

Taming Intractable Branching in Qualitative Simulation

Benjamin Kuipers

Department of Computer Sciences
University of Texas at Austin
Austin, Texas 78712

Charles Chiu

Department of Physics and Artificial Intelligence Laboratory
University of Texas at Austin
Austin, Texas 78712


Abstract

Qualitative simulation of behavior from structure is a valuable method for reasoning about partially known physical systems. Unfortunately, in many realistic situations, a qualitative description of structure is consistent with an intractably large number of behavioral predictions. We present two complementary methods, representing different trade-offs between generality and power, for taming an important case of intractable branching. The first method applies to the most general case of the problem. It changes the level of the behavioral description to aggregate an exponentially exploding tree of behaviors into a few distinct possibilities. The second method draws on additional mathematical knowledge, and assumptions about the smoothness of partially known functional relationships, to derive a correspondingly stronger result. Higher-order derivative constraints are automatically derived by manipulating the structural constraint model algebraically, and applied to eliminate impossible branches. These methods have been implemented as extensions to QSIM and tested on a substantial number of examples. They move us significantly closer to the goal of reasoning qualitatively about complex physical systems.

1 Introduction

Qualitative simulation is a promising method for reasoning about the behavior of physical systems, starting from incomplete knowledge of the structure and initial state [Kuipers, 84, 86; de Kleer and Brown, 84; de Kleer and Bobrow, 84; Forbus, 84; Williams 84a, 86]. Incompletely known values may be described qualitatively in terms of their relations with a discrete set of *landmark values*. Incompletely known functional relations may be described qualitatively as monotonically increasing or decreasing, and passing through certain corresponding landmark values. Methods of qualitative simulation have demonstrated promising results on a variety of small and moderate-sized examples [Kuipers, 84, 85, 86, 87; Forbus 84, 86; de Kleer, 84; de Kleer and Brown, 84; Williams, 84a, 86].

In attempting to extend these techniques to simulate the continuous behavior of larger and more tightly interacting systems, however, certain problems have been encountered, resulting in a proliferation of predicted behaviors. Under the different representations for qualitative behavior, this proliferation is manifest in different ways.




- (a) A tree of behaviors for the cascaded tank system produced by QSIM.
- (b) A transition graph for damped spring system from [de Kleer and Brown, 1984].

*This research was supported in part by the National Science Foundation through grants MCS-8303640, DCR-8417934, and DCR-8512779.

Figure 1: Intractable branching in tree and transition graph representations.

- QSIM [Kuipers, 85, 86] produces a tree of possible behaviors, where each path down the tree is a sequence of qualitative states. Proliferation is manifest as an intractably branching behavior tree. (Figure 1a)
- Representations that do not create new landmark values [de Kleer and Brown, 84; de Kleer and Bobrow, 84; Forbus, 84; Williams, 84a, 86] are able to enumerate all possible qualitative states in advance, and thus produce a *total envisionment*, or transition graph on the qualitative states. Proliferation is exhibited in the transition graph as branches and loops, making an infinite set of possible paths through the graph. (Figure 1b)

An important class of proliferation problems arises with *coupled systems* such as two tank problems (Figure 2). (These problems represent general classes of important applications problems (e.g. [Sachs, et al., 1986], etc.).)



$inflowA = IF^*$	$inflowA = IF^*$
$pressureA = M^+(amountA)$	$pressureA = M^+(amountA)$
$pressureB = M^+(amountB)$	$pressureB = M^+(amountB)$
$outflowA = M^+(pressureA)$	$pressureDiff = pressureA - pressureB$
$outflowB = M^+(pressureB)$	$crossflowAB = M^+(pressureDiff)$
$netflowA = inflowA - outflowA$	$outflowB = M^+(pressureB)$
$netflowB = outflowA - outflowB$	$netflowA = inflowA - crossflowAB$
$\frac{d}{dt}amountA = netflowA$	$netflowB = crossflowAB - outflowB$
$\frac{d}{dt}amountB = netflowB$	$\frac{d}{dt}amountA = netflowA$
	$\frac{d}{dt}amountB = netflowB$

(a)
(b)

Figure 2: Two-tank systems and their constraint models.

- (a) The Cascaded Tank system.
- (b) The Coupled Tank system.

