Automatica, Vol. 25, No. 4, pp. 571–585, 1989
Printed in Great Britain.

0005-1098/89 $3.00 + 0.00
Pergamon Press plc

Qualitative Reasoning: Modeling and Simulation with Incomplete Knowledge*

BENJAMIN KUIPERS†

Incomplete knowledge during diagnosis or design can make ODE models difficult to apply; a qualitative model can express incomplete knowledge of landmark values and functional relationships, and allow useful qualitative predictions to be derived.

Key Words—Artificial intelligence; modeling; simulation; qualitative reasoning; qualitative simulation.

Abstract—Recently developed methods for qualitative reasoning may fill an important gap in the modeling and control toolkit. Qualitative reasoning methods provide greater expressive power for states of incomplete knowledge than differential or difference equations, and thus make it possible to build models without incorporating assumptions of linearity or specific values for incompletely known constants. Even with incomplete knowledge, there is enough information in a qualitative description to support qualitative simulation, predicting the possible behaviors of an incompletely described system. We survey results from several approaches to qualitative reasoning, and provide a detailed example of the application of these methods to a simple problem. The mathematical validity of qualitative simulation is also assessed. Initial results have been encouraging, and steps are now being taken to develop additional mathematical power, hierarchical decomposition methods, and incremental quantitative constraints, to make qualitative reasoning into a formal reasoning method useful on realistic problems.

1. INTRODUCTION

THE GOAL of artificial intelligence research is to develop symbolic computational methods for representing knowledge and processes of inference.

One of the fruits of AI research is the rule-based expert system. The MYCIN system (Shortliffe, 1976), applied to the problem of diagnosis, is the best and clearest early example of this method.

1.1. *The need for model-based reasoning*

While the rule-based approach yielded impressive results, further experience and analysis showed that there were serious limitations to an inference method based on empirical associations between observable findings and diagnostic hypotheses. Within the rule-based approach, it is difficult to reason about interactions between diseases, or idiosyncratic responses to change. New methods were proposed for incorporating causal relations as well as empirical associations (Patil *et al.*, 1981; Pople, 1982).

This line of investigation produced various approaches to model-based diagnosis, using an interplay between an empirical association-based reasoner which proposes hypotheses, and a model-based reasoner which elaborates and evaluates them (Patil *et al.*, 1981; Simmons and Davis, 1987).

Empirical studies of causal reasoning by expert physicians showed that one important representation for the model of a mechanism consisted of qualitative descriptions of continuous variables, their directions of change, and the constraints among them (Kuipers and Kassirer, 1984).

This observation supported a convergence between research on model-based diagnosis and research on qualitative reasoning about physical systems (de Kleer, 1977; de Kleer and Brown, 1984; Forbus, 1984; Kuipers, 1984). The model-based reasoner has the task of taking a qualitative description of the mechanism and its current state, and producing a qualitative description of its possible behaviors.

Current AI research in model-based reasoning is attempting to develop the framework sketched in Fig. 1, extending the power of qualitative model-based reasoning methods, and developing

* Received 30 June 1988; revised 6 October 1988; received in final form 18 January 1989. The original version of this paper was presented at the IFAC Symposium on Modelling and Control in Biomedical Systems which was held in Venice, Italy, during April 1988. The Published Proceedings of this IFAC meeting may be ordered from: Pergamon Press plc, Headington Hill Hall, Oxford OX3 0BW, U.K. This paper was recommended for publication in revised form by Editor A. P. Sage.

† Department of Computer Sciences, University of Texas at Austin, Austin, TX 78712, U.S.A.

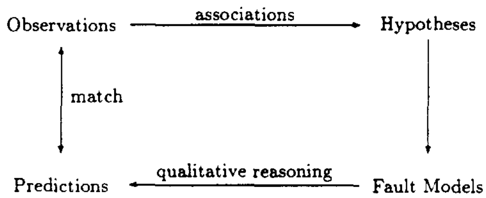

FIG. 1. Model-based reasoning in a generate-and-test cycle for diagnostic problem-solving.

the appropriate style of interaction between the two reasoning methods. Although this framework is defined for the problem of diagnosis, the problems of monitoring and control are closely related. We may informally paraphrase the problems of monitoring and control in terms of diagnosis and prediction, as follows.

• Monitoring: "What is really going on, given what I observe?"
• Control: "What should I do next?"→ "What will happen if I do X?"

In this overview, I will concentrate on the problem of qualitative model-based reasoning, and particularly on the problem of prediction. I will attempt to show the relations between the different major approaches to this problem, and how they address the concerns of the systems modeling community. I will conclude by discussing current research directions and progress toward application to realistic systems.

1.2. *The need for incomplete knowledge*

How, then, can we represent models of mechanisms capable of providing useful predictions about possible behaviors? We are faced with several conflicting requirements.

(1) The model should express what we know about a mechanism.
(2) The model should not require assumptions beyond what we know.
(3) It must be mathematically and computationally feasible to derive predictions.
(4) It should be possible to match predictions against observations.

There is a spectrum of current approaches to building models, with advantages and disadvantages to each point on the spectrum. Qualitative reasoning fills a gap in this spectrum between some of the more traditional methods.

• Numerical simulation of difference equations.
• Analytical solution of differential equations.
• (Gap.)
• Influence diagrams.

Numerical difference equations are precise and computationally tractable. They require us to specify computable functions relating the variables in the equations, which frequently involves using linear approximations to unknown or non-linear functions, and providing numerical values for incompletely known constants. These functional relations and constant values amount to additional assumptions beyond what we know. A numerical prediction never matches the observed values precisely, but there is an enormous literature of methods for describing and quantifying the match between prediction and data.

Differential equations have many of the same strengths, and in addition make it possible to withhold commitment on the numerical values of constants. However, functions relating variables must be specified explicitly, even where they are not known with confidence. The most serious problem with this representation is that many differential equations cannot be solved analytically at all, especially if some of the functional relations are nonlinear.

Influence diagrams (signed directed graphs) are used to describe how the direction of certain changes influence the direction of change of other variables. Their strength is in representing a state of partial knowledge during model creation. An influence diagram specifies the relevant variables and the relevant functional relations, without making any commitment to their precise form. Sophisticated methods exist for assessing the stability or instability of complex feedback systems expressed as influence diagrams (Puccia and Levins, 1985). In case the system is stable, the effects of perturbations can be predicted, within a quasi-static equilibrium assumption.

1.3. *The need for qualitative reasoning*

Qualitative reasoning methods provide an intermediate point on this spectrum. They provide more expressive power for states of incomplete knowledge than differential or difference equations, by means of a low-resolution *quantity space* representation for values, and classes of monotonic functions for functional relations. They provide more inferential power than influence diagrams by applying limit analysis to changing variables, and by having a mathematically precise semantics as an abstraction of differential equations. Matching and prediction are both facilitated because the landmarks of the quantity space are defined by semantically important points where important changes take place.

Several distinct approaches to qualitative reasoning have developed, and the relations among them are not always clear.

• de Kleer and Brown, "A qualitative physics based on confluences" (de Kleer and Brown, 1984).

Johan de Kleer first explored the properties of qualitative representations of mechanisms. His method using confluences assumes that the mechanism is always in, or very near, a state of equilibrium.

• Forbus, "Qualitative process theory" (Forbus, 1984).

Qualitative process theory is a model-building methodology, which recognizes the elements of a model from a physical description of a system, then applies a closed-world assumption to create the appropriate set of constraints.

• Kuipers, "Qualitative simulation" (Kuipers, 1984, 1986).

Qualitative simulation starts with a set of qualitative constraints and an initial state, and predicts the set of possible futures for the system. This approach has a precise mathematical semantics as an abstraction of differential equations.

In the following sections of this paper, we will present the concepts of qualitative reasoning, and clarify the relations among the different approaches through a simple example. Finally, we will discuss the mathematical strengths and limitations of this approach, and briefly sketch current research directions.

2. CONCEPTS OF QUALITATIVE SIMULATION

In order to introduce the concepts of qualitative simulation, we will go through an elementary modeling exercise, illustrating how qualitative modeling differs from previous methods.

We can explore these concepts by looking at the U-tube: a simple, two-tank fluid flow system. The structure of the U-tube is simple and clear, and its behavior is easily deduced. Nonetheless, there is enough complexity to motivate the different features of qualitative simulation systems, and we will use the U-tube as a model for more complex systems. In particular, the U-tube is a simple analog for flow between two physiological compartments.

