using the same pedagogical technique (Cronback & Snow, 1977). Second, they provide students with a sense of autonomy that may be a crucial factor in motivation. Finally, we argue that experiences in managing their own learning, in an environment that makes explicit the evolution of domain knowledge and the possible pedagogical strategies, should play a valuable role in helping students to develop the skills that are needed to become expert learners.

4.2. Instructional Effectiveness

The learning environment was tried out on seven high school students who had had no formal instruction in circuit theory. (The results are further described and discussed in White & Frederiksen (1986a).) 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 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 a set of circuit problems as a pretest, and asked to explain the behavior of each circuit as the states of devices within it were manipulated. As described earlier in the paper, the students initially 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 six circuit problems they had attempted in the pretest along with three troubleshooting problems, and were asked to explain the behavior of the circuit or to troubleshoot.

After five hours of working within the learning environment on an individual basis, all seven of the students were able to make accurate model-based 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 nine 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 (involving the possibility that the fault could be either a short or an open), and all of the students were able to get the correct answer.

In their use of the system, all of the students were remarkably conservative. 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 within the progression, and went through the set of models 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 reasons why the students employed this "conservative" learning strategy. First, the fact that the system presented a curriculum to the students implied that its designers believed it was a good idea to progress through the models in a linear order. If, instead, students 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 might have been quite different. Second, the students' conception of how one learns, derived from their school environment, is primarily that of following a curriculum by hearing explanations and 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 mode! of learning. Finally, when interacting with the system, the students were being observed. This may have inhibited their exploratory behavior and increased their desire to "do the right thing", i.e., get correct answers to problems as quickly as possible. They thus relied on the explanations, rather than on discovering model transformations for themselves.

The implications of these results are that in our future research, we should explicitly teach and encourage alternative learning strategies. For instance we could apply the model progression concept to students' mental models of their own learning processes. Using the flexibility of QUEST, students could start with a familiar learning strategy, i.e. learning from examples and explanations, gradually progress to learning via induction and feedback, and finally be introduced to exploration and discovery learning. Through this progression, students could acquire these alternative learning strategies, as well as the skills needed for independent scientific inquiry.

5. Summary

The design of our intelligent learning environment is based upon a theory of expertise and its acquisition. We have argued 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, causal models. Algebraic, constraint-based reasoning should be introduced only after the domain is understood in causal terms.

In this article, we focused primarily on qualitative, behavioral models of electrical circuit operation which have been designed to make the causality of circuit behavior derivable, in a clear-cut manner, from basic physical principles. The constraints on model evolution, in terms of causal consistency, modifiability, and learnability, were discussed and a sequence of models that embody a possible transformation from novice to expert status was 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 causal 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 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 directed either by the system or by the student. Students are provided with 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. Thus students can be actively engaged not only in problem solving but also in managing their learning.

ACKNOWLEDGEMENT

This research was funded by the Army Research Institute, under contract MDA-903-87-C-0545, and 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. We are grateful to Judith Oransanu of ARI, as well as to Marshall Farr and Susan Chipman of ONR for supporting and encouraging this work. The initial system prototype was programmed by Frank Ritter, who helped with its design, particularly the circuit orientation scheme and the user interface. The current implementation was programmed by Eric Cooper and Frank Ritter. We wish to thank Bill Clancey and two anonymous reviewers for their constructive comments on this paper. In addition, we are grateful to Allan Collins, Wally Feurzeig, Paul Horwitz, Ed Smith, Kathy Speohr, and Mark Burstein for many stimulating and helpful discussions.

REFERENCES

1. Anderson, J.R., Farrell, R., & Sauers, R. (1984). Learning to program in LISP. Cognitive Science, 8, 87-129.
2. Bobrow, D.G. (Ed.) (1985). Qualitative Reasoning about Physical Systems. Cambridge, MA: MIT Press.
3. Brown, J.S., & Burton, R.R. (1978). Diagnostic models for procedural bugs in basic mathematical skills. Cognitive Science, 2, 155-192.
4. Brown, J.S., & Van Lehn, K. (1980). Repair theory: A generative theory of bugs in procedural skills. Cognitive Science, 4, 379-426.
5. 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.
6. 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.
7. Bruner, J.S., Goodnow, J.J., & Austin, G.A. (1956). A Study of Thinking. New York: Wiley.
8. Chi, M., Feltovich, P., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. Cognitive Science, 5, 121-152.

