Tasks, Models, Perspectives, Dimensions

Zheng-Yang Liu

EDS Research, Austin Laboratory
1601 Rio Grande, Ste 451
Austin, Texas 78701
liu@austin.eds.com

Arthur M. Farley

Computer and Information Science
University of Oregon
Eugene, Oregon 97403
art@cs.uoregon.edu

Abstract

This paper presents a task-driven approach to perspective-taking for qualitative reasoning. Central to our approach is the notion that model formulation and selection is an integral part of reasoning about complex physical systems. Using the domain of electronic circuits as an example, we investigate three perspective-taking dimensions in qualitative reasoning: topological configuration, structural aggregation, and ontological choice. We show that our approach can extend the range of automated qualitative reasoning about complex physical systems.

Introduction

When reasoning about the physical world, we adopt some point of view or perspective. To study commonsense reasoning and write programs that reason effectively about the physical world, we must understand the nature of perspective-taking in qualitative reasoning — which perspectives to take, how to represent them, and when to shift from one perspective to another.

In this paper, we describe an implemented framework, ARC, based on a task-driven, perspective-taking approach to formulating and selecting models of a target system for a given task. Each model embodies a specific perspective which can involve choices made on several dimensions. We investigate three such dimensions for automated model formulation and selection: topological configuration, structural aggregation, and ontological choice.

We begin by considering an example. Figure 1 presents a half-wave rectifier, a simple nonlinear circuit. Given an *ac* voltage input, it produces a *dc* output. The explanation below is typical of the standard electronics textbook treatment of how the circuit works:

> "The capacitor becomes charged up almost to the input peak voltage when the diode is forward biased. When the diode is reverse biased, the capacitor partially discharges through the load. Since the capacitor always has some positive charge, the diode becomes forward biased only near the peaks of input voltage. At this time it passes a current pulse to the capacitor to replace the charge lost to the load..." (page 45, Bell 1980)

The explanation is purely in qualitative, causal terms. The circuit's overall behavior is inferred from the behaviors of individual components and their interconnections. Abstracting from this example, we see three dimensions of "perspective-taking" at work:

1. Topological configuration - focusing on a subset of the structure of a system. For example, the explanation deals with nonlinear behavior of the circuit based on two unstated configurations of the circuit, reflecting distinct device states.
2. Structural aggregation - abstracting structural elements into composite constructs to suppress uninteresting detail. For example, the explanation suppresses mentioning *R1* and *R2* when the diode is forward-biased but still implies their functionality.
3. Ontological choice - using distinct vocabularies to describe the behavior of a physical system. For example, the explanation integrates both macroscopic (device-level) and microscopic (charge-level) descriptions, such as "voltage or current", and "charges".

Our research focuses on these three dimensions of perspective-taking for reasoning about electronic circuits. Figure 2 shows eight such models of the rectifier circuit generated by ARC for various tasks. Each model embodies a distinct perspective of one or several dimensions. In general, a model of a circuit can be formulated that embodies an N-dimensional choice of perspective. For example, a description of a circuit's active topological configuration in a specific qualitative state stands as a model which embodies a 1D perspective (e.g., models 3 and 5 in Figure 2), revealing the dynamic structure of the circuit topology in the given state. A description of the structural aggregation over the active configuration stands as a model embodying a 2D perspective (e.g., models 4 and 8 in Figure 2). This perspective reveals not only the dynamic configuration of the circuit, but also a series of two black boxes encapsulating component devices in the configuration. The QDE (qualitative differential equations)

This work was supported in part by a Tektronix Fellowship. The research was conducted while the first author was at the University of Oregon.

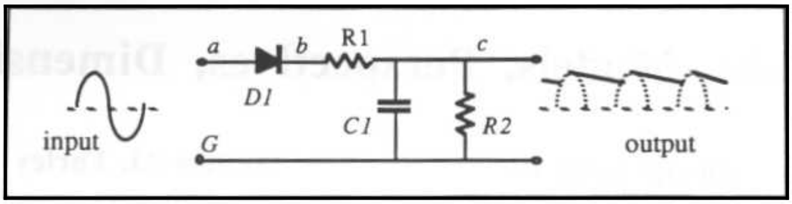

Figure 1: A half-wave rectifier with capacitor smoothing

description of an active configuration stands as a model embodying a 3D perspective (e.g., models 1, 2, 6, and 7 in Figure 2), with the additional dimension reflecting ontological choice. For example, models 1, 2, and 7 are in the standard device ontology [de Kleer & Brown 1984, Williams 1984], while model 6 is in the charge-carrier ontology [Liu and Farley 1990].

In this work, we define model generation as representation reformulation of general domain knowledge specialized for a particular task at hand. The system topology of a circuit as a symbolic description of the schematics of the circuit, is the only circuit-specific information provided. It represents the static structural connectivity of circuit component devices and remains unchanged throughout the reasoning process. No assumptions are made about whether the circuit is a linear or nonlinear system or about possible current flow paths. No assumptions are made about which individual components should constitute certain black boxes. Given a general domain theory of electricity, the system topology of a target circuit, and a specific task definition, ARC generates one or several models of the circuit and dynamically switch between them to carry out the task while suppressing unnecessary detail. Figure 3 presents the architecture of ARC. The input to ARC is a task definition of perturbation analysis. The output from ARC is a qualitative, causal explanation of circuit behavior as a result of the input perturbation.