The U-tube consists of two tanks (named A and B) connected by a flow channel (Fig. 2). We assume that the momentum of water flowing through the channel is not a significant factor.

2.1. *The qualitative structure of the U-tube*

Each tank holds a certain amount of water, which produces a certain pressure at the bottom.

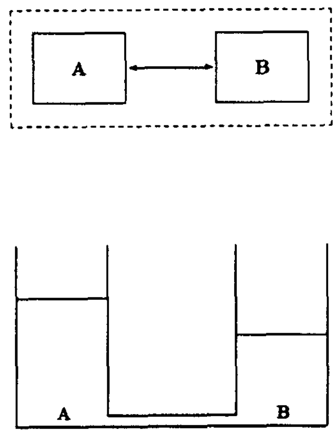

FIG. 2. The U-tube is a simple example of a two-compartment system.

Since the U-tube is a system which may change over time, these magnitudes are continuous real-valued functions of time: amtA(t), amtB(t), pressureA(t), and pressureB(t).

Reasoning qualitatively, and not knowing the exact properties of each tank, we know that the pressure increases with the amount of water in the tank, but we do not know the exact relationship. We express this state of incomplete knowledge by saying that there is some unspecified monotonically increasing function relating *amount* to *pressure*

$$\text{pressure} = M^+(\text{amount}). \qquad (1)$$

The term M^+ refers to an unspecified member of the class of monotonically increasing functions.

We also have incomplete, qualitative knowledge about the values taken on by the parameters amtA and pressureA at any given time. Instead of specifying numerical values in the real number line, the value of a parameter is described qualitatively in terms of its *quantity space*. Each quantity space is defined by an ordered set of *landmark values*. The quantity spaces for amtA and pressureA are

$$\begin{array}{ll} \text{amtA} & 0 \cdots \text{AMAX} \cdots \infty \\ \text{pressureA} & 0 \cdots\cdots\cdots\cdots \infty \end{array}$$

The landmarks $-\infty$, 0, and ∞ have known properties and may be used in any quantity space. Other landmarks are symbols the meanings of which are specific to the particular quantity space, and are defined only by their relations with other landmarks. Thus, the landmark AMAX, representing the maximum capacity of tank A, is somewhere in the interval $(0, \infty)$, but otherwise unspecified. Since amtA and pressureA are necessarily non-negative, the low bounds of the quantity spaces are both 0, and negative values cannot even be described.

In the case of the U-tube, we have additional knowledge about the M^+ function relating amtA and pressureA. We know that it passes through the origin: there is a correspondence between amtA = 0 and pressureA = 0. (Following normal mathematical usage, we will write the correspondence between 0 and 0 as $(0, 0)$.) We also assert that the relation pressureA = M^+(amtA) has a correspondence at (∞, ∞), which has the effect of excluding a horizontal or vertical asymptote.

There is a similar (but not necessarily identical) monotonically increasing function relating amtB to pressureB, so we can write

$$\text{pressureB} = M^+(\text{amtB}). \tag{2}$$

This illustrates that M^+ is not the name of a function, but the name of a class of functions.

The pressure difference between the two tanks is another important continuous function

$$p\text{AB} = \text{pressureA} - \text{pressureB}. \tag{3}$$

The flow between A and B is a monotonic function of the pressure difference, with corresponding values $(0, 0)$ and (∞, ∞)

$$\text{flowAB} = M^+(p\text{AB}). \tag{4}$$

Finally, we observe that flowAB represents the rate of change of amtB and the inverse of the rate of change of amtA

$$\text{flowAB} = \frac{\mathrm{d}}{\mathrm{d}t}\text{amtB} \tag{5}$$

$$\text{flowAB} = -\frac{\mathrm{d}}{\mathrm{d}t}\text{amtA}. \tag{6}$$

These *qualitative constraints* (equations (1)–(6)), along with their corresponding values and the quantity spaces of their parameters, describe the qualitative structure of the U-tube system. Figure 3 shows a graphical form of this set of constraints. Since some of our computations later on will involve propagation of information across constraints, a graphical notation helps clarify the information flow.

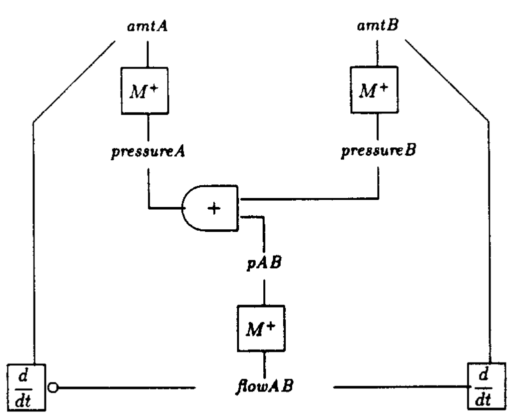

FIG. 3. U-Tube constraints.

The quantity spaces for the parameters of this system are

amtA	$0 \cdots \text{AMAX} \cdots \infty$
pressureA	$0 \cdots \infty$
amtB	$0 \cdots \text{BMAX} \cdots \infty$
pressureB	$0 \cdots \infty$
pAB	$-\infty \cdots 0 \cdots \infty$
flowAB	$-\infty \cdots 0 \cdots \infty$.

As will be discussed below, the quantity spaces define a descriptive "language" representing the set of qualitative distinctions that can be expressed by this model, at least at its initial state.

The constraints and corresponding values are:

constraints	*corresponding values*
pressureA = M^+(amtA)	$(0, 0), (\infty, \infty)$
pressureB = M^+(amtB)	$(0, 0), (\infty, \infty)$
pressureA − pressureB = pAB	
flowAB = M^+(pAB)	$(0, 0), (\infty, \infty)$
flowAB = $\frac{\mathrm{d}}{\mathrm{d}t}$amtB	
flowAB = $-\frac{\mathrm{d}}{\mathrm{d}t}$amtB.	

We will also consider that these constraints might be written as a *qualitative differential equation* (QDE)

$$\frac{\mathrm{d}}{\mathrm{d}t}\text{amtB} = f(g(\text{amtA}) - h(\text{amtB})), f, g, h \in M^+$$

by collapsing the constraint equations together. Since f, g, and h are unknown monotonic functions, and may well be nonlinear, this equation is intractable analytically. Later, we will discuss how a set of constraints like this can be derived from a physical description of a mechanism.

3. QUALITATIVE KNOWLEDGE OF STATE

Like our structural knowledge—the qualitative constraints and quantity spaces—we have incomplete knowledge of the state of the U-tube at any given moment. Its state is described by the qualitative values of the parameters: amtA, amtB, pressureA, pressureB, pAB, and flowAB.

3.1. *Qualitative state is dynamic*

As the system changes with time, we will need to describe the sequences of qualitative states that it may pass through. Suppose that I am filling tank A until it overflows, ignoring tank B. The sequence of qualitative magnitudes of the parameter amtA will be

$$\text{amtA}(t):\quad 0 \rightarrow (0, \text{AMAX}) \rightarrow \text{AMAX}$$

meaning that after the time that $\text{amtA}(t) = 0$ and until $\text{amtA}(t) = \text{AMAX}$ the value of $\text{amtA}(t)$ is somewhere in the open interval $(0, \text{AMAX})$. This is not the whole story, however.

In order to predict a behavior like that of amtA, we will need to determine the transitions from one qualitative state to another. In addition to qualitative magnitude in the quantity space, we need to know the direction of change of each parameter. Thus, for each parameter, we describe its qualitative state in terms of its magnitude (in the quantity space), and its direction of change: increasing, decreasing, or steady. According to the sign of the derivative of a parameter at a particular instant, we describe the direction of change as inc, dec, or std.

Thus, the qualitative state of amtA in the midst of filling has two components—magnitude and direction of change—and is written $\langle (0, \text{AMAX}), \text{inc} \rangle$. When we need to refer to the direction of change alone, we will write, e.g. $\text{qdir}(\text{amtA}) = \text{inc}$. When tank A is filled to overflowing, the behavior of amtA over time can be described as

$$\text{amtA}(t):\quad \langle 0, \text{inc} \rangle \rightarrow \langle (0, \text{AMAX}), \text{inc} \rangle \rightarrow \langle \text{AMAX}, \text{inc} \rangle.$$

We can capture the same description of the behavior of an individual parameter with a *qualitative graph,* in which the landmark values in the quantity space are arranged on the vertical axis, and the distinguished time points along the horizontal axis (cf. Fig. 5). Qualitative values are plotted at, or midway between, landmark points, and the symbol plotted $(\uparrow, \downarrow, \ominus)$ represents the direction of change.