9. Clancey, W.J. (1986). Qualitative student models. Ann. Rev. Comput. Sci., 1, 381-450.
10. 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.
11. Collins, A. (1985). Component models of physical systems. Proceedings of the Seventh Annual Conference of the Cognitive Science Society, University of California, Irvine.
12. Cronbach, L.J., & Snow, R.E. (1977). Aptitudes and Instructional Methods: A Handbook for Research on Interactions. New York: Irvington.
13. Davis, R. (1983). Reasoning from first principles in electronic troubleshooting. International Journal of Man-Machine Studies, 19, 403-423.
14. deKleer, J. (1979). Causal and teleological reasoning in circuit recognition. TR-529, MIT Artificial Intelligence Laboratory, Cambridge, MA.
15. deKleer, J. & Brown, J.S. (1983). Assumptions and ambiguities in mechanistic mental models. In D. Gentner & A. Stevens (Editors), Mental Models. Hillsdale, NJ: Lawrence Erlbaum Associates.
16. deKleer, J. (1985). How circuits work. In Bobrow, D.G. (Ed.). Qualitative Reasoning about Physical Systems. Cambridge, MA: MIT Press.
17. diSessa, A. (1983). Phenomenology and the evolution of intuition. In D. Gentner & A. Stevens (Editors), Mental Models. Hillsdale, NJ: Lawrence Erlbaum Associates.
18. Feltoovich, P.J., Spiro, R.J., & Coulson, R.L. (in press). The nature of conceptual understanding in biomedicine: The deep structure of complex ideas and the development of misconceptions. In D. Evans & V. Patel (Editors), The Cognitive Sciences in Medicine. Cambridge, MA: M.I.T. Press.
19. Feurzeig, W. & White, B.Y. (1983). Development of an articulate instructional system for teaching arithmetic procedures. BBN Report No. 5484, BBN Laboratories, Cambridge, Massachusetts.
20. Feurzeig, W. & Ritter, F. (in press). Understanding reflective problem solving. To appear in Psotka, J., Massey, L.D., & Mutter, S. (Eds.). Intelligent Tutoring Systems: Lessons Learned. Lawrence Erlbaum Associates, Hillsdale, New Jersey.
21. Forbus, K., & Stevens, A.S. (1981). Using qualitative simulation to generate explanations. BBN Report No. 4490. Cambridge, MA: BBN Laboratories.
22. Forbus, K.D. (1985). Qualitative process theory. In Bobrow, D.G. (Ed.). Qualitative Reasoning about Physical Systems. Cambridge, MA: MIT Press.
23. Forbus, K. & Gentner, D. (1986). Causal reasoning about quantities. In the Proceedings of the Eighth Annual Conference of the Cognitive Science Society, Amherst, Massachusetts.
24. Frederiksen, J., & White, B. (in press). "Intelligent tutors as intelligent testers." To appear in Diagnostic Monitoring of Skill and Knowledge Acquisition. N. Frederiksen, R. Glaser, A. Lesgold, & M. Shafto (Editors), Lawrence Erlbaum Associates, Hillsdale, New Jersey.
25. Frederiksen, J., White, B., Collins, A., & Egan, G. (in press). "Intelligent tutoring systems for electronic troubleshooting." In Intelligent Tutoring Systems: Lessons Learned. Psotka, J., Massey, D., & Mutter, S. (Editors), Lawrence Erlbaum Associates, Hillsdale, New Jersey.
26. Frederiksen, J., & White, B. (1988). "Cognition and explanation: Designing reductionistic, causal models of electricity". Report No. 6689, BBN Laboratories, Cambridge, Massachusetts.
27. Gentner, D., & Stevens, A. L. (1983). Mental Models. Hillsdale, NJ: Lawrence Erlbaum Associates.