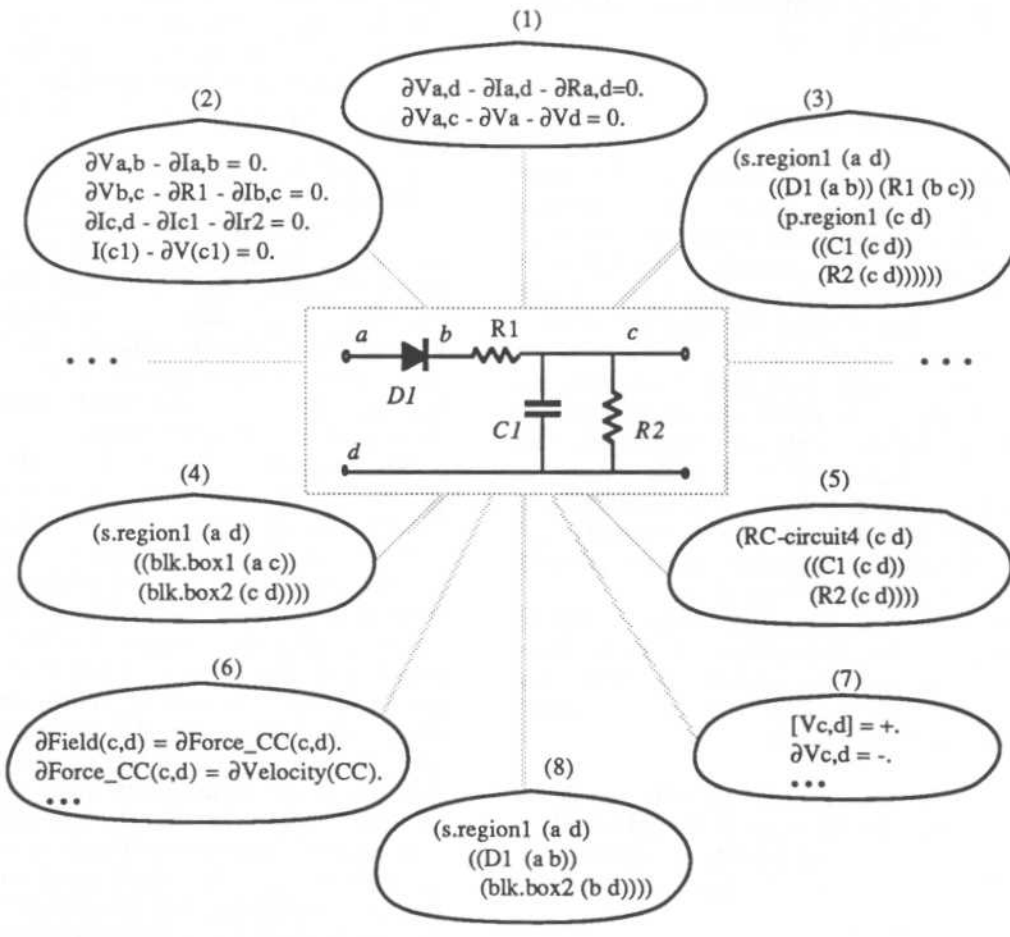

Figure 2: Model formulation from different perspectives of a physical system

Section 2 introduces a simple task definition language for circuit perturbation analysis. We indicate what kind of information we need from a task definition to allow perspective-taking and dynamically formulation of models. Section 3 illustrates the three dimensions we consider in this paper. We demonstrate topological configuration, structural aggregation, and ontological shift as driven by the task at hand. Finally, we discuss related work and suggest several questions for future research.