But before we can predict behavior, we must define an initial state.

4. PREDICTING BEHAVIOR FROM INITIAL CONDITIONS

A qualitative simulation problem specifies a structure in terms of parameters, their quantity spaces, and their constraints with corresponding values. Then an initial state is given, and we want to know the possible behaviors.

Suppose we start with tank A full and tank B empty. We thus begin with the following information:

$$t = t_0 \Rightarrow \boxed{\begin{array}{l} \text{amtA} = \langle \text{AMAX}, ? \rangle \\ \text{amtB} = \langle 0, ? \rangle. \end{array}}$$

In order to simulate the system, we need a *complete* qualitative state description: a qualitative magnitude and direction of change for each parameter.

4.1. *Propagating to the complete initial state*

Given the initially specified information, we can propagate locally across constraints to determine more about the qualitative state (Fig. 4). Here we are able to propagate to determine the initial state of the U-tube completely.

(1) Because of the corresponding value at $(0, 0)$

$$\text{amtB} = 0 \Rightarrow \text{pressureB} = 0.$$

(2) Because of corresponding values at $(0, 0)$ and (∞, ∞), we can conclude that

$$\text{amtA} = \text{AMAX} \Rightarrow \text{pressureA} = (0, \infty).$$

Because of the absence of landmarks, this is the best description for the value of pressureA expressible in its quantity space.

(3) An addition (or subtraction) constraint has an implicit set of corresponding values at $(0, 0, 0)$, so

$$\text{pressureA} = (0, \infty) \wedge \text{pressureB} = 0 \Rightarrow p\text{AB} = (0, \infty).$$

(4) The corresponding value at $(0, 0)$ of the constraint $\text{flowAB} = M^+(p\text{AB})$ then gives us

$$p\text{AB} = (0, \infty) \Rightarrow \text{flowAB} = (0, \infty).$$

(5) The derivative constraints can now determine the directions of change of amtA and amtB

$$\text{flowAB} = (0, \infty) \Rightarrow \text{qdir}(\text{amtA}) = \text{dec} \wedge \text{qdir}(\text{amtB}) = \text{inc}.$$

(6) The directions of change propagate through monotonic functions

$$\text{qdir}(\text{amtA}) = \text{dec} \Rightarrow \text{qdir}(\text{pressureA}) = \text{dec}$$
$$\text{qdir}(\text{amtB}) = \text{inc} \Rightarrow \text{qdir}(\text{pressureB}) = \text{inc}.$$

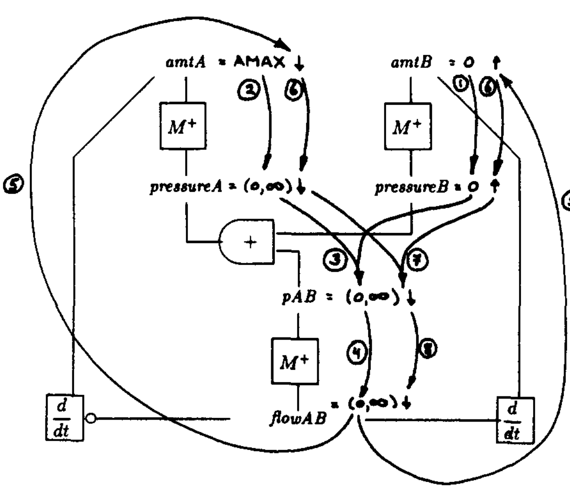

FIG. 4. Propagation through the constraint model.

(7) Directions of change also propagate through the addition constraint

$$\text{qdir(pressureA)} = \text{dec} \wedge \text{qdir(pressureB)} = \text{inc} \Rightarrow \text{qdir}(p\text{AB}) = \text{dec}.$$

(8) The last propagation completes the state

$$\text{qdir}(p\text{AB}) = \text{dec} \Rightarrow \text{qdir(flowAB)} = \text{dec}.$$

The final result of this propagation is a complete, qualitative, description of the state of the U-tube at the initial instant, t_0.

$t = t_0 \Rightarrow$

amtA	$= \langle \text{AMAX}, \text{dec} \rangle$
pressureA	$= \langle (0, \infty), \text{dec} \rangle$
amtB	$= \langle 0, \text{inc} \rangle$
pressureB	$= \langle 0, \text{inc} \rangle$
pAB	$= \langle (0, \infty), \text{dec} \rangle$
flowAB	$= \langle (0, \infty), \text{dec} \rangle$.

Although these inferences propagate around a network of constraints, they do not correspond to changes taking place over time. They derive mathematical consequences of certain observations about a particular instant in time to the rest of the state description at that same instant.

This gives us a complete initial state for simulation. The qualitative state description shows us that several of the parameters are changing. We can predict the evolution of the U-tube system over time by looking at the nature of the possible qualitative changes.

4.2. *Predicting the next state*

Each parameter changes continuously with time, so it is relatively easy to predict the successor(s) to a given state. Since at the initial instant t_0, amtB = 0 and increasing, then in the next qualitative state, amtB > 0 and still increasing. Similarly with pressureB. The same approach handles amtA = $\langle$AMAX, dec$\rangle$. However, pressureA is within an open interval, so it would require a finite amount of time to reach the boundary of its interval, while amtA or amtB can move off a landmark value instantaneously. Thus we get the next state description, specifying the qualitative values the parameters must have over some open time-interval following t_0. The subsequent simulation will show what event determines the time-point t_1 that terminates this interval.

$t \in (t_0, t_1) \Rightarrow$

amtA	$= \langle (0, \text{AMAX}), \text{dec} \rangle$
pressureA	$= \langle (0, \infty), \text{dec} \rangle$
amtB	$= \langle (0, \text{BMAX}), \text{inc} \rangle$
pressureB	$= \langle (0, \infty), \text{inc} \rangle$
pAB	$= \langle (0, \infty), \text{dec} \rangle$
flowAB	$= \langle (0, \infty), \text{dec} \rangle$.

Before we continue, there are a few technicalities we need to observe.

- We are actually presuming that not only the parameters, but their derivatives, vary continuously with time. That is, the parameters of the system are *continuously differentiable* functions of time.
- Although we speak of the "next state", since the underlying process we are describing is continuous, strictly speaking there *is* no next state. However, the sequence of distinct qualitative state descriptions is a discrete sequence, so the "next state" of a mechanism refers to the next distinct qualitative description in that sequence.

Time is modeled as an alternating sequence of time-points and open time-intervals. Our first state description represented the qualitative state of the U-tube at the initial instant t_0. The change from that point in time puts us into an open interval on the time-line, during which the qualitative description remains fixed, until we reach another distinguished time-point when a qualitative change takes place. Thus, the second state description applies, not to a single instant in time, but to every time t in the interval (t_0, t_1). Since several parameters are changing, the mechanism is changing during that interval, and the description reflects that fact. However, since it has not crossed any qualitative boundary, the *qualitative description* of the changing system remains constant during (t_0, t_1).

4.3. *Moving to a limit*

The previous qualitative state description applied to all time-points in a time-interval, $t \in (t_0, t_1)$. Now we need to determine the qualitative change that defines the time $t = t_1$ that terminates this interval. There are several kinds of qualitative changes:

- a parameter that is moving toward a limit may reach it;
- a parameter that is equal to a landmark value may move off it;
- a parameter may change its direction of change.

In this example, all six parameters are moving toward various limits, for a theoretical maximum of $2^6 = 64$ possible combinations of qualitative changes. The corresponding values on the addition and monotonic function constraints, however, greatly reduce this set. In the end, we are left with only a single unresolved ambiguity: the race between amtB $\rightarrow$ BMAX and flowAB $\rightarrow$ 0. There are no constraints to resolve this ambiguity, so our prediction must branch nondeterministically. (Physical intuition confirms that these possibilities are reasonable, since we do not know the relative sizes of tanks A and B.)

A non-deterministic prediction simply means that the qualitative description of the system does not contain enough information to specify its future state uniquely. The three possible states at t_1 are given below.