The underlying problem, illustrated here in terms of QSIM behaviors, is that when two distinct processes produce qualitative parameters that are changing in the same direction, their difference is unconstrained, except by continuity. Figure 3 shows one such behavior, representing a qualitative phenomenon we call “chatter”. Here the parameter *netflowB* chatters, while the other parameters move without changing direction. Figure 4a shows the transition graph representation of a single unconstrained parameter. For the examples shown here, the chattering behaviors consist of unconstrained wandering among the qualitative states $\langle +, inc \rangle$, $\langle +, std \rangle$, and $\langle +, dec \rangle$.¹ The problem is that, with no information about the actual shapes of functional relationships such as $outflowA = M^+(pressureA)$, all of the predicted behaviors are *real possibilities*. But an exponentially growing set of behaviors so obscures the actual qualitative properties of the system as to eliminate the value of qualitative simulation.

We have developed two distinct methods for solving this problem. One method applies to the general problem, and produces a slightly weaker qualitative description that collapses unimportantly distinct branches into a single history. The other method takes advantage of additional knowledge or assumptions about the system, and produces a correspondingly stronger result.





Figure 3: One “chattering” behavior of the Cascaded Tanks


2 Ignoring Irrelevant Distinctions

Consider the general case of the two cascaded tanks (Figure 2a). If we are filling the system from an initially empty state by providing a fixed input to the up-stream tank, the flow from tank A to tank B will increase monotonically with time. However, since the relation $outflowA = M^+(amountA)$ is incompletely known, $outflowA$ can “wobble” considerably while increasing. Depending on how this interacts with the behavior of $outflowB$, which is monotonically related to $amountB$, the variable $netflowB(t)$ chatters, rising and falling arbitrarily until finally returning to zero. This gives a large set of behaviors, distinguished only by the behaviors of $netflowB(t)$.

Suppose, for a particular application, we are not concerned about the detailed behavior of $netflowB(t)$, but only its sign. We would like to modify the QSIM algorithm to simulate the mechanism using only the qualitative magnitude of $netflowB$, and ignoring its direction of change. We can do this by adding a new term, *ign*, as a possible description of a direction of change. We then create a new set of qualitative state transitions, corresponding to the transitions in Figure 4b. The effect of these transitions is that all the behaviors wandering among the qualitative states $\langle +, inc \rangle$, $\langle +, std \rangle$, and $\langle +, dec \rangle$, are collapsed into a single behavior with the qualitative state $\langle +, ign \rangle$. Its eventual transition to $\langle 0, ign \rangle$ brings the system to quiescence and ends the behavior.



(a)



(b)

- (a) The full qualitative transition graph is adequate to capture continuity constraints, but permits “chattering” behaviors.
- (b) The collapsed transition graph, ignoring direction of change, eliminates chatter, but fails to detect discontinuous change.

Figure 4: Transition graphs for a single unconstrained qualitative parameter.

Unfortunately, this is not sufficient. The transitions in Figure 4b fail to capture the constraint that the derivative of a changing parameter must change continuously. For example, the transition

$$\langle +, ign \rangle \longrightarrow \langle 0, ign \rangle$$


should be excluded in case the only possible complete value for the second state is $\langle 0, inc \rangle$. Therefore, we apply a global *satisfiability filter* after each step of the prediction when an *ign* value is used. The satisfiability filter checks:

1. Whether a consistent state exists with all *ign* values replaced by one of $\{inc, std, dec\}$. This is accomplished by treating *ign* as unknown, and propagating to determine whether there is one or more complete states, consistent with the other known values.
2. If so, whether there is a consistent successor of the previous state. This is done by checking, for each parameter, whether its pair of qualitative values is consistent with a transition.


The result of qualitative simulation of the cascaded tanks, ignoring direction of change for $netflowB$, is two distinct behaviors (Figure 5). In one case both tanks reach equilibrium at the same time, while in the other, tank A reaches equilibrium before tank B.² Naturally, since direction of change is ignored, QSIM does not detect critical points or create new landmark values for $netflowB$, though it still does so for the other parameters.

This method eliminates intractable branching by changing the level of qualitative description of behavior, collapsing the *descriptions* of the different real possibilities. However, this multitude of real possibilities only arises in case of pathological interactions between the “wiggles” of different partially known M^+ relations. If we know that the relations are reasonably well-behaved, we would like to be able to take advantage of this knowledge to eliminate the chattering behaviors as impossible, and produce a stronger description of the real possible behaviors. Our second method gives us this power.

²Strictly speaking, there is only one real behavior, with both tanks reaching equilibrium at $t = \infty$. QSIM normally treats exponential approach to a limit like any other move-to-limit, and considers the possibility of reaching the limit in finite time, which corresponds to the physical perception of such a process. QSIM can be restricted to produce only the single mathematically correct prediction.



(a)



(b)

- (a) The two tanks reach equilibrium at the same time.
- (b) Tank A reaches equilibrium before tank B.

Figure 5: Two IGN behaviors of the Cascaded Tanks.

3 Applying Higher-Order Derivatives

By examining the detailed structure of a branch point in the arbitrarily chattering behavior of the parameter *netflowB*, we can obtain a clearer picture of the nature of the proliferation problem. The typical branch is a three-way branch from a point where *netflowB* has a critical point, i.e. its derivative is zero. The standard QSIM transition table [Kuipers, 1986] provides three possibilities after such a point: the parameter can increase, remain steady, or decrease. (Figure 6a)

In a system sufficiently well-understood and well-behaved to be described by a linear ordinary differential equation, the unique behavior is determined at a critical point by the first non-zero higher-order derivative at that point. For these coupled tank systems, this is the second derivative, or curvature, of the parameter in question (Figure 6b). The usual qualitative value for a parameter consists of a qualitative description of its magnitude and its direction of change. Its curvature is not made explicit, so the spurious possibilities cannot be filtered out and a proliferation of behaviors results.

The value of higher-order derivative information has been recognized previously. Williams [1984a, 1984b] showed that higher-order derivative information could disambiguate certain branching behaviors. De Kleer and Bobrow [1984] then presented methods for explicitly deriving qualitative descriptions of higher-order derivatives from the original confluences for a mechanism. Our approach starts from these correct observations, but overcomes two limitations of the methods presented in previous papers.³

- The paper [de Kleer and Bobrow, 1984] gives very little guidance on when to apply higher-order derivative information, and how to derive it for a general constraint model when it is needed.
- The straight-forward approach extends the constraint model to include terms for the higher-order derivatives and constraints linking them to the previous terms. Unfortunately, this simply pushes the problem into the higher-order terms, while adding parameters whose distinctions may cause new qualitative branches in the behavior tree.



- (a) *netflowB(t)* has a three-way branch from a critical point where *netflowB'(t)* = 0.
- (b) In case we know that *netflowB''(t)* < 0, we get only one branch where *netflowB'(t)* = 0.

Figure 6: Three-way and one-way branch

³De Kleer and Bobrow have independently identified and corrected certain errors in their paper [J. de Kleer, personal communication].

3.1 Our Approach

Our approach permits automatic identification of the problem, and automatic derivation of the appropriate constraints, which we call *curvature constraints*.⁴ We focus our attention on the parameter describing the *Highest Order Derivative (HOD)* in the system. As we shall see, it is possible for a system to have more than one *HOD*.

In outline, the algorithm consists of two steps:

1. Identify the highest-order derivative(s), *HOD* or *HODs*, in the system.
 - This method is necessary and applicable in case there is intractible branching from points where $HOD' = 0$.
2. For each *HOD*, algebraically derive an expression, valid where $HOD' = 0$, for HOD'' in terms of the other parameters of the system.
 - Use this expression to determine the sign of HOD'' at critical points, and generate only transitions consistent with this curvature.

3.2 Identifying the HOD

We can identify all the *HODs* in a general constraint model by viewing constraints as paths linking parameters. We are looking for the maximal points on chains of derivative and other two-argument constraints. We use the QSIM notation [Kuipers 1984, 1986] for constraint models, though our techniques apply generally.

- The derivative constraint leads *upward* from a parameter to its derivative.
- The two-argument constraints, M^+ , M^- , and *minus*, connect parameters *horizontally*.
- The three-argument constraints, *ADD* and *MULT*, terminate a chain and block passage. In order to avoid unnecessary blockage, it may be necessary to apply algebraic simplification rules such as those in [Kuipers, 1984, Appendix D], e.g.

$$Y = M^-(X) - M^+(X) \longrightarrow Y = M^-(X).$$

- Starting from each derivative constraint, move upward or horizontally until no further progress is possible. The maximal derivatives are the *HODs*.
- If the chain forms a closed loop (e.g. the frictionless spring) the analysis is unnecessary and may be terminated, since the loop already encodes the desired curvature constraint.

Since this process may yield several maximum points, there may be multiple *HODs*, possibly of different orders. Both the Cascaded Tanks and the Coupled Tanks have two distinct *HODs*: *netflowA* and *netflowB*. However, in the Coupled Tanks both *HODs* exhibit chatter, while in the Cascaded Tanks system only *netflowB* does.

⁴There can certainly be cases where both HOD' and HOD'' vanish, and the constraint must be stated in terms of even higher-order derivatives. However, for many models, including the damped-spring case discussed in [de Kleer and Bobrow, 1984], applying the second derivative is sufficient, so we confine our attention to that case. Extension to the higher-order case is straight-forward.

3.3 Deriving the Curvature Expression

3.3.1 The Smoothness Assumption

In order to derive an expression for the sign of the curvature, HOD'' , while still using incompletely known and possibly non-linear monotonic function constraints, we need to assume that the system is reasonably well-behaved in any local neighborhood.

Suppose $Y = M^+(X)$. This means that there is a monotonically increasing function M such that for all t , $Y(t) = M(X(t))$. The relationship between the first derivatives of X and Y is

$$Y'(t) = M'(X(t)) * X'(t),$$

where M' is the derivative of the monotonic function M . This tells us only that $Y'(t)$ and $X'(t)$ must have the same sign, since all we know about M' is that it is positive. Since M'' is unconstrained, the relationship between second derivatives is even weaker:

$$Y''(t) = M''(X(t)) * (X'(t))^2 + M'(X(t)) * X''(t).$$

The *Smoothness Assumption* says that in any local neighborhood, $M(X)$ is approximately a linear function. In practice, as we evaluate the curvature constraints, we assume that the M'' term is sufficiently small that we may regard $Y''(t)$ and $X''(t)$ as having the same sign.

3.3.2 Rules for Curvature

We can derive a set of rules for reasoning about the sign of the curvature of a particular parameter after applying the Smoothness Approximation. By treating arithmetic relations as qualitative relations on signs and allowing multiple solutions, we may write $\text{sign}(x + y) = \text{sign}(x) + \text{sign}(y)$.

For brevity of notation, we define:

$$\text{sd}2(X) = \text{sign}\left(\frac{d^2}{dt^2}X\right), \quad \text{sd}(X) = \text{sign}\left(\frac{d}{dt}X\right).$$

For a parameter X , $\text{sd}(X)$ is just its direction of change, and so is explicitly represented by QSIM at each time-point. $\text{sd}2(X)$ is its curvature, so the curvature expression solves for curvature in terms of explicitly available information.

The rules for deriving the explicit curvature expression are the following:

1. Start with the expression $\text{sd}2(HOD)$.
2. Apply the following rules for qualitative curvature in depth-first order wherever they are applicable, to propagate $\text{sd}2$ terms through all possible constraints.

$$\begin{aligned} \text{sd}2(M(x)) &= \text{sd}2(x) \\ \text{sd}2(x + y) &= \text{sd}2(x) + \text{sd}2(y) \\ \text{sd}2(x * y) &= y * \text{sd}2(x) + x * \text{sd}2(y) + 2\text{sd}(x) * \text{sd}(y) \\ \text{sd}2(-x) &= -\text{sd}2(x) \end{aligned}$$

The first rule depends on the Smoothness Assumption. The others are straight-forward consequences of the theorems of differential calculus. This process terminates when every explicit parameter is either an exogenous variable, or explicitly linked to its derivative.

3. Apply the final transformations:

$$\begin{aligned}sd2(constant) &= 0 \\sd2(x) &= sd(x') \\sd(HOD) &= 0\end{aligned}$$

and simplify the result.

The result is an expression for $sd2(HOD)$ in terms of explicitly available information. It is used as a constraint, applied only at critical points of the HOD, to select among branches such as that in Figure 6.

3.4 The Coupled Tanks example

Let us follow the derivation of the curvature expression on the Coupled Tanks (Figure 2b).

The parameters $netflowA$ and $netflowB$ are both HODs in this system, and both exhibit the chattering behavior. We will explicitly derive an expression for the sign of the curvature of $netflowB$, $sd2(netflowB)$, valid at critical points, i.e. where $sd(netflowB) = 0$.

$$sd2(netflowB) = sd2(crossflowAB) - sd2(outflowB)$$

$$\begin{aligned}sd2(crossflowAB) &= sd2(pressureDiff) \\&= sd2(pressureA) - sd2(pressureB) \\&= sd2(amountA) - sd2(amountB) \\&= sd(netflowA) - sd(netflowB) \\&= sd(netflowA)\end{aligned}$$

$$\begin{aligned}sd2(outflowB) &= sd2(pressureB) \\&= sd2(amountB) \\&= sd(netflowB) \\&= 0.\end{aligned}$$

The resulting constraint, $sd2(netflowB) = sd(netflowA)$, tells us that the sign of the curvature of $netflowB$ at a critical point is the same as the sign of the slope of $netflowA$ at the same point in time. A similar derivation gives us the constraint $sd2(netflowA) = sd(netflowB)$, applying at critical points of $netflowA$. Applying these two constraints, we get the following single behavior for the scenario of filling the Coupled Tank system from empty (Figure 7).

We predict an unambiguous behavior: $amountA$ and $amountB$ increase monotonically from zero to their equilibrium values; $netflowA$ decreases monotonically to zero; and $netflowB$ increases monotonically from zero to some maximum value, then decreases back to zero. This description is stronger than that produced by our first method, in that all parameters now have complete qualitative descriptions. In the method of ignoring irrelevant distinctions the HODs, and certain other closely related parameters exhibiting chatter, would be described in terms of magnitude only.

4 Conclusions

Although the class of qualitative simulation problems we treat here has considerable importance in its own right, the methods we have developed have a more general significance. A fundamental decision in the modeling of a system is selection of the level of detail for the model. Both methods explicitly manipulate that level of description.

- The method of ignoring irrelevant distinctions, by ignoring direction of change for certain parameters, represents a change in the level of detail of the model. In its current form, it is a knowledge engineering method, to be applied explicitly as a model is being developed and debugged. We expect that, using techniques similar to our second method, it can be extended to recognize "chattering" situations automatically, and decide when to ignore which distinctions.
- The higher-order derivative method, by applying the smoothness assumption and a more powerful inference about the algebraic structure of the constraint model, allows the simulation to produce a complete qualitative description of all the parameters, including the highest-order derivatives. It also preserves QSIM's ability to create new landmark values.

Chiu and Kuipers [1987] present the details of our methods for automatically deriving and applying the curvature constraint. While these methods handle important cases of coupled systems, there are additional cases of intractable branching in the more general second-order system. Lee, Chiu, and Kuipers [1987] extend this work by applying the curvature method and two additional constraints to handle an important case of the damped spring system.




Figure 7: Single HOD behavior of the Coupled Tanks.

Thus, our two methods are able to solve the “chattering” problem for *coupled systems*, in the most general case by ignoring certain distinctions, and, when more knowledge is available, by applying more powerful inferences to produce a stronger result. These methods provide important mathematical and computational tools for qualitative simulation, and move us toward our goal of reasoning qualitatively about complex physical systems.

5 References

1. C. Chiu and B. Kuipers. 1987. The curvature constraint algorithm and a complete simulation of a two tank system. University of Texas Computer Science TR-xx, forthcoming.
2. J. de Kleer and D. G. Bobrow. 1984. Qualitative reasoning with higher-order derivatives. In *Proceedings of the National Conference on Artificial Intelligence (AAAI-84)*. Los Altos, CA: Morgan-Kaufman Publishers.
3. J. de Kleer and J. S. Brown. A qualitative physics based on confluences. *Artificial Intelligence* 24: 7 - 83, (1984).
4. J. de Kleer. How circuits work. *Artificial Intelligence* 24: 205 - 280, (1984).
5. K. D. Forbus. Qualitative process theory. *Artificial Intelligence* 24: 85 - 168, (1984).
6. B. J. Kuipers. 1984. Commonsense reasoning about causality: deriving behavior from structure. *Artificial Intelligence* 24: 169 - 204.
7. B. J. Kuipers. 1985. The limits of qualitative simulation. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence (IJCAI-85)*. William Kaufman, Los Altos, CA.
8. B. J. Kuipers. 1986. Qualitative simulation. *Artificial Intelligence* 29: 289 - 338.
9. B. Kuipers. 1987. Qualitative Simulation as Causal Explanation. To appear in *IEEE Transactions on Systems, Man, and Cybernetics* 17, No. 3, 1987; special issue on Causal and Strategic Aspects of Diagnostic Reasoning.
10. W. W. Lee, C. Chiu, and B. Kuipers. 1987. Steps toward constraining qualitative simulation. University of Texas Computer Science TR-87-44.
11. P. A. Sachs, A. M. Paterson, and M. H. M. Turner. 1986. Escort - an expert system for complex operations in real time. *Expert Systems* v. 3, n. 1, pp. 22-29.
12. B. Williams. 1984. Qualitative analysis of MOS circuits. *Artificial Intelligence* 24: 281 - 346. (a)
13. B. Williams. 1984. The use of continuity in a qualitative physics. In *Proceedings of the National Conference on Artificial Intelligence (AAAI-84)*, 350 - 354. (b)
14. Brian Williams. 1986. Doing time: putting qualitative reasoning on firmer ground. In *Proceedings of the Fifth National Conference on Artificial Intelligence (AAAI-86)*. Los Altos, CA: Morgan Kaufman Publishers, pp. 105-112.




Figure A: Cascaded Tank structure




Figure B: Coupled Tank structure