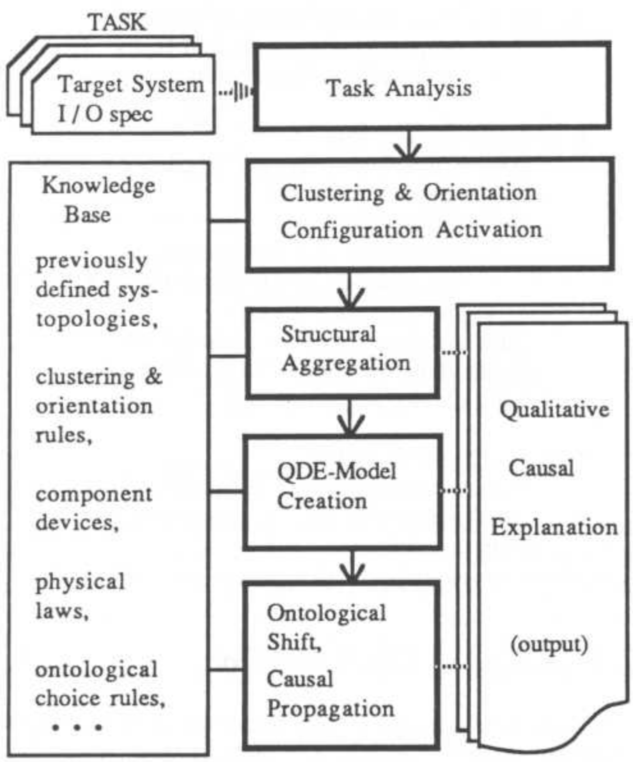

Figure 3: ARC Architecture

Task Definition

For qualitative causal analysis of electronic circuits, our current implementation focuses on standard perturbation analysis [de Kleer & Brown 1984] and supports a simple language to define such tasks. Three items comprise a task definition in ARC: (1) Name of a target system, (2) Specification of input perturbations, and (3) Specification of output desired.

The name of a target system in a task definition provides ARC with access to a symbolic description of the system topology of the target circuit — circuit component devices and their physical connectivity. To define the system topology of a circuit, one specifies the types of component devices and how they are connected, via nodes, into a device network. Figure 4 shows the system-topology specification of the half-wave rectifier in ARC.

Each device and node in the network is represented as an object that has its own internal state and procedures to interface with the outside world. Given a task, ARC compiles the circuit topology into an object-oriented representation by setting up links as attributes of the objects. ARC maintains a catalog of device models as well

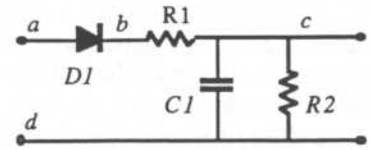


```
(define-circuit! Half-Wave-Rectifier
  ( ; Components:
    (Diode:     D1)
    (Resistor:   R1  R2)
    (Capacitor:  C1) )
  ( ; Connections:
    (a (terminal1 D1))
    (b (terminal2 D1) (terminal1 R2))
    (c (terminal2 R1) (terminal1 C1) (terminal1 R2))
    (d (terminal2 C1) (terminal2 R2))))1
```

Figure 4: ARC's representation of system topology

as system-topology specifications for circuits created in the simulation environment. Each can be referenced by name in a particular task definition.

Note that no current flow path or structural aggregation are implied in the system topology of a circuit. We cannot do so without knowing the input and desired output.

To specify input, the traditional qualitative simulation paradigm is primarily concerned with how a physical system, in some equilibrium state, responds to a single input perturbation, such as ∂V_{in} = +, or ∂V_{in} = -. This method, called "small signal analysis" [de Kleer 1984, Williams 1984], works under the hidden assumption that the active configuration of the target system is already identified.

ARC extends this single-perturbation method by allowing a sequence of input perturbations of arbitrary length based on a discrete set of time points. In contrast to the small signal analysis, each perturbation value in a sequence now consists of two qualitative values - (Qval Qdir). For example, one may give a sequence of perturbations to specify a cycle of sinusoidal voltage at input as in Figure 5.

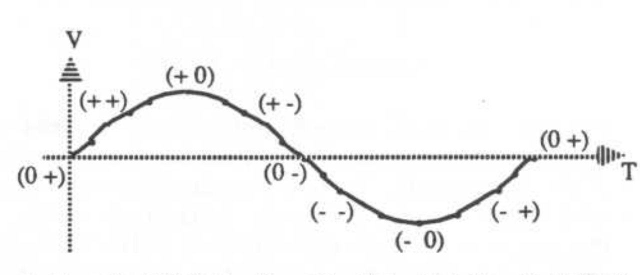

(voltage (in GND) (0 +) (+ +) (+ 0) (+ -) (0 -) (- -) (- 0) (- +))

Figure 5: A sequence of perturbations as input

In this example, each of the perturbing signals in a sequence specifies a qualitative state of the system parameter "$\text{voltage}_{in,GND}$". A sequence of such perturbations drives the qualitative simulation of the circuit. Each signal's processing is performed in the context created by those signals early in the sequence and early in time. With a sequence of signals we can simulate configuration switching as the individual devices change states. Perturbation for small signal analysis is easily represented as a sequence of length one. For zero-order analysis, we just let Qdir = 0 in (Qval Qdir) of a sequence.

To specify output, ARC allows the user to indicate the behavior of a specific system variable as desired output. For example, one can specify to analyze voltage, current, or charge flow with respect to the nodes in the system topology of a target circuit. The nodes specified define the output structural unit, encapsulating either a single device or a group of devices.

An output variable specification is best viewed as placing a "probe" in the circuit under analysis. In this respect, the current or charge-flow "probe" are placed in series with the output structural unit, and the voltage "probe" placed in parallel with the unit where the voltage difference is of interest. If a voltage "probe" is specified with only one node, ARC automatically assumes "ground" for the other node. The output that ARC produces is a qualitative causal explanation of the behavior as a result of the input perturbations.