- amtB and flowAB reach their limits together, so the U-tube reaches equilibrium with tank B brimming full.
- flowAB = 0 while amtB < BMAX, so we reach an equilibrium state with tank B partially filled. (Figure 5(a) shows this behavior.)
- amtB reaches BMAX while flowAB > 0, so tank B overflows. (The current model of the U-tube ceases to apply, so we must make a transition to a new model. Figure 5(b) shows a transition to a model in which tank B has burst, and all the water in the system drains out.)

This non-deterministic prediction is represented as a branching tree of qualitative states:

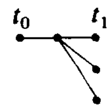

(a) tank B stops at top
(b) partially full
(c) overflow.

This type of non-deterministic prediction is an important feature of qualitative simulation. A specific U-tube is a deterministic mechanism, and cannot choose arbitrarily which behavior to take. However, our qualitative description is incomplete, and does not contain enough information to determine which behavior will be the real one. There are different physical U-tubes, all of which satisfy the same qualitative constraints and initial state, which exhibit each of the three possible behaviors.

The set of possible qualitative states at t_1 is:

t	$t_{1(a)}$	$t_{1(b)}$	$t_{1(c)}$
amtA	$\langle(0, \text{AMAX}), \text{std}\rangle$	$\langle(0, \text{AMAX}), \text{std}\rangle$	$\langle(0, \text{AMAX}), \text{dec}\rangle$
pressureA	$\langle(0, \infty), \text{std}\rangle$	$\langle(0, \infty), \text{std}\rangle$	$\langle(0, \infty), \text{dec}\rangle$
amtB	$\langle\text{BMAX}, \text{std}\rangle$	$\langle(0, \text{BMAX}), \text{std}\rangle$	$\langle\text{BMAX}, \text{inc}\rangle$
pressureB	$\langle(0, \infty), \text{std}\rangle$	$\langle(0, \infty), \text{std}\rangle$	$\langle(0, \infty), \text{inc}\rangle$
*p*AB	$\langle 0, \text{std}\rangle$	$\langle 0, \text{std}\rangle$	$\langle(0, \infty), \text{dec}\rangle$
flowAB	$\langle 0, \text{std}\rangle$	$\langle 0, \text{std}\rangle$	$\langle(0, \infty), \text{dec}\rangle$

One might hope that qualitative simulation would give us exactly the set of possible behaviors consistent with the partial knowledge captured by the qualitative description. In other words, every behavior of any real mechanism would be predicted, and every predicted behavior would be real for some mechanism satisfying the constraints. As we have seen, this is true for the U-tube. However, the first half of this can be guaranteed in general, but not the second half.

4.4. *Creating new landmark values*

In the second case ($t_{1(b)}$) described above, the U-tube system reaches equilibrium while amtB is still in the open interval (0, BMAX), so its qualitative state might be described as

$$\text{amtB}(t_{1(b)}) = \langle(0, \text{BMAX}), \text{std}\rangle .$$

The value of amtB when $t = t_{1(b)}$ is a critical value of the function amtB(t); its value when its derivative becomes zero. The critical value may be sufficiently important that it should be given a name so it can be referred to later. It may represent an important qualitative distinction that will be useful in other contexts. This is exactly what we mean by a "landmark value".

Since amtB($t_{1(b)}$) lies strictly between two existing landmarks, we can give it a name, e.g. B^*, and insert it as a new landmark into the quantity space for amtB

$$\text{amtB:}\quad 0 \cdots B^* \cdots \text{BMAX} \cdots \infty .$$

Following this strategy for the other parameters that reach critical values between landmarks, we augment the quantity spaces for the U-tube system

amtA	$0 \cdots A^* \cdots \text{AMAX} \cdots \infty$
pressureA	$0 \cdots PA^* \cdots \infty$
amtB	$0 \cdots B^* \cdots \text{BMAX} \cdots \infty$
pressureB	$0 \cdots PB^* \cdots \infty$
*p*AB	$-\infty \cdots 0 \cdots \infty$
flowAB	$-\infty \cdots 0 \cdots \infty$.

Several constraints can now define new corresponding values from the new equilibrium state, representing the meaning of these new landmarks through their relationships with the

Structure: U-tube,
Initialization: Tank A full; B empty (S-2735)
Behavior 2 of 3: (S-2735 S-2736 S-2738).
Final state: (QUIESCENT), NIL, NIL.

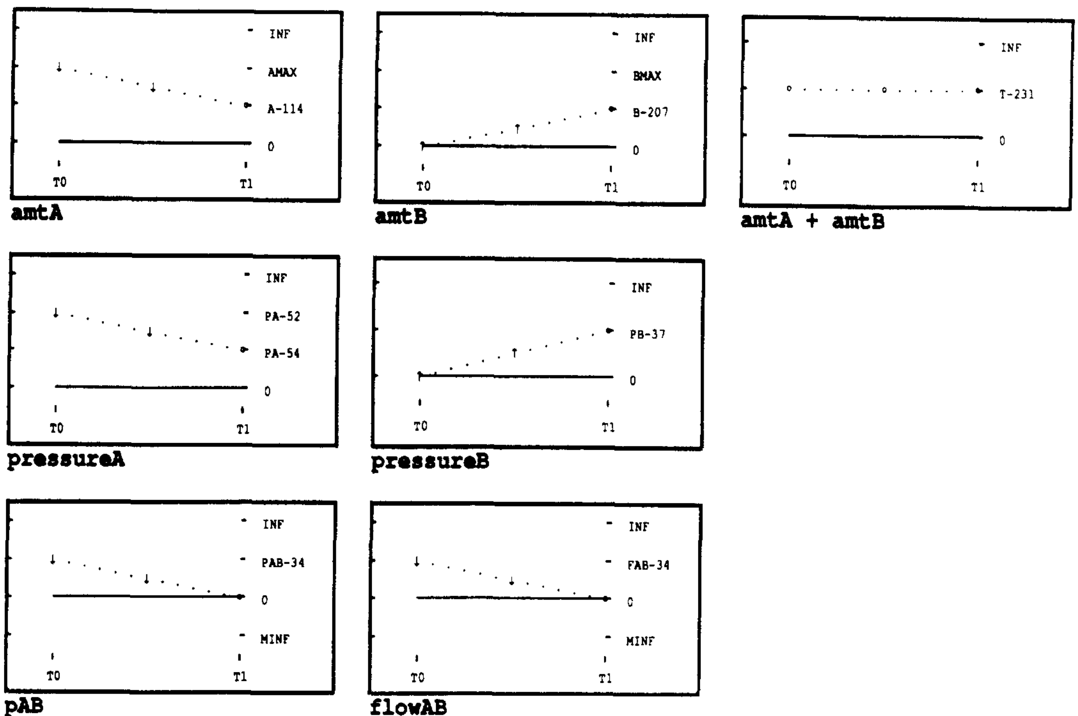

(a)

Structure: U-tube,
Initialization: Tank A full; B empty (S-2735)
Behavior 3 of 3: (S-2735 S-2736 S-2739 S-2740 S-2741 S-2742 S-2744 S-2745).
Final state: (QUIESCENT), NIL, NIL.

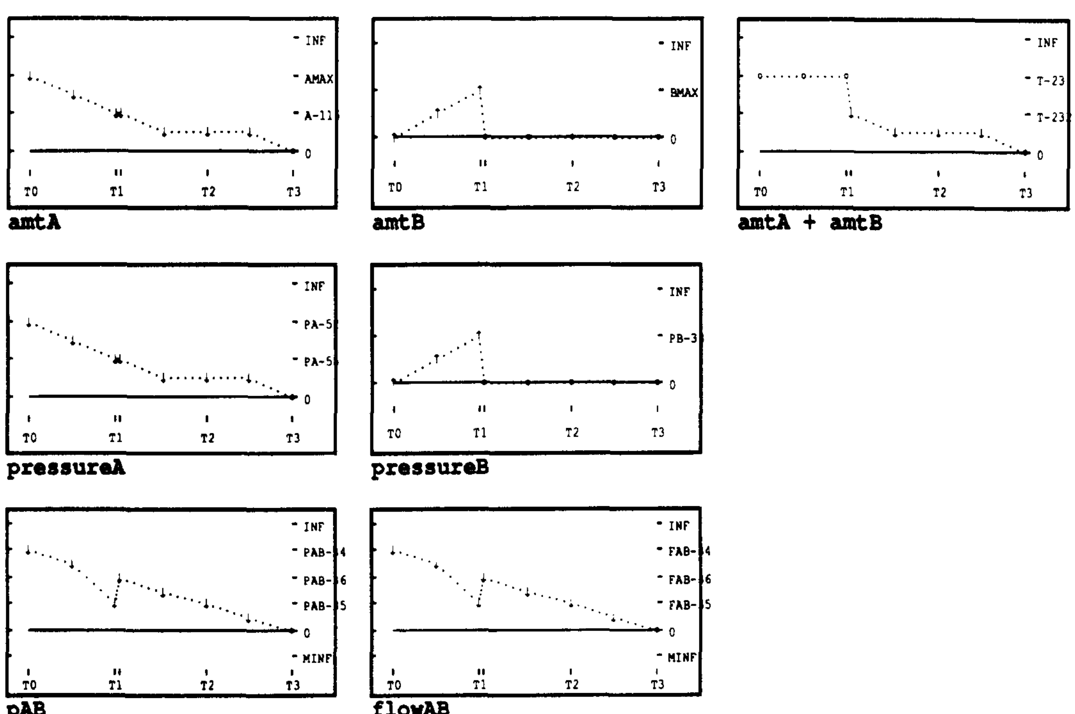

(b)