28. 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.
29. Haertel, H. (1987). "A qualitative approach to electricity". Report No. IRL87-0001, Xerox Palo Alto Research Center, Palo Alto, California.
30. Horowitz, P., & Hill, W. (1980). The Art of Electronics. Cambridge, England: Cambridge University Press.
31. Kieras, D. & Bovair, S. (1984). the role of a mental model in learning to operate a device. Cognitive Science, 8, 255-273.
32. 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.
33. Larkin, J.H., McDermott, J., Simon, D.P., & Simon, H.A. (1980). Expert and novice performance in solving physics problems. Science, 208, 1335-1342.
34. 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.
35. 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.
36. Ritter, F. (1987). OREO -- Adding orientation to a dynamic qualitative simulation. Report No. 6560, BBN Laboratories, Cambridge, Massachusetts.
37. 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.
38. Smith, E. E., & Goodman, L. (1984). Understanding instructions: The role of an explanatory schema. Cognition and Instruction, 1, 359-396.
39. Sleeman, D., & Brown, J. S. (Eds.) (1982). Intelligent Tutoring Systems. London: Academic Press.
40. Soloway, E. (1986). Learning to program = learning to construct mechanisms and explanations. Communications of the ACM, 29-9, 850-858.
41. Soloway, E., Rubin, E., Woolf, B., Bonar, J., & Johnson, W.L. (1983). Memo-II: An AI-based programming tutor. Journal of Computer-Based Instruction, 10, 1.
42. Steinberg, M.S. (1983). Reinventing electricity. In Proceedings of the International Seminar, Misconceptions in Science and Mathematics, Ithaca, New York.
43. Weld, D. (1983). Explaining complex engineering devices. BBN Report No. 5489, BBN Laboratories, Cambridge, Massachusetts.
44. Weld, D. (1986). The use of aggregation in causal simulation. In Artificial Intelligence, 30, 1-34.
45. 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.
46. White, B. Y., & Frederiksen, J. R. (1985). QUEST: Qualitative understanding of electrical system troubleshooting. ACM SIGART Newsletter, 93, 34-37.
47. White, B., & Frederiksen, J. (1986a). "Progressions of Qualitative Models as a Foundation

- for Intelligent Learning Environments." Report No. 6277, BBN Laboratories, Cambridge, Massachusetts.
48. White, B., & Frederiksen, J. (1986b). "Intelligent tutoring systems based upon qualitative model evolutions." In the Proceedings of the Fifth National Conference on Artificial Intelligence, Philadelphia, Pennsylvania.
 49. White, B., & Frederiksen, J. (1987). "Qualitative models and intelligent learning environments." In AI and Education. R. Lawler & M. Yazdani (Editors), Ablex Publishing Corporation.
 50. White, B., & Horwitz, P. (1987). "ThinkerTools: Enabling Children to Understand Physical Laws." Report No. 6470, BBN Laboratories, Cambridge, Massachusetts.
 51. Williams, B.C. (1985). Qualitative analysis of MOS circuits. In Bobrow, D.G. (Ed.) (Ed.). Qualitative Reasoning about Physical Systems. Cambridge, MA: MIT Press.

Received March 1986; revised version received November 1987

Running heading: CAUSAL MODEL PROGRESSIONS

Figure Captions

Figure 1: An elementary circuit used in studying naive models of electricity.

Figure 2: An example of a circuit used to motivate model transformations.

Figure 3: Troubleshooting a simple series circuit with a test light.

Figure 4: A circuit containing a capacitor, used to illustrate causal propagations.

Figure 5: A history of device data tables that represents the behavior of the circuit shown in Figure 4.

Figure 6: A summary of the behavior of the circuit shown in Figure 4, as derived from the history of device data tables shown in Figure 5.

Figure 7: A simple common-emitter amplifier. (V_{cc} stands for the positive side of the voltage source.)