For example, "Given the half-wave rectifier with a sinusoidal input voltage, what is the behavior of the output voltage between nodes c and d?"
This task is specified as follows:

- Target System: [half-wave-rectifier1]
- Input Specification: (voltage (a d) (0 +) (+ +) (+ 0) (+ -) (0 -) (- -) (- 0) (- +)...)
- Output Specification: (voltage (c d))

Table 1 summarizes the information extracted from a specific task definition to provide ARC with guidance for perspective-taking. This information enables ARC to formulate appropriate circuit models which suffice to carry out the given task while minimizing unnecessary detail. Below, we describe how ARC dynamically formulates models of the target system to adapt to the task at hand for qualitative causal analysis and produces output answers.

Perspective-Taking

Given a task, ARC reasons from structure to behavior, formulating appropriate models from the system topology of the target circuit. To formulate configuration models, ARC uses a graph search method [Arnborg & Proskurowski 1989] to determine the active topological configuration in the given qualitative state of the circuit. For structural aggregation, ARC extends the notion of *slices* [Sussman & Steele 1980] by automatically aggregating components into black boxes to suppress irrelevant details for the task at hand. For ontological shift, ARC creates QDE models in chosen ontologies and uses a set of ontological choice rules to control ontological shift if needed [Liu & Farley 1990].

TABLE 1. Information Extracted by ARC from a Task Definition

Target System Name:
System Topology (Circuit Schematics)
- component devices & their types
- connectivity of devices via nodes

Input Specification:
System parameter being perturbed
Pole nodes and polarities
Perturbation type
- a single quantity perturbation
- a sequence of quantity perturbations
- a mathematical function

Qualitative value of each perturbation
Qualitative derivative of each perturbation
Input ontology

Output Specification:
System variable desired in output
Structural unit of the output variable
Output ontology

Topological Configuration

In case of topological configuration of a circuit, each model of a configuration involves a specific qualitative state of the circuit as the underlying assumption for the configuration. Figure 6 shows three topological configurations of the half-wave rectifier circuit when the simulation is analyzing a sinusoidal input voltage to the circuit. Each configuration involves only a subset of the system components. If the assumption becomes violated as a new qualitative state arises during simulation, the program identifies and shifts to the new configuration model which satisfies the state and where reasoning continues.

The configuration process performs the crucial zero-order reasoning of current flow. One important property of a configuration is that the individual devices within the configuration, linear or nonlinear, behave linearly over the duration of the given qualitative state. This provides a configuration-wise linearization method for *automating* qualitative analysis of nonlinear systems. Although nonlinear devices have been modeled piecewise linearly in previous systems [de Kleer 1984, Williams 1984], the correct configuration cannot be determined in the simulation with "small signal analysis". Thus, the choice of which causal rules to us in any state can only be made under the assumption that the device is operating with that particular state [de Kleer 1984]. This is partly due to the lack of control in the envisioning process.

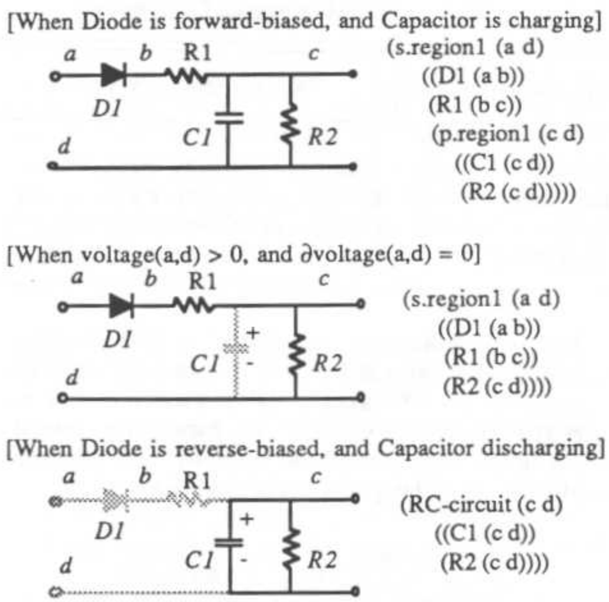

Figure 6: Shifting in the dimension of topological configurations

In our work, the representation of an active configuration of a circuit stands as a model which embodies a 1D perspective, revealing the dynamic configuration of the system topology in a given qualitative state. ARC saves the configuration model together with the assumption about the states of nonlinear devices in the circuit. For future tasks, ARC can select a configuration model previously created if the qualitative state of the target circuit satisfies the modeling assumption.

Structural Aggregation

To derive appropriate structural granularity of a circuit for a given task, two aspects are of relevance to qualitative simulation. First, the inherent structural hierarchy of the system topology of the circuit is revealed by clustering individual devices at various levels of physical organization [Farley 1988]. Reasoning can then traverse the hierarchy. At any given level, lower-level details are suppressed.

Second, different tasks may suggest grouping subsets of structural units as black boxes to suppress extraneous details for the task at hand. As a result, there are three types of regions in a given configuration: serial regions, parallel regions, and regions of individual devices.