FIG. 5. Two qualitative behaviors of the U-tube: (a) equilibrium with tank B partially full; (b) tank B overflows and bursts, losing its contents immediately and allowing the contents of tank A to drain away as well.

constraints:

constraints	*new corresponding values*
pressureA = M^+(amtA)	(A^*, PA^*)
pressureB = M^+(amtB)	(B^*, PB^*)
pressureA − pressureB = pAB	$(PA^*, PB^*, 0)$
flowAB = M^+(pAB)	
flowAB = $\frac{d}{dt}$ amtB	
flowAB = $-\frac{d}{dt}$ amtB	

The final equilibrium state of the system along this branch of the tree can now be described more precisely in terms of the new landmarks.

$t = t_{1(b)} \Rightarrow$

amtA	$= \langle A^*, \text{std} \rangle$
pressureA	$= \langle PA^*, \text{std} \rangle$
amtB	$= \langle B^*, \text{std} \rangle$
pressureB	$= \langle PB^*, \text{std} \rangle$
pAB	$= \langle 0, \text{std} \rangle$
flowAB	$= \langle 0, \text{std} \rangle$

New landmarks and corresponding values are only meaningful on their own branch of the state tree. For example, the landmark A^*, representing the amount of water in tank A when the U-tube reaches equilibrium, has no significance along the branch in which tank B overflows. Along that branch, there would be some other new landmark for amtA representing its value of that parameter when the overflow occurred.

The three paths through the tree of qualitative states represent the three possible behaviors of the system. Each qualitative graph represents a single behavior, so we need three qualitative graphs to capture these possibilities. Figure 5 shows the qualitative behavior on two branches of the tree. Figure 5(a) shows the behavior where equilibrium is achieved before tank B becomes full. Figure 5(b) shows qualitative behavior where tank B overflows; in this scenario, overflow results in tank B bursting, losing its own contents immediately and allowing the contents of tank A to drain away as well, emptying the system.

4.5. *The QSIM representation and algorithm*

The representation and inference process we have discussed here has been embodied in an efficient constraint-filtering algorithm called QSIM (Kuipers, 1985, 1986), which generates the tree of possible behaviors following from a given initial state and a QDE.

QSIM has been applied to a large collection of mechanisms from simple physical systems to electronic circuits to physiological systems. It is the foundation for ongoing research on qualitative reasoning methods by our group at the University of Texas at Austin.

5. THE QUASI-EQUILIBRIUM ASSUMPTION

In our previous example, we assumed that the U-tube began in a state where amtA(t_0) = AMAX and amtB(t_0) = 0. But how could it ever get into such a state? In commonsense terms, if the channel between the tanks is small, and I quickly pour a bucket of water into tank A, that faster process is complete before the U-tube can respond significantly. From the U-tube's point of view, it suddenly finds itself in a non-equilibrium state and moves toward equilibrium again. Thus, we can treat a non-equilibrium initial state as the result of a process acting at a much faster time-scale than the one we are considering.

What if the U-tube is driven by a much slower process? Then we can make the *quasi-equilibrium assumption*: that the system is always in, or infinitely close to, equilibrium. Consider the resistor, which moves from one equilibrium state to another at electronic speeds. Since the process that changes a voltage is likely to be much slower than the response of the resistor to a perturbation of its equilibrium, we can assume that it is *always* in equilibrium, and do inference with Ohm's law, which describes the equilibrium states. Frequently, but not always, solving the equations can be done by local propagation through the constraints. The effect of changes to such a system can be deduced simply by solving the equilibrium equations for the new state. This is essentially the inference method used by de Kleer and Brown (1984). Puccia and Levins (1985) and Iwasaki and Simon (1986) also discuss methods for applying quasi-equilibrium reasoning.

In the U-tube, where we have a non-equilibrium model, we can express the quasi-equilibrium assumption by asserting that flowAB = 0 and by deleting the derivative constraints (Fig. 6).

If we now assert that amtA is positive, we can propagate that fact through the remaining constraints to deduce (using de Kleer and Brown's notation in which $[X]_0 = \text{sign}(X)$)

$$[\text{flowAB}]_0 = 0 \Rightarrow [p\text{AB}]_0 = 0$$
$$[\text{amtA}]_0 = + \Rightarrow [\text{pressureA}]_0 = + \Rightarrow [\text{pressureB}]_0 = + \Rightarrow [\text{amtB}]_0 = +.$$

We can carry out a similar inference for perturbations around an equilibrium state,

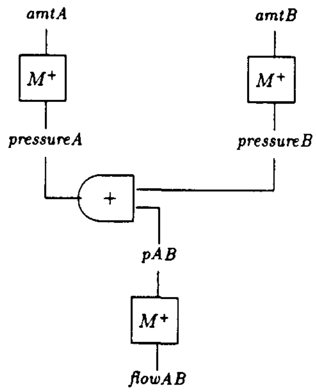

FIG. 6. The U-tube under the quasi-equilibrium assumption.

where $[X]_*$ gives the sign of X relative to the equilibrium value X_*

$$[\text{flowAB}]_* = 0 \Rightarrow [p\text{AB}]_* = 0$$
$$[\text{amtA}]_* = + \Rightarrow [\text{pressureA}]_* = +$$
$$\Rightarrow [\text{pressureB}]_* = +$$
$$\Rightarrow [\text{amtB}]_* = +.$$

We can also deduce the results of changing values within the quasi-equilibrium assumption, as long as the changes are slow with respect to the intrinsic time-scale of the mechanism. If amtA increases slowly, amtB will increase slowly right along with it

$$[\text{flowAB}]_* = \langle 0, \ominus \rangle \Rightarrow [p\text{AB}]_* = \langle 0, \ominus \rangle$$
$$[\text{amtA}]_* = \langle +, \uparrow \rangle \Rightarrow [\text{pressureA}]_* = \langle +, \uparrow \rangle$$
$$\Rightarrow [\text{pressureB}]_* = \langle +, \uparrow \rangle$$
$$\Rightarrow [\text{amtB}]_* = \langle +, \uparrow \rangle.$$

Limit analysis still applies, so although flowAB = 0 and constant, amtA → AMAX and amtB → BMAX. As before, there are three possible next states, depending on the order in which these parameters reach their limits.

In domains where the quasi-equilibrium assumption is realistic, this kind of near-equilibrium reasoning is computationally efficient and avoids some of the sources of ambiguity encountered in the transient states of a dynamic simulation.

6. BUILDING MODELS WITH QUALITATIVE PROCESS THEORY

We have been discussing methods for qualitative simulation of a model once it has been created. However, a prior question is, "How is the model created?" That is, how are the elements of a model identified, and how are they assembled to create a meaningful and useful model?

Qualitative process theory (Forbus, 1984) provides us with a promising initial approach to this very difficult problem. There are two fundamental types of description in QP theory: *Individual Views*, which represent objects or sets of objects viewed in a particular way, and *Processes*, which represent active changes taking place. All change is considered to emanate from processes. Table 1 gives examples of the individual view Contained-liquid and the process Fluid-flow involved in building the U-tube model.