Figure 8: A simple common-collector amplifier. (V_{cc} stands for the positive side of the voltage source.)






Figure 1.



(a)



(b)

Figure 2.

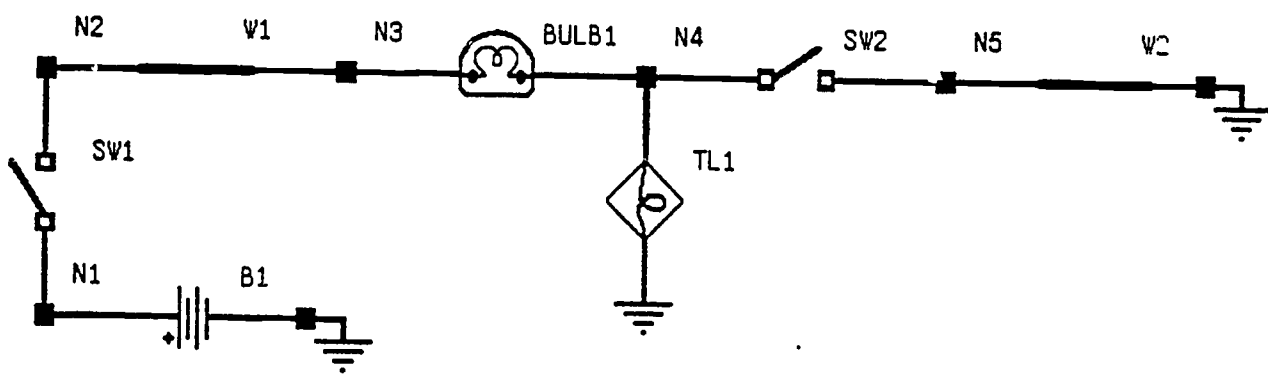


Figure 3.




Figure 4.

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 ND	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 ND	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 ND	HOT C-R NO	HOT C-R NO	HOT C-R NO	COLD C-R NO	COLD C-R ND
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


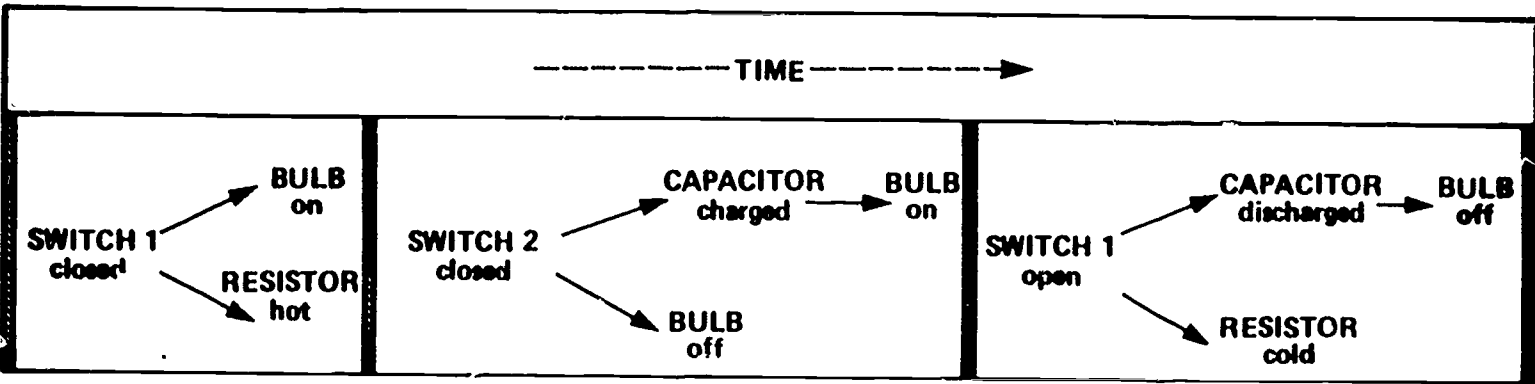


Figure 6




Figure 7.




Figure 8.