For example, suppose the active topological configuration of the rectifier circuit is the one when the diode is forward-biased and the capacitor is charging. Figure 7 shows two ways of black boxing this configuration for two different tasks.

(1) If the output of a task indicates the behavior of voltage between nodes c and d, then ARC aggregates the structure of the circuit and creates two black boxes, as shown in Figure 6.1.

(2) If the output of a task indicates the behavior of the diode in this configuration, then R1, R2, and C1 are grouped as a black box, as shown in Figure 6.2.

Black boxes play the role of "equivalent circuits", as are commonly used in electrical engineering. A complex circuit can be viewed as a simpler one for analysis if they

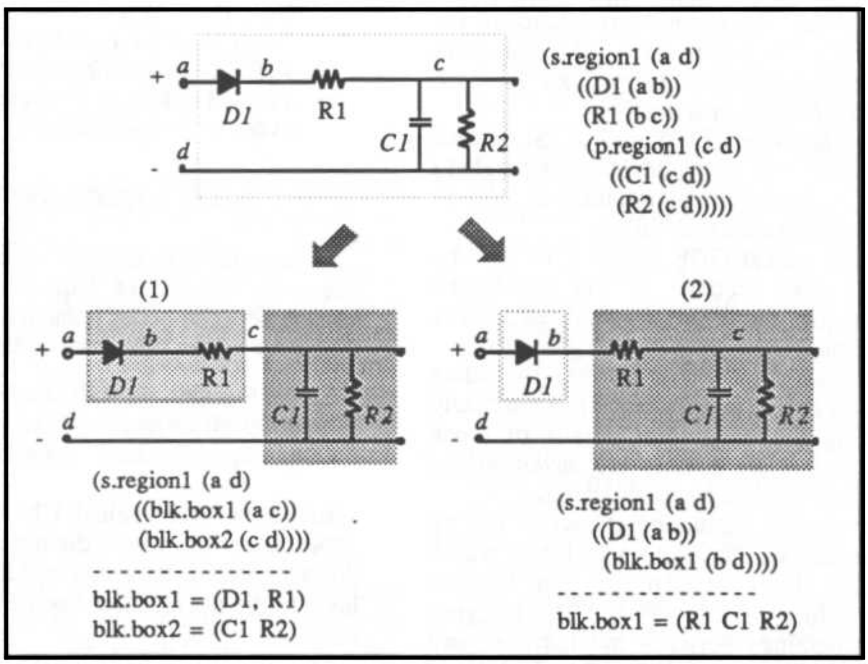

Figure 7: Shifting in the dimension of structural aggregation

are considered equivalent. Sussman & Steele (1980) introduced a constraint language to declare slices to aggregate individual devices to avoid algebra. But the responsibility to create instances of slices in their system rests with the user. ARC extends the notion of slices by automatically creating appropriate structural granularities as driven by the task at hand.

A description of the structural aggregation over a topological configuration stands as a model embodying a 2D perspective. This perspective reveals not only the dynamic configuration of the circuit, but also the chosen structural granularity with appropriate black boxes encapsulating individual devices in the configuration.

Ontological Choice

In ARC, models along the ontological choice dimension are represented as QDE models. We use the standard device ontology at the macroscopic level and the charge-carrier (CC) ontology at the microscopic level to reason about electronics circuits. For ontological shift, we use a set of domain-independent ontological-choice rules based on the task at hand to control ontological choice, as described in [Liu & Farley 1990]. The input and output specifications in a task definition are stated in chosen language terms, which indicate the input and output ontologies. This information prompts ARC to select appropriate language to formulate QDE models of the active configuration of the target circuit.

Construction of QDEs for a parallel or serial region involves three steps:

- First, the region is treated as a single structural unit. As such, the descriptions of Ohm's Law and voltage compatibility for the unit are generated.
- Second, the relationships among the sub-regions in the region are described. For example, resistance of the region is summarized of its sub-regions. The continuity condition describing Kirchhoff's Current Law for each connection node in a serial region is described.
- Finally, each sub-region itself is described. Since each of the sub-regions is either a single device, or a parallel / serial region, it can recursively use the same procedure as for the given region to create its own QDEs.

Figure 8 shows the actual QDE model generated by ARC for the half-wave rectifier in the aggregated configuration where the diode is forward-biased. For different tasks, different QDE models may be generated.

Black-boxes are treated as individual devices. The hiding of irrelevant structural details inside a black box ultimately leads to suppressing the behavioral details of those structural units. The QDE models are saved in the knowledge base with the associated configurations.

Ontological shift during simulation is carried out by switching QDE models via "bridging relations", which associate comparable structural elements from the two related ontologies [Liu & Farley 1990]. The bridging relations involve mappings between macroscopic and microscopic concepts of electronics. Table 2 presents some bridging relations between the device ontology and the CC ontology.

TABLE 2. Mapping Between Macroscopic and Microscopic Descriptions

Macroscopic	bridge	Microscopic (CC)
voltage	<=>	intensity of field in region;
resistance	<=>	physical characteristics of region;
current	<=>	charge-carrier movement in region;
capacitance	<=>	pole size and distance in region.