TABLE 1. THE CONTAINED-LIQUID VIEW AND THE FLUID-FLOW PROCESS

Individual-view Contained-liquid
Individuals:
c: a container
s: a piece-of-liquid
Preconditions:
in(s, c)
Quantity-conditions:
amt(s) > 0
Relations:
pressure(s) $\propto_{Q^+}$ amt(s)
Process Fluid-flow(X, Y)
Individuals:
X: a contained-liquid
Y: a contained-liquid
Preconditions:
flow-connected(X, Y)
Quantity-conditions:
pressure(X) > pressure(Y)
Relations:
flow(X, Y) $\propto_{Q^+}$ pressure(X)
flow(X, Y) $\propto_{Q^-}$ pressure(Y)
Influences:
I^+(flow(X, Y), amt(Y))
I^-(flow(X, Y), amt(X))

The model building process can be summarized as follows.

(1) Each individual view and process checks various conditions about the world to determine whether it should have one or more active *instances*.

(2) The active view and process instances are grouped into sets of mutually consistent elements, called *View-process structures*.

(3) Each view or process in a view-process structure contributes fragments, called direct or indirect *influences*, to the constraint model.

(4) The Closed World Assumption is applied within each view-process structure. This makes it possible to determine the set of all influences, direct or indirect, applying to any given quantity, so the influences can be translated to constraints.

(5) The resulting constraint model can then be simulated as discussed above.

The U-tube example might be physically described as follows:

container(A)	container(B)
piece-of-water(WA)	piece-of-water(WB)
in(WA, A)	in(WB, B)
amt(WA) > 0	amt(WB) > 0.

The description of water in a container as a "piece of water" is taken from Hayes' (1985) ontology for liquids. This allows the water in a container to be treated as a single object of varying size, even as water flows in or out of the container, even if every molecule of water is actually replaced. Such an ontology is assumed by any macroscopic mass-balance model, but must be carefully analyzed and made explicit if such models are to be constructed automatically.

Table 1 shows (simplified) descriptions of the individual view for Contained-liquid and the process for Fluid-flow. Given the above information, we can activate two instances of Contained-liquid, one for each container, and two potential instances of Fluid-flow: flow(A, B) and flow(B, A). The two process instances are not compatible, so there are three possible view-process structures, one for each Fluid-flow instance, and one where neither flow is active.

Let us consider the view-process structure that corresponds to the scenario we have looked at above

Contained-liquid(A)
Contained-liquid(B)
Fluid-flow(A, B).

This contributes a set of influences to the model

$$\text{pressure(A)} \propto_{Q^+} \text{amt(A)}$$
$$\text{pressure(B)} \propto_{Q^+} \text{amt(B)}$$
$$\text{flowAB} \propto_{Q^+} \text{pressure(A)}$$
$$\text{flowAB} \propto_{Q^-} \text{pressure(B)}$$
$$I^+(\text{flowAB}, \text{amt(B)})$$
$$I^-(\text{flowAB}, \text{amt(A)}).$$

An influence is not the same as a constraint. An indirect influence $Y \propto_{Q^+} X$ asserts that for a given direction of change of X, Y will change in the same direction, *all else being equal*. It is only with a closed world assumption, after the entire view-process structure has been identified, that the program can determine whether all else is equal or not.

• Indirect influences (e.g. $\propto_{Q^+}$) are translated into monotonic function constraints. Since only amt(A) influences pressure(A), the direct influence pressure(A) $\propto_{Q^+}$ amt(A) can be translated directly into the monotonic function constraint

$$\text{pressure(A)} = M^+(\text{amt(A)}).$$

• Since flowAB is positively influenced by pressure(A) and negatively influenced by pressure(B), the two influences must be combined (by default, additively) before being translated to constraints

$$\text{flowAB} = M^+(\text{pressure(A)}) - M^+(\text{pressure(B)}).$$

• Direct influences are translated into derivative constraints, with the same handling of multiple influences. Here, however, the two directly influenced quantities have only one influence each (from the same quantity)

$$\frac{d}{dt}\text{amt(A)} = -\text{flowAB}$$

$$\frac{d}{dt}\text{amt(B)} = +\text{flowAB}.$$

The result of this is a set of constraints almost identical to the U-tube model in Fig. 3

$$\text{pressure(A)} = M^+(\text{amt(A)})$$
$$\text{pressure(B)} = M^+(\text{amt(B)})$$
$$\text{flowAB} = M^+(\text{pressure(A)}) - M^+(\text{pressure(B)})$$
$$\frac{d}{dt}\text{amt(A)} = -\text{flowAB}$$
$$\frac{d}{dt}\text{amt(B)} = +\text{flowAB}.$$

Since the Fluid-flow(A, B) process asserts that pressure(A) > pressure(B), this system is changing dynamically. Qualitative simulation shows that it approaches a state where pressure(A) = pressure(B). Examination of the view-process structure consistent with that state shows that no processes are active, so the system is quiescent.

There remain open problems hidden in this sketch of the interface between QP theory as a model-builder, and qualitative simulation as a model-simulator. Falkenhainer and Forbus (1988) describe recent progress on the use of Qualitative Process Theory for automatic creation of qualitative models, particularly in manipulating explicit modeling assumptions.

7. THE MATHEMATICAL VALIDITY OF PREDICTED BEHAVIORS

The qualitative modeling methods we have been discussing are intended to express partial knowledge of the kind of systems that would traditionally be modeled by differential equations. In order to have confidence in our methods, we need to determine the validity of the predictions made by QSIM or other qualitative reasoning systems. The results discussed here are presented in detail in Kuipers (1986).

A set of qualitative constraints may be regarded as a *qualitative differential equation* (or QDE). A given QDE may be an abstraction of one or more ordinary differential equations. In particular, where a QDE may assert that $y = M^+(x)$, corresponding ODEs would state that $y = f(x)$ for various specific monotonic

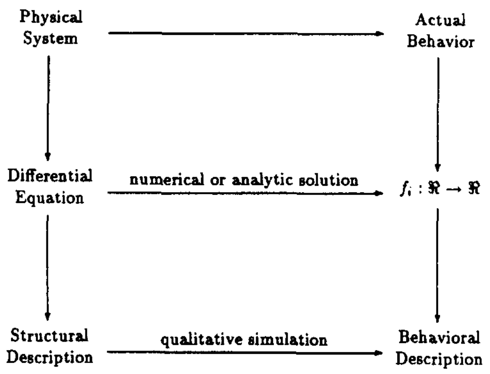

FIG. 7. Qualitative simulation and differential equations are both abstractions of actual behavior.

functions f. Thus, we can identify an abstraction relation (Fig. 7) from the physical world itself, to an ODE description of the world, to a more abstract QDE description of the same situation.

Note that in practice, QDEs are not constructed by abstracting previously-known ODEs. Rather, in modeling a physical situation, it may be that not enough is known to construct an ODE that is both valid and tractable, so the only useful model that can be created is a QDE model. Nonetheless, we presume that, known or not, and tractable or not, there *exists* a valid ODE model of such a situation, and the predictions of that ODE are the gold standard against which we evaluate our quantitative simulation algorithm, QSIM.

QSIM takes as input a QDE and a description of its state at time t_0. It then predicts the possible behaviors of the QDE as a (possibly branching) tree of states. A behavior is a sequence of qualitative descriptions of states

Behavior
$= [\text{state}(t_0), \text{state}(t_0, t_1), \text{state}(t_1), \ldots, \text{state}(t_n)]$.

QSIM predicts a set of possible behaviors, which is interpreted as a disjunction

$$\text{QSIM: QDE, state}(t_0) \rightarrow \text{or}(B_1, \ldots, B_k).$$

That is, starting in state(t_0), QSIM predicts that one of the behaviors $B_1, \ldots, B_k$ will describe the actual behavior of the system.

This inference is sound, but incomplete.

- *Soundness.* If the QDE is an abstraction of a certain ODE, then every solution to the ODE corresponds to some B_i in the QSIM prediction. That is, if the QDE is valid, the disjunctive prediction is valid.
- *Incompleteness.* If QSIM predicts the disjunction or($B_1, \ldots, B_k$) from a given QDE, it is possible for some disjunct, say B_k, to fail to correspond to any solution of any ODE that abstracts to the QDE. That is, QSIM fails to prove the stronger disjunct, or($B_1, \ldots, B_{k-1}$).
- *Corollary.* If the QDE is valid, and QSIM predicts a single behavior, or(B_1), then B_1 is a valid description of the behavior of the system.

Soundness is proved because the QSIM algorithm implicitly defines an enormously large product space guaranteed to include all valid behaviors, and filters out impossible behaviors with probably correct filters. Incompleteness is proved by exhibiting a QDE such that the solutions to all ODEs abstracting to that QDE correspond to only one of several predicted behaviors. The others are unnecessary disjuncts (sometimes called *spurious predictions*).

The reason for the spurious predictions that have been identified thus far is the local view of the qualitative behavior of a system taken by qualitative state descriptions and the limit-analysis approach to prediction.

8. SCALING UP TO REALISTIC SYSTEMS

There are three key steps that are required to scale this technology up to handle simulation of realistic systems. A very active community of researchers around the world, including our group at the University of Texas at Austin, are working to take these steps.

(1) More powerful mathematical methods.
(2) Hierarchical decomposition of complex models.
(3) Incremental application of quantitative knowledge.

8.1. *More mathematical power*

More powerful mathematical methods make it possible to select more appropriate levels of qualitative description, and to filter out additional inconsistent predictions. Directions of recent work include the following.

- Methods for changing the level of description to aggregate large sets of behaviors the distinctions of which are real, but not considered important, without sacrificing the ability to make valid qualitative predictions (Kuipers and Chiu, 1987).
- Methods for reasoning with higher-order derivatives, including the use of symbolic algebra on QDEs, along with the required assumptions of smoothness of certain relations (de Kleer and Bobrow, 1984; Kuipers and Chiu, 1987).
- Methods for comparing the energy and phase properties of an oscillatory system at successive extreme points (Lee *et al.*, 1987).
- More general application of geometric phase space concepts and Lyapunov (generalized "energy") functions to a larger class of

second-order QDEs (Lee and Kuipers, 1988; Struss, 1988).

• Methods for aggregating and abstracting the behavior of a repetitive process (e.g. an oscillatory process consuming a finite resource) to reason about its limit (Weld, 1986).

• Methods for qualitative perturbation analysis showing how a predicted behavior changes given some change to the QDE or the initial conditions (Weld, 1987).

• Methods for emulating the qualitative reasoning methods of mathematicians and engineers investigating dynamical systems (Sacks, 1985, 1987a,b).

8.2. *Hierarchical decomposition*

Realistic systems are typically much larger than second-order, making them normally intractable to qualitative or analytic analysis. The best methods for coping with such complexity impose a hierarchical structure to decompose a large, complex system into a collection of simpler systems and their interactions.

We have developed a method called *time-scale abstraction* for decomposing a complex process operating on a widely separated time-scale into several simple processes, each at its own time-scale (Kuipers, 1987, 1988). A given process in a time-scale hierarchy can view a slower process as constant (or as driving a change under the quasi-equilibrium assumption), and a faster one as essentially instantaneous. The key step is to view the result of the behavior of a fast mechanism as implementing a monotonic functional constraint (M^+ or M^-) from the point of view of the slower mechanism.

This decomposition method has the potential for allowing us to give qualitative predictions of the behavior of complex system models by decomposing by time-scale, and transforming a complex model into a hierarchy of simpler models. The more powerful the mathematical methods we can apply to the component models, the less decomposition is required in the hierarchy. An effective method for deriving the time-scale abstraction hierarchy from a linear model of a nearly-decomposable system has recently been developed by Iwasaki and Bhandari (1988), building on work by Simon and Ando (1961).

8.3. *Incremental quantitative constraints*

Qualitative reasoning methods are applied to purely qualitative descriptions of the world, in which a value is described purely in terms of its ordinal relations with other values. This has been a surprisingly fruitful research strategy: ordinal relations are a major part of human commonsense knowledge about quantities, and it is remarkable how many useful conclusions can be drawn from such an apparently weak description. However, it is clear that humans normally do have some useful knowledge about quantitative magnitudes, though seldom actual numerical values, and that this knowledge plays a role in qualitative predictions.

One approach is to include qualitative *order of magnitude* relations among values. This approach has been fruitfully explored by Raiman (1986), Dague *et al.* (1987), and Mavrovouniotis and Stephanopoulos (1987).

We have developed methods for incrementally adding quantitative constraints to qualitative predictions (Kuipers and Berleant, 1988). A QDE includes quantity spaces, the landmarks of which are essentially *names* for values on the real number line. We can assert quantitative information, if it is known, about the numerical value of each of those landmarks, about the lengths of intervals between landmarks, and about the functional constraints linking parameters.

There are many different representations for incompletely known quantities, including intervals specified by upper and lower bounds, and distributions specified by mean and variance. Known information is propagated along constraint to provide a more complete quantitative description of the elements of the qualitative behavior. A particularly desirable outcome is for this propagation process to yield a contradiction, demonstrating that a particular branch of the behavior tree is incompatible with quantitative knowledge and can be eliminated.

Figure 8 shows the result of considering quantitative ranges for the landmarks AMAX and BMAX which imply that tank B is strictly larger than tank A. There are also piecewise linear envelopes that constrain the behavior of the monotonic function (M^+) constraints. Two of the three qualitatively possible behaviors are inconsistent with the quantitative information provided, and the remaining behavior is annotated with quantitative range information.

9. CONCLUSION

This paper provides a tutorial overview of a collection of qualitative reasoning methods that may fill an important gap in the modeling and control toolkit. Qualitative reasoning methods provide greater expressive power for states of incomplete knowledge than differential or difference equations, and thus make it possible to build models without incorporating assumptions of linearity or specific values for incompletely known constants. Even with incomplete

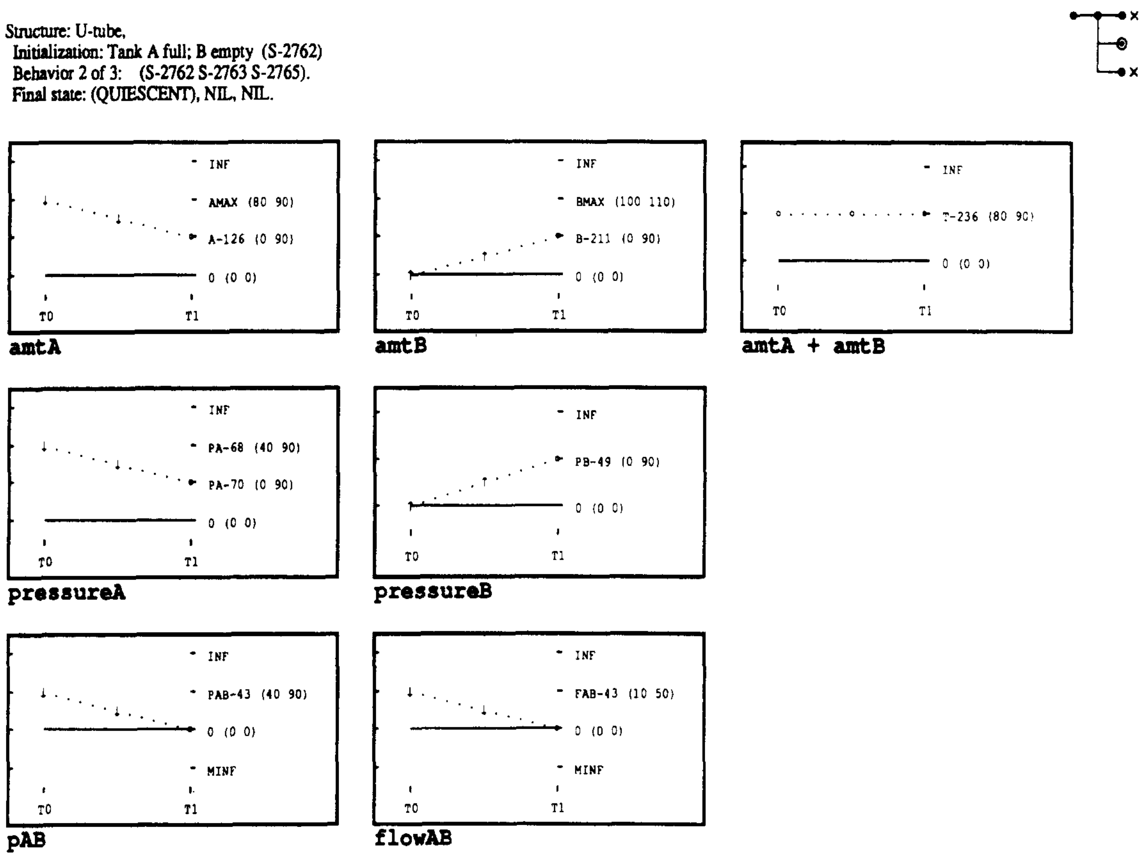

FIG. 8. Only one out of three qualitatively possible behaviors of the U-tube is consistent with the given incomplete quantitative information.

knowledge, there is enough information in a qualitative description to make meaningful prediction feasible.

Current results with small examples have been encouraging, and steps are now being taken toward additional mathematical power, hierarchical decomposition, and incremental quantitative constraints. Some applications to realistic systems have begun to appear in the literature. Dague *et al.* (1987) have applied order-of-magnitude reasoning to a model-based troubleshooting expert system for complex analog circuits. Dalle Molle *et al.* (1988b) have developed QSIM models for the two-input, two-output mixing tank, with and without a proportional controller. Dalle Molle *et al.* (1988a) present a QSIM model of a radial-flow plasma etcher.

We believe that qualitative model and simulation are becoming valuable formal reasoning methods.

Acknowledgments—The research reported here was supported in part by the National Science Foundation through grants DCR-8512779 and IRI-8602665, by the National Library of Medicine through NIH grants LM 04374 and LM 04515, and by the National Aeronautics and Space Administration through grant NAG 2-507.

REFERENCES

Dague, P., O. Raiman and P. Deves (1987). Troubleshooting: when modeling is the trouble. *Proc. Nat. Conf. on Artificial Intelligence* (*AAAI-87*). Morgan Kaufmann, Los Altos, California.

Dalle Molle, D. T., T. F. Edgar and I. Trachtenberg (1988a). Qualitative modeling of a plasma etcher. Papers of the AAAI Workshop on AI in Process Engineering, Minneapolis, Minnesota, 25 August.

Dalle Molle, D. T., B. J. Kuipers and T. F. Edgar (1988b). Qualitative modeling and simulation of dynamic systems. *Comput. Chem. Engng*, **12**, 853–866.

de Kleer, J. (1977). Multiple representations of knowledge in a mechanics problem-solver. In *Proc. Fifth Int. Joint Conf. on Artificial Intelligence* (*IJCAI-77*). Morgan Kaufmann, Los Altos, California.

de Kleer, J. and D. G. Bobrow (1984). Qualitative reasoning with higher-order derivatives. *Proc. Nat. Conf. on Artificial Intelligence* (*AAAI-84*). Morgan Kaufmann, Los Altos, California.

de Kleer, J. and J. S. Brown (1984). A qualitative physics based on confluences. *Artif. Intell.*, **24**, 7–83.

Falkenhainer, B. and K. D. Forbus (1988). Setting up large-scale qualitative models. *Proc. Nat. Conf. on Artificial Intelligence* (*AAAI-88*). Morgan Kaufmann, Los Altos, California.

Forbus, K. D. (1984). Qualitative process theory. *Artif. Intell.*, **24**, 85–168.

Hayes, P. J. (1985). Naive physics I: ontology for liquids. In J. R. Hobbs and R. C. Moore (Eds), *Formal Theories of the Commonsense World*. Ablex, Norwood, New Jersey.

Iwasaki, Y. and I. Bhandari (1988). Formal basis for commonsense abstraction of dynamic systems. *Proc. Nat. Conf. on Artificial Intelligence* (*AAAI-88*). Morgan Kaufmann, Los Altos, California.

Iwasaki, Y. and H. A. Simon (1986). Causality in device behavior. *Artif. Intell.*, **29**, 3–32.
Kuipers, B. J. (1984). Commonsense reasoning about causality: deriving behavior from structure. *Artif. Intell.*, **24**, 169–204.
Kuipers, B. J. (1985). The limits of qualitative simulation. *Proc. Nineth Int. Joint Conf. on Artificial Intelligence* (*IJCAI-85*). Morgan Kaufmann, Los Altos, California.
Kuipers, B. J. (1986). Qualitative simulation. *Artif. Intell.*, **29**, 289–338.
Kuipers, B. (1987). Abstraction by time-scale in qualitative simulation. *Proc. Nat. Conf. on Artificial Intelligence* (*AAAI-87*). Morgan Kaufmann, Los Altos, California.
Kuipers, B. (1988). Abstraction by time-scale in qualitative simulation for biomedical modeling. In C. Cobelli and L. Mariani (Eds), *Modelling and Control in Biomedical Systems*, Preprints of the IFAC Symposium, Venice, Italy, 6–8 April 1988. Pergamon Press, Oxford.
Kuipers, B. and D. Berleant (1988). Using incomplete quantitative knowledge in qualitative reasoning. *Proc. Nat. Conf. on Artificial Intelligence* (*AAAI-88*). Morgan Kaufmann, Los Altos, California.
Kuipers, B. J. and J. P. Kassirer (1984). Causal reasoning in medicine: analysis of a protocol. *Cognitive Sci.*, **8**, 363–385.
Kuipers, B. and C. Chiu (1987). Taming intractable branching in qualitative simulation. *Proc. Tenth Int. Joint Conf. on Artificial Intelligence* (*IJCAI-87*). Morgan Kaufmann, Los Altos, California.
Lee, W. W., C. Chiu and B. J. Kuipers (1987). Developments towards constraining qualitative simulation. University of Texas at Austin, Artificial Intelligence Laboratory AI TR 87-44, January.
Lee, W. W. and B. Kuipers (1988). Non-intersection of trajectories in qualitative phase space: a global constraint for qualitative simulation. *Proc. Nat. Conf. on Artificial Intelligence* (*AAAI-88*). Morgan Kaufmann, Los Altos, California.
Mavrovouniotis, M. L. and G. Stephanopoulos (1987). Reasoning with orders of magnitude and approximate relations. *Proc. Nat. Conf. on Artificial Intelligence* (*AAAI-87*). Morgan Kaufmann, Los Altos, California.
Patil, R. S., P. Szolovits and W. B. Schwartz (1981). Causal understanding of patient illness in medical diagnosis. *Proc. Seventh Int. Joint Conf. on Artificial Intelligence* (*IJCAI-81*).
Pople, H. E., Jr. (1982). Heuristic method for imposing structure on ill structured problems: the structuring of medical diagnostics. In P. Szolovits (Ed.), *Artificial Intelligence in Medicine*. AAAS/Westview Press, Boulder, Colorado.
Puccia, C. J. and R. Levins (1985). *Qualitative Modeling of Complex Systems*. Harvard University Press, Cambridge, Massachusetts.
Raiman, O. (1986). Order of magnitude reasoning. *Proc. Nat. Conf. on Artificial Intelligence* (*AAAI-86*). Morgan Kaufmann, Los Altos, California.
Sacks, E. (1985). Qualitative mathematical reasoning. *Proc. Nineth Int. Joint Conf. on Artificial Intelligence* (*IJCAI-85*). Morgan Kaufmann, Los Altos, California.
Sacks, E. (1987a). Hierarchical reasoning about inequalities. *Proc. Nat. Conf. on Artificial Intelligence* (*AAAI-87*). Morgan Kaufmann, Los Altos, California.
Sacks, E. (1987b). Piecewise linear reasoning. *Proc. Nat. Conf. on Artificial Intelligence* (*AAAI-87*). Morgan Kaufmann, Los Altos, California.
Shortliffe, E. H. (1976). *Computer-based Medical Consultations*: *MYCIN*. Elsevier, New York.
Simmons, R. and R. Davis (1987). Generate, test and debug: combining associational rules and causal models. *Proc. Tenth Int. Joint Conf. on Artificial Intelligence* (*IJCAI-87*). Morgan Kaufmann, Los Altos, California.
Simon, H. A. and A. Ando (1961). Aggregation of variables in dynamic systems. *Econometrica*, **29**, 111–138.
Struss, P. (1988). Global filters for qualitative behaviors. *Proc. Nat. Conf. on Artificial Intelligence* (*AAAI-88*). Morgan Kaufmann, Los Altos, California.
Weld, D. S. (1986). The use of aggregation in causal simulation. *Artif. Intell.*, **30**, 1–34.
Weld, D. S. (1987). Comparative analysis. *Proc. Tenth Int. Joint Conf. on Artificial Intelligence* (*IJCAI-87*). Morgan Kaufmann, Los Altos, California.
Williams, B. (1986). Doing time: putting qualitative reasoning on firmer ground. *Proc. Fifth Nat. Conf. on Artificial Intelligence* (*AAAI-86*), pp. 105–112. Morgan Kaufmann, Los Altos, California.