

Qualitative Simulation*

Benjamin Kuipers[†]

April 5, 2001

Qualitative simulation predicts the set of possible behaviors consistent with a qualitative differential equation model of the world. Its value comes from the ability to express natural types of incomplete knowledge of the world, and the ability to derive a provably complete set of possible behaviors in spite of the incompleteness of the model.

A qualitative differential equation model (QDE) is an abstraction of an ordinary differential equation, consisting of a set of real-valued variables and functional, algebraic and differential constraints among them. A QDE model is qualitative in two senses. First, the values of variables are described in terms of their ordinal relations with a finite set of symbolic landmark values, rather than in terms of real numbers. Second, functional relations may be described as monotonic functions (increasing or decreasing over particular ranges) rather than by specifying a functional form. These purely qualitative descriptions can be augmented with semi-quantitative knowledge in the form of real bounding intervals around unknown real values and real-valued bounding envelope functions around unknown real-valued functions. Qualitative and semi-quantitative models can be derived by composing model fragments and collecting the associated modeling assumptions.

Qualitative simulation starts with a QDE and a qualitative description of an initial state. Given a qualitative description of a state (called a qstate), it predicts the qualitative state descriptions that can possibly be direct successors of the current state description. Repeating this process produces a graph of qualitative state descriptions, in which the paths starting from the root are the possible qualitative behaviors. The graph of qualitative states is pruned according to criteria derived from the theory of ordinary differential equations, in order to preserve the guarantee that all possible behaviors are predicted. Abstraction methods have also been developed to simplify the resulting qualitative behaviors.

The resulting graph of qualitative states (the behavior graph) can still be quite large, requiring automated methods based on temporal logic model-checking to determine whether the qualitative prediction implies a desired conclusion. Conclusions derived in this way can be used in the design and validation of dynamical systems such as controllers. A set of qualitative models and their associated predictions can also be unified with a stream of observations to monitor an ongoing dynamical system or to do system identification on a partial model.

Ongoing research topics include qualitative simulation and abstraction methods, the use of various types of quantitative knowledge, automated ways to determine the conclusions to draw

*This work has taken place in the Qualitative Reasoning Group at the Artificial Intelligence Laboratory, The University of Texas at Austin. Research of the Qualitative Reasoning Group is supported in part by NSF grants IRI-9504138 and CDA 9617327, by NASA grant NAG 9-898, and by the Texas Advanced Research Program under grants 003658-242 and 003658-347.

[†]Computer Science Department, University of Texas at Austin, Austin, Texas 78712 USA. kuipers@cs.utexas.edu.

from a predicted behavior graph, design and verification methods, online monitoring frameworks, and modeling methods suited for particular application domains. The specific notations in this article are those used in the QSIM representation [20, 21], but the concepts covered include the related ideas from [8, 12, 30, 31].

1 The Qualitative Model Representation

Like an *ordinary* differential equation, a *qualitative* differential equation model consists of a set of variables related by constraints. (Figure 1 shows an example of the QSIM code for a QDE describing a simple U-tube system consisting of two tanks, A and B, connected by a thin channel.) A variable represents a continuously differentiable function over the *extended* real number line, $v : \mathbb{R}^* \rightarrow \mathbb{R}^*$, including $\pm\infty$. However, in a QDE model, the range of each variable, including the independent variable time, is described qualitatively by a *quantity space*. A quantity space is a finite, totally ordered set of symbolic *landmark values* representing qualitatively important values in the real number line (see figure 1). Every quantity space includes landmarks for zero and positive and negative infinity. A purely qualitative model specifies only the ordinal relations among landmarks, though as we shall see below, semi-quantitative extensions may provide bounds on the possible real values corresponding to a landmark.

The algebraic and differential constraints in a QDE are simple and familiar equations, universally quantified over t .


$$\begin{aligned} (\text{add } x \ y \ z) &\equiv x(t) + y(t) = z(t) \\ (\text{mult } x \ y \ z) &\equiv x(t) \cdot y(t) = z(t) \\ (\text{minus } x \ y) &\equiv y(t) = -x(t) \\ (\text{d/dt } x \ y) &\equiv \frac{d}{dt}x(t) = y(t) \\ (\text{constant } x) &\equiv \frac{d}{dt}x(t) = 0 \end{aligned}$$

Since they are asserted as individual constraints, rather than composed as hierarchical expressions in traditional algebra, a QDE must include explicit variables for subexpressions. However, a QDE may also include constraints representing unknown functions in the set M^+ of monotonically increasing continuously differentiable functions (satisfying additional technical conditions discussed in [21]).

$$\begin{aligned} (M^+ \ x \ y) &\equiv y(t) = f(x(t)), f \in M^+ \\ (M^- \ x \ y) &\equiv y(t) = -f(x(t)), f \in M^+ \end{aligned}$$

The M^+ and M^- constraints make it possible to express a QDE model including functions whose explicit form is not known, and which are only described in terms of monotonicity. An algebraic or functional constraint may specify corresponding values, which are tuples of landmark values known to satisfy the constraint. A QDE may also explicitly describe the boundaries of its domain of applicability by specifying transition conditions that carry the behavior into a different model. The U-tube model in Figure 1 illustrates each of these features.

The qualitative magnitude of a variable is described either as a landmark value or as an open interval between two adjacent landmarks in the quantity space of that variable. The qualitative value of a variable is described as its qualitative magnitude and the sign of its derivative (its direction of change: inc, std, or dec). (Note that the antecedents of the transition conditions in



```

(define-QDE U-Tube
  (quantity-spaces
    (amtA      (minf 0 AMAX inf))
    (pressureA (minf 0 inf))
    (amtB      (minf 0 BMAX inf))
    (pressureB (minf 0 inf))
    (pAB       (minf 0 inf))
    (flowAB    (minf 0 inf))
    (-flowAB   (minf 0 inf))
    (total     (minf 0 inf)))
  (constraints
    ((M+ amtA pressureA)      (0 0) (inf inf))
    ((M+ amtB pressureB)      (0 0) (inf inf))
    ((add pAB pressureB pressureA))
    ((M+ pAB flowAB)          (minf minf) (0 0) (inf inf))
    ((minus flowAB -flowAB))
    ((d/dt amtB flowAB))
    ((d/dt amtA -flowAB))
    ((add amtA amtB total))
    ((constant total)))
  (transitions
    ((amtA (AMAX inc)) -> tank-A-overflow)
    ((amtB (BMAX inc)) -> tank-B-bursts)))

```

Figure 1: QSIM code for the U-tube model.

This QDE model can be written in the form of an ODE

$$\frac{d}{dt}B = -\frac{d}{dt}A = f(p_1(A) - p_2(B))$$

except that the functions $f, p_1, p_2 \in M^+$ are only qualitatively described. All three M^+ constraints include tuples of corresponding landmark values. Their significance in this case is to exclude horizontal and vertical asymptotes.

Figure 1 are specified by the qualitative values of particular variables.) A qualitative state of a model is a tuple of associations of qualitative values to each variable in the model.

Time is described in the same way as every other variable. Since its direction of change is always inc, time progresses through an alternating sequence of landmark values (called time-points) and open intervals between adjacent time-points. The time-points are defined as those points in time when the qualitative state of the model (i.e., the qualitative value of any variable) changes.

A qualitative behavior is a sequence of qualitative states, where each state is the immediate successor of the one before it. Because of the qualitative representation, it is possible for a finite sequence of qualitative states to represent the behavior of a system from its initial state at $t = 0$ to a final state at $t = \infty$. For example, one possible behavior of the U-tube model in Figure 1, initialized with Tank A full and Tank B empty, is the following three-state behavior concluding with a state where both Tank A and Tank B are partly full. We see new landmark values being created and inserted into quantity spaces when new critical values are defined; i.e., when a qualitative magnitude lies in an open interval, but direction of change is *std*.

t	t_0	(t_0, t_1)	t_1
$amtA$	$\langle AMAX, dec \rangle$	$\langle (0, AMAX), dec \rangle$	$\langle (0, AMAX), std \rangle$ $= \langle A_0, std \rangle$
$pressureA$	$\langle (0, \infty), dec \rangle$	$\langle (0, \infty), dec \rangle$	$\langle (0, \infty), std \rangle$ $= \langle P_0, std \rangle$
$amtB$	$\langle 0, inc \rangle$	$\langle (0, BMAX), inc \rangle$	$\langle (0, BMAX), std \rangle$ $= \langle A_1, std \rangle$
$pressureB$	$\langle 0, inc \rangle$	$\langle (0, \infty), inc \rangle$	$\langle (0, \infty), std \rangle$ $= \langle P_1, std \rangle$
pAB	$\langle (0, \infty), dec \rangle$	$\langle (0, \infty), dec \rangle$	$\langle 0, std \rangle$
$flowAB$	$\langle (0, \infty), dec \rangle$	$\langle (0, \infty), dec \rangle$	$\langle 0, std \rangle$
$total$	$\langle (0, \infty), std \rangle$ $= \langle TO_0, std \rangle$	$\langle (0, \infty), std \rangle$ $= \langle TO_0, std \rangle$	$\langle (0, \infty), std \rangle$ $= \langle TO_0, std \rangle$

Figure 2(a) shows a plot of this qualitative behavior (each qualitative value is plotted at a landmark, or midway between two landmarks). Because the tanks have unknown sizes (the landmarks AMAX and BMAX) and geometries (the monotonic functions linking amount and pressure), the behavior graph for this model is a tree of three behaviors.

There are extensions to the representation not discussed here, including the region transition and discontinuous changes shown in Figure 2(b). See [21] for details.

2 Qualitative Simulation

The QSIM algorithm [20, 21] performs qualitative simulation by deriving the immediate successors of each qualitative state, repeating this step to grow the behavior graph from the initial state at its root. In order to guarantee that all possible behaviors are predicted, we require first that all possible qualitative value transitions are predicted, and second, that combinations of qualitative values are only deleted when they are inconsistent.




Figure 2: Qualitative simulation of the U-tube.

Two of the three qualitatively distinct behaviors predicted by the U-tube model, from an initial state $AmtA(t_0) = AMAX$, $AmtB(t_0) = 0$.

- (a) The system reaches equilibrium with both Tank A and Tank B partially full.
- (b) Tank B overflows and bursts, causing a transition to a new model where there is no backpressure to impede flow out of tank A, so the system drains to empty.

In the third behavior (not shown), Tank B reaches equilibrium at $AmtB(t_1) = BMAX$.

Table 1 enumerates all transitions from each qualitative value description to its possible successors. The validity of this table follows directly from the Intermediate Value and Mean Value Theorems from elementary calculus.

The successor generation phase of QSIM consists of the following steps, given a QDE and a current state S .

1. (Value generation.) For each variable in S , generate all possible successor values using Table 1.
2. (Constraint filtering.) For each constraint in the QDE, which applies to a tuple of variables, generate all corresponding tuples of successor values. Delete each tuple that violates its constraint.
3. (Local consistency filtering.) For each pair of constraints that are adjacent, in the sense that they share a variable v , and for each tuple of one constraint that assigns a value, say x , to v , delete that tuple if there is no tuple associated with the other constraint that also assigns the value x to v .
4. (State generation.) From the remaining tuples of values associated with constraints, exhaustively enumerate all consistent complete assignments of values to variables. Create a successor state for S from each of these assignments.

Once successor states have been added to the behavior graph, they can be analyzed and the description augmented in several ways. In some cases inconsistencies can be identified that were not visible at the successor-generation level, allowing states to be pruned from the graph.

State filters consider information local to the current state and perhaps its immediate predecessor. Inconsistency can be propagated from a state to its predecessors.

- A quiescent state (fixed point) can be recognized because all directions of change are *std*. In some cases, its stability can also be determined.
- Transitions to infinite values and infinite times must satisfy additional constraints.
- Higher-order derivative constraints can sometimes be derived algebraically from the QDE and applied to eliminate certain successor states.
- New landmarks and new corresponding value tuples can be created explicitly to describe critical values and other uniquely determined values in quantity spaces.
- A region transition is identified when the current state matches the antecedent to a transition rule. The current state is linked to a new state, created with respect to the QDE for the new operating region, with values mapped from the current state. The transition may represent a discontinuous change, or it may represent a re-description of the current state within a new model.

Behavior filters derive properties of the entire behavior terminating in the current state, to augment the behavior description and sometimes determine its inconsistency.

- A periodic behavior can be identified by matching the new state to one of its predecessors in the behavior graph.

We can enumerate the possible successor relations from one qualitative value to the next for a continuously differentiable variable $v : [a, b] \rightarrow \mathfrak{R}^*$. There are two sets of successor rules, depending on whether the antecedent state is a time-point or a time-interval. Let $l_{j-1} < l_j < l_{j+1}$ be three adjacent landmarks in the quantity space for v .

- **P-Successors:** point to interval.

$QV(v, t_i)$	\Rightarrow	$QV(v, t_i, t_{i+1})$
$\langle l_j, std \rangle$		$\langle l_j, std \rangle$
$\langle l_j, std \rangle$		$\langle (l_j, l_{j+1}), inc \rangle$
$\langle l_j, std \rangle$		$\langle (l_{j-1}, l_j), dec \rangle$
$\langle l_j, inc \rangle$		$\langle (l_j, l_{j+1}), inc \rangle$
$\langle l_j, dec \rangle$		$\langle (l_{j-1}, l_j), dec \rangle$
$\langle (l_j, l_{j+1}), inc \rangle$		$\langle (l_j, l_{j+1}), inc \rangle$
$\langle (l_j, l_{j+1}), dec \rangle$		$\langle (l_j, l_{j+1}), dec \rangle$
$\langle (l_j, l_{j+1}), std \rangle$		$\langle (l_j, l_{j+1}), std \rangle$
$\langle (l_j, l_{j+1}), std \rangle$		$\langle (l_j, l_{j+1}), inc \rangle$
$\langle (l_j, l_{j+1}), std \rangle$		$\langle (l_j, l_{j+1}), dec \rangle$

- **I-Successors:** interval to point.

$QV(v, t_i, t_{i+1})$	\Rightarrow	$QV(v, t_{i+1})$
$\langle l_j, std \rangle$		$\langle l_j, std \rangle$
$\langle (l_j, l_{j+1}), inc \rangle$		$\langle l_{j+1}, std \rangle$
$\langle (l_j, l_{j+1}), inc \rangle$		$\langle l_{j+1}, inc \rangle$
$\langle (l_j, l_{j+1}), inc \rangle$		$\langle (l_j, l_{j+1}), inc \rangle$
$\langle (l_j, l_{j+1}), inc \rangle$		$\langle (l_j, l_{j+1}), std \rangle$
$\langle (l_j, l_{j+1}), dec \rangle$		$\langle l_j, std \rangle$
$\langle (l_j, l_{j+1}), dec \rangle$		$\langle l_j, dec \rangle$
$\langle (l_j, l_{j+1}), dec \rangle$		$\langle (l_j, l_{j+1}), dec \rangle$
$\langle (l_j, l_{j+1}), dec \rangle$		$\langle (l_j, l_{j+1}), std \rangle$
$\langle (l_j, l_{j+1}), std \rangle$		$\langle (l_j, l_{j+1}), std \rangle$

Table 1: The qualitative successor rules

- The phase space trajectory of a dynamical system can only intersect itself if the behavior is periodic. Qualitative behaviors that self-intersect without creating a cycle are inconsistent.
- Terms equivalent to potential and kinetic energy, and conservative and non-conservative work, can be derived from some QDE models, and tested for consistency.
- When semi-quantitative information is associated with the QDE, it can be propagated to refine or refute the behavior description. (Section 3.)

The QSIM Guaranteed Coverage Theorem states that the QSIM behavior graph describes all real solutions to ODE models consistent with the given QDE and initial qualitative state. This follows directly from the fact that all possible successor values are generated, and that states and behaviors are deleted only when proved to be inconsistent.

There is no converse guarantee that every predicted qualitative behavior corresponds to a real solution to some ODE described by the QDE. While the constraint satisfaction algorithm in QSIM is sound and complete, there may well be a Gödel-like incompleteness theorem stating that the properties of real dynamical systems are too rich to be captured by any finite set of symbolic constraints.

2.1 Tractability

The set of behaviors generated by qualitative simulation may include many distinctions unimportant to the model-builder, due to the fixed level of description implied by the qualitative value and qualitative state representations. There are two classes of such unimportant distinctions. In the first, there is a region of the state space of the QDE where the qualitative behavior is unconstrained. Recently, methods have been developed for identifying such a “chattering” region and replacing it with an abstract state whose predecessors and successors describe the trajectories into and out of that region [4].

In the second class, two or more events (qualitative value transitions) take place, but their temporal order is not constrained by the QDE. Based on the concept of interacting histories [14, 32] methods have been developed for qualitative simulation of a QDE decomposed into sub-models [5]. The interactions between the histories of the sub-models are considered only when they are needed to permit simulation of the sub-model. Events internal to separate sub-models are not explicitly related, so they do not require explicit branches in the graph of qualitative behaviors.

While the behavior graph predicted for a complex QDE model may still be quite large, these two methods have, in principle, eliminated the problem of intractable branching in qualitative simulation by eliminating the two sources of explicit distinctions unimportant to the model-builder.

2.2 Querying the behavior graph

An individual qualitative behavior is a description of the time-evolution of the variables in the QDE, and is not difficult to interpret, whether it is purely qualitative or if it is augmented with bounding intervals and envelopes. However, a large behavior graph represents a disjunctive prediction with many disjuncts, and requires automated interpretation tools.

A particularly interesting tool that has been developed recently is temporal logic model-checking, applied to the behavior graph output by QSIM [29]. The branching-time temporal logic CTL* [11]

is particularly well suited to expressing statements of interest about the QSIM behavior graph. We can express:

QSIM predicates on states	qvalue, quiescent, cycle, etc.
Logical connectives	and, or, not, implies.
Temporal path relations	eventually, always, next, until, etc.
Modal quantifiers	necessarily, possibly.

The behavior graph output by QSIM can be interpreted as a branching-time temporal model, against which a temporal assertion can be checked for validity. There are efficient incremental model-checking algorithms that can be used to check whether a temporal model structure is an interpretation of a given statement in CTL* [2]. The model-checking algorithm is sound and complete. However, the QSIM Guaranteed Coverage Theorem provides only a one-sided guarantee about the relation between the QSIM behavior graph and the set of predicted behaviors: QSIM predicts every real behavior, but some predictions could be spurious, and not correspond to any real behavior. Therefore, model-checking can prove a universal statement in temporal logic (one of the form *necessarily*(P)), but not an existential statement (one of the form *possibly*(P)), since the behavior that is identified as the interpretation for P could be a spurious behavior [29]. Temporal logic model checking can be used to prove properties of dynamical systems such as non-linear controllers, even with incomplete knowledge [22].

2.3 Guided Simulation

The relation between temporal logic and qualitative simulation can be carried one step farther, to allow assertions in temporal logic to be treated as part of the model [3]. The extended qualitative simulator, TeQSIM, generates only qualitative behaviors that satisfy the temporal logic assertions as well as the constraints in the QDE and the requirements of continuity.

This approach has two major uses. First, it extends the expressive power available to the model-builder to state properties of the system that are difficult to capture in the constraint language of the QDE. An example is the ability to describe time-varying behavior of exogenous variables, including specifying bounds on the time of occurrence of discrete events. The second use is to allow the model-builder to focus the simulator's attention on a subset of the state space of the model described by the temporal logic assertions, rather than to explore the larger space of all possible behaviors.

3 Semi-Quantitative Simulation

Partial knowledge can be quantitative as well as purely qualitative. The QDE and the qualitative behaviors produced by QSIM can serve as a symbolic and algebraic framework for reasoning with several representations of incomplete quantitative knowledge.

A landmark value is a symbolic name for an unknown real number, described in terms of its ordinal relations with other landmark values, and the corresponding value tuples it participates in. A natural form of partial quantitative knowledge about the unknown real number corresponding to a landmark is a bounding interval, whose endpoints can be real numbers or $\pm\infty$. Two assertions of bounding intervals for the same landmark can be combined simply by intersecting the intervals. A smaller resulting interval corresponds to more precise knowledge about the value of that landmark.

An empty intersection means that no value can be consistently assigned to that landmark, so the current qualitative behavior is refuted.

A monotonic function constraint ($M+ \ x \ y$) is a qualitative description of an unknown function $y = f(x)$, describing the shape of f only as monotonically increasing. A natural form of partial quantitative information about f is to provide a pair of real-valued “static envelope” functions \bar{f} and \underline{f} that bound f above and below: $\underline{f}(x) \leq f(x) \leq \bar{f}(x)$ for all x (Figure 3). It can also be useful to assert bounds on the slope $f'(x)$ of a monotonic function.

The QDE, augmented with bounds on landmark values and static envelopes on monotonic function constraints, is referred to as a “semi-quantitative differential equation” or SQDE. Qualitative simulation augmented with semi-quantitative inference is called “semi-quantitative simulation” or SQSIM. Figure 3 shows a SQDE model of a water tank.

Purely qualitative simulation of the water tank SQDE predicts three qualitative behaviors: equilibrium partly full, overflow, and equilibrium exactly at the brim. The following table of qualitative values shows the three states of the equilibrium-partly-full behavior, including the creation of new landmarks for uniquely specified values.

t	t_0	(t_0, t_1)	t_1
<i>amount</i>	$\langle 0, inc \rangle$	$\langle (0, FULL), inc \rangle$	$\langle (0, FULL), std \rangle$ $= \langle a_0, std \rangle$
<i>outflow</i>	$\langle 0, inc \rangle$	$\langle (0, \infty), inc \rangle$	$\langle (0, \infty), std \rangle$ $= \langle o_1, std \rangle$
<i>inflow</i>	$\langle if^*, std \rangle$	$\langle if^*, std \rangle$	$\langle if^*, std \rangle$
<i>netflow</i>	$\langle (0, \infty), dec \rangle$ $= \langle n_0, dec \rangle$	$\langle (0, \infty), dec \rangle$ $= \langle (0, n_0), dec \rangle$	$\langle 0, std \rangle$ $\langle 0, std \rangle$

3.1 Propagating Interval Bounds

The simplest semi-quantitative extension to QSIM, called Q2, is based on interval arithmetic [21, chapter 9]. A qualitative behavior can be interpreted as a set of algebraic and functional constraints among landmark values. The following constraints, called the “Q2 equations”, are derived from the corresponding landmark value tuples implied by time-point qstates in the qualitative behavior above, from the bounding envelopes on monotonic function constraints, and from time-interval states via the Mean Value Theorem.

$$\begin{aligned}
 o_1 &= f(a_0) \\
 if^* &= n_0 + 0 \\
 if^* &= 0 + o_1 \\
 span(n_0, 0) &= d(0, a_0)/d(t_0, t_1) \\
 slope(f) &= d(0, o_1)/d(0, a_0) \\
 d(0, a_0) &= a_0 - 0 \\
 d(t_0, t_1) &= t_1 - t_0 \\
 d(0, o_1) &= o_1 - 0
 \end{aligned}$$

```

(define-QDE Water-tank
  (quantity-spaces
    (amount (0 full inf))
    (outflow (0 inf))
    (inflow (0 if* inf))
    (netflow (minf 0 inf)))
  (constraints
    ((M+ amount outflow) (0 0) (inf inf))
    ((add netflow outflow inflow))
    ((d/dt amount netflow))
    ((constant inflow)))
  (transitions ((amount (full inc)) -> t))
  (envelopes
    ((M+ amount outflow) (upper ue) ; ue, ui, le, and li
                          (u-inv ui) ; are the names of the
                          (lower le) ; envelope functions
                          (l-inv li)))
  (initial-ranges ((amount full) (80 100))
                  ((inflow if*) (4 8))
                  ((time t0) (0 0))))

```




Figure 3: Semi-quantitative model of a water tank: QDE plus bounds on landmarks and envelopes around monotonic functions. Envelope functions and their inverses are specified, to allow propagation in both directions.

-
- Where x is a primitive term (landmark, $d(p, q)$, $span(p, q)$, or $slope(f)$) the bounding interval $range(x) = [\underline{x}, \bar{x}]$ is retrieved from a table and updated by propagation.
 - Where x and y are expressions evaluating to intervals, $k \in \mathfrak{R}$, and $f \in M^+$, the interval bound on a complex expression is computed from its parts.

$$\begin{aligned}
range(k) &= [k, k] \\
range(x + y) &= [\underline{x} + \underline{y}, \bar{x} + \bar{y}] \\
range(x - y) &= [\underline{x} - \bar{y}, \bar{x} - \underline{y}] \\
range(span(x, y)) &= [min(\underline{x}, \underline{y}), max(\bar{x}, \bar{y})] \\
range(|x|) &= \begin{cases} [\underline{x}, \bar{x}] & \text{if } \underline{x} \geq 0 \\ [-\bar{x}, -\underline{x}] & \text{if } \bar{x} \leq 0 \\ [0, max(\bar{x}, -\underline{x})] & \text{otherwise} \end{cases} \\
range(x \cdot y) &= [\underline{x} \cdot \underline{y}, \bar{x} \cdot \bar{y}], \text{ if } \underline{x} \geq 0 \text{ and } \underline{y} \geq 0 \\
range(1/x) &= [1/\bar{x}, 1/\underline{x}], \text{ if } \underline{x} > 0 \\
range(x/y) &= range(x \cdot (1/