The perturbation is propagated in QDE models by QUALEX [Douglas & Liu 1989]. As a causal inference engine, QUALEX takes a QDE model and a single perturbation to one of the circuit parameters, propagates the change throughout the circuit, and generates a qualitative, causal explanation of the behavior.

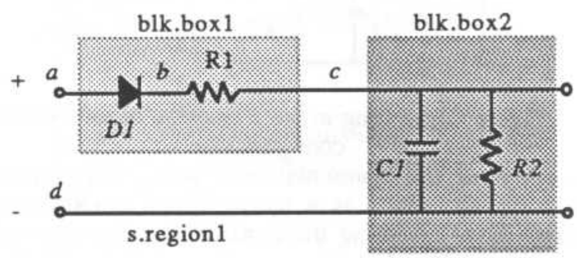

$\partial V@\text{a\_d} - \partial I@\text{a\_d} - \partial R@\text{s.region1} = 0.$

$\partial V@\text{a\_d} - \partial V@\text{a} + \partial V@\text{d} = 0.$

$\partial R@\text{s.region1} - \partial R@\text{blk.box1} - \partial R@\text{blk.box2} = 0.$

$\partial I@\text{a\_d} - \partial I@\text{a\_c} = 0.$

$\partial I@\text{a\_c} - \partial I@\text{c\_d} = 0.$

$\partial V@\text{a\_c} - \partial I@\text{a\_c} - \partial R@\text{blk.box1} = 0.$

$\partial V@\text{a\_c} - \partial V@\text{a} + \partial V@\text{c} = 0.$

$\partial V@\text{c\_d} - \partial I@\text{c\_d} - \partial R@\text{blk.box2} = 0.$

$\partial V@\text{c\_d} - \partial V@\text{c} + \partial V@\text{d} = 0.$

Figure 8: A QDE model generated by ARC

As an example, suppose one is interested to ask the question: "Given the half-wave rectifier, if the input voltage between nodes *a* and *d* increases, what happens to the charge carrier flow from nodes *c* to *d*?"

- Target System: [half-wave-rectifier1]
- Input Specification: (voltage (a d) (+ +))
- Output Specification: (cc-flow (c d))

Based on Ontological Choice Rules [Liu & Farley 1990], this task, where the input and output are stated in different ontologies, requires shifting from the macroscopic level to the microscopic level during reasoning. The rule says:

If the input and the output variables are from different ontologies, proceed with the input ontology until causal propagation comes to the region of the output variable and then shift to the output ontology to complete the reasoning.

Accordingly, ARC generates two QDE models from the active configuration of the circuit, one in the device ontology and the other in the CC ontology. Reasoning proceeds in the device-ontology model until it comes to the region of the output variable where it switches to the CC-ontology model, shifting in the dimension of ontological choices. The actual output generated by ARC is as follows:

- In the serial region between nodes (a d), since resistance@s.region1 remains constant, voltage@a_d's increase causes current@a_d to increase.
- current@a_d's increase means current@a_c's increase because they two originate from the same source and share the same path in the configuration.
- According to KCL with regard to node c for a serial connection, we know that current@a_c's increase leads to the increase of current@c_d.
- For blk-box2 between nodes (c d), since resistance@blk.box2 remains constant, current@c_d's increase causes voltage@c_d to increase.

—— Cross a bridge to charge-carrier-ontology ——

- In the region between nodes (c d), voltage@c_d's increase is equivalent to field@c_d's increase.
- field@c_d's increase causes force-on-cc@c_d to increase.
- force-on-cc@c_d's increase causes cc-velocity@c_d to increase.
- cc-velocity@c_d's increase causes cc-flow@c_d to increase.

Discussion

We have presented ARC, a task-driven, perspective-taking approach to formulating models of a target system for qualitative causal analysis. Each model embodies a perspective, which is a position taken in one or more possible dimensions. ARC is able to reason about a circuit from different perspectives of a circuit and to switch among them as needs dictate. It is precisely this ability of task-driven model formulation that offers ample opportunity for improving the flexibility and efficiency of automated reasoning systems.

Related Work

A number of researchers in qualitative physics have recently described systems that reason with multiple models to great advantage. The graph-of-models approach by Addanki et al. (1989) and Weld (1989) is effective to the extent that all the possible models required by future tasks can be pre-defined, but incurs a large amount of storage and trial-and-error search for a valid model.

The approaches by Sussman & Steele (1980), Davis (1984) and Genesereth (1985) used multiple levels of structural descriptions to control search. When reasoning concerns high-level components, the internal detail of the components is suppressed. Our approach shares the same motivation as theirs. However, we are able to automate structural aggregation, rather than resort to a predefined structural hierarchy.

The approach closest to ours is compositional modeling by Falkenhainer & Forbus (1988, 1991). This technique uses a fine-grained modular approach to modeling. The fragments of a general domain model are attached with explicit modeling assumptions, each describing various aspects of the domain. Based on a query, it uses "consider-assumptions" to trim the domain model to generate one scenario model, which suffices to answer the query while minimizing extraneous detail. As in ARC's task definition, the terms in the query provide significant constraint in identifying a set of modeling assumptions and associated model fragments. By contrast, ARC can generate, not just one, but several models for a single task and reason in the space of possible dimensions, shifting perspective as commensurate with the needs of the task. In addition, the modeling assumption for a model in ARC is generated in the simulation process and then stored with the model in the knowledge base for possible re-use by future tasks. Models formulated are guaranteed consistent because only a single position is taken on each possible dimension. If there are more than one position taken in some dimension, the model may become inconsistent or intractable.

Limitations and Future Directions

We have just begun to study the problem of reasoning with task-driven perspective-taking. We have applied the techniques described in this paper to the domains of electronic circuits and hydraulic systems. Modeling in ARC critically depends on the implicit assumption that influences are confined to known connections between components. While ARC shows a promising direction for automated reasoning, there are several important limitations in the current work. Some of them were deliberately chosen so we could focus on the central issues of this research. Some of them are due to the present implementation choice within ARC.

This research has focused entirely on qualitative causal reasoning from structure to behavior. The teleological aspect of physical systems regarding their purposes or functions are not addressed. We have only examined structural aggregation based on parallel-serial reductions. While series-parallel circuits are the norm, there are circuits that cannot be aggregated using parallel-serial reduction. As such, ARC cannot aggregate structural components that form a functional unit. For example, a circuit can be aggregated as a black box and described as an *adder* or a *CPU*. Aggregation based on teleology on function remains a challenge. de Kleer (1984) noted that a single resister could have multiple purposes in a circuit. Experts also use the geometrical features [Palies, et al. 1986] of the circuit diagram to recognize functional components, such as an amplifier, in circuit analysis. We realized that a complete causal explanation of how a physical device works requires the knowledge of purpose [Downing 1990]. One feasible way for doing this would be to access the design knowledge to guide the structural aggregation.

The tasks given to ARC are queries made by the user. Although ARC specifies tasks with both input and output, capable of "large signal analysis", the current version of the task definition language is not sufficient to express tasks beyond the parameter-perturbation analysis problems.

Improvement of a task definition language like the one developed in this work will be tightly coupled with the advances of theories and techniques in automated reasoning research. In ARC, a task consists of terms in chosen ontologies with nodes implying the structural units in a target circuit. This suggests an important framework for task analysis with improved task definition languages in the future. Given a more general task definition language, how can its form be used to suggest appropriate perspectives for modeling and reasoning in automated reasoning systems?

The representation of charge carriers as pieces of stuff is rather limited in ARC. Complex analytic tasks require spatial and temporal reasoning about charge carriers as pieces of stuff. We suspect that this direction for future research offers grounds to develop new ontologies to reason about complex behaviors of real-world systems. We only considered two related ontologies. The interesting question to ask is how we can extend the current framework to more than two ontologies. In the presence of multiple ontologies, do we need pair-wise bridges between each pair of ontologies? How do we control ontological shift among multiple ontologies?

The current implementation of ARC does not generate complete envisionment of all possible behaviors when ambiguity occurs. Instead of envisioning, ARC handles the ambiguity problem by turning the control over to the user and letting the user make a choice so that only behavior interesting to the user is examined. ARC only allows the user to make one choice in the situation and then focuses reasoning on this particular choice. The current implementation does not try to predict all the possible future behaviors. Part of the reason for doing so is that ARC can follow a given sequence of perturbations that drive the qualitative simulation. This helps remedy the envisioning process that ignores control issues in qualitative simulation by explicitly generating the entire search space [Forbus 1988]. In some situations, however, it was felt that the user should be given the choice to turn on an envisioner to see all possible future behaviors of the target system at a time.

Conclusion

In this research, we have investigated the general problem of how to automate reasoning about physical systems with multiple perspectives. we have begun specifying a theory of task-driven perspective-taking for effective reasoning about complex physical systems. The results of this research show that by using an integrated framework of topological configuration, structural aggregation, and ontological shift, we can extend the range of qualitative causal reasoning about complex physical systems.

References

[Addanki, et al. 1989] S. Addanki, R. Cremonini, & J. S. Penberthy, "Reasoning about assumptions in graphs of models". In *Proceedings of IJCAI-89*, pp. 1432-1438, Detroit, MI, August 1989.

[Arnborg & Proskurowski 1989] S. Arnborg & A. Proskurowski, *Canonical Representations of Partial 2- and 3-trees*. CIS-TR-89-11, Department of Computer and Information Science, University of Oregon. July, 1989.

[Bell 1980] D. A. Bell. *Electronic Devices and Circuits.* Reston Publishing Company, Inc., Reston, VA, 2 edition, 1980.

[Collins & Forbus 1987] J. W. Collins and K. D. Forbus, "Reasoning about Fluids via Molecular Collections". In *Proceedings AAAI-87*, pp. 590-594, Seattle, WA, 1987.

[Davis & Hamscher 1988] R. Davis & W. Hamscher, "Model-based reasoning: Troubleshooting". In H. E. Shrobe, editor, *Exploring Artificial Intelligence*, chapter 8, pp. 297-346, Morgan Kaufmann, San Mateo, CA, 1988.

[de Kleer 1984] J. de Kleer, "How circuits work". *Artificial Intelligence*, 24:205-280, 1984.

[de Kleer & Brown 1984] J. de Kleer & J. S. Brown, "A qualitative physics based on confluences". *Artificial Intelligence*, 24:7-83, 1984.

[Douglas & Liu 1989] Douglas, S. A. and Liu, Z. Y., "Generating Causal Explanation from a Cardio-Vascular Simulation". In Proceedings IJCAI-89, pp. 489-494. Detroit, MI., 1989.

[Downing 1990] K. L. Downing, *The Use of Teleology in the Qualitative Evaluation and Explanation of Circulatory Systems*, Ph.D Thesis. University of Oregon, 1990.

[Falkenhainer & Forbus 1988] Falkenhainer, B. & Forbus, K. D. 1988, "Setting Up Large Scale Qualitative Models". In *Proceedings AAAI-88*, pp. 301-306, Saint Paul, Minn.

[Falkenhainer & Forbus 1991] Falkenhainer, B. & Forbus, K., "Compositional Modeling: Finding the Right Model for the Job". *Artificial Intelligence*, special issue on qualitative physics, Fall, 1991.

[Farley 1988] A. M. Farley, "Cluster-based representation of hydraulic systems". In *Proceedings of the 4th Conference on Artificial Intelligence Applications*, pp. 358-364, San Diego, CA, March 1988.

[Forbus 1988] K. D. Forbus, "Qualitative Physics: Past, Present, and Future". In H. E. Shrobe and AAAI, editors, *Exploring Artificial Intelligence*, chapter 7, pp. 239--296. Morgan Kaufmann, Inc, San Mateo, California, 1988.

[Genesereth 1985] M. Genesereth. "The Use of Design Description in Automated Diagnosis", in Bobrow, editor, *Qualitative Reasoning about Physical Systems*. pp. 411-436. MIT Press, Cambridge, Massachusetts, 1985.

[Hayes 1985a] P. Hayes, "Naive Physics I: Ontology for Liquids". In J. R. Hobbs and R. C. Moore, editors, *Formal Theories of the Commonsense world*, chapter 3, pp. 71-107, Ablex, Norwood, NJ., 1985

[Hobbs 1985] J. R. Hobbs, "Granularity". In *Proceedings of IJCAI-85*, pp. 432--435.

[Iwasaki & Bhandari 1988] Y. Iwasaki and I. Bhandari, "Formal Basis for Commonsense Abstraction of Dynamic Systems". In *Proceedings of AAAI-88*, pp.307-312, Saint Paul, Minnesota, August 1988.

[Kuipers 1986] B. Kuipers, "Qualitative simulation". *Artificial Intelligence*, (29):289-338, 1986.

[Kuipers 1987] B. Kuipers, "Abstraction by Time-Scale in Qualitative Simulation". In *Proceedings of AAAI-87*, pp. 621-625.

[Liu & Farley 1990] Z. Y. Liu & A. M. Farley, "Shifting Ontological Perspectives in Reasoning about Physical Systems". In *Proceedings of AAAI-90*. pp. 395-400. Boston, August, 1990.

[Liu & Farley 1990] Z. Y. Liu & A. M. Farley, "Structural Aggregation in Commonsense Reasoning". In *Proceedings of AAAI-91*. Anaheim, California, July, 1990.

[Liu 1991] Z. Y. Liu, *Qualitative Reasoning about Physical Systems with Multiple Perspectives*. Ph.D Dissertation, University of Oregon. March, 1991.

[Palies et al. 1986] O. Palies, M. Caillot, E. Cauzinille-Marmeche, J. Lauriere and J. Mathieu, "Student Modelling by a Knowledge-Based System". *Computational Intelligence*, (2):99-107, January 1986.

[Smith & Davis 1989] B. Smith and R. Davis, "Expert Systems: How Far Can They Go?" *AI Magazine*, 10(2):65-77, Summer (A panel session at IJCAI-85 in Los Angeles.), 1989.

[Sussman & Steele 1980] G. J. Sussman & G. L. Steele, Constraints - a language for expressing almost-hierarchical descriptions. *Artificial Intelligence*, 1(36):1-39, 1980.

[Weld 1989] D. S. Weld, "Automated model switching". In *Proceedings of the 3rd Int. Qualitative Physics Workshop*, Palo Alto, CA, August 1989.

[Weld & Addanki 1990] D. S. Weld & S. Addanki, "Task-Driven Model Abstraction", In *Proceedings of the 4th Int' Workshop on Qualitative Physics*. Lugano, Switzerland, July, 1990.

[Williams 1984] B. C. Williams, "Qualitative analysis of MOS circuits". *Artificial Intelligence*, 24:281-346, 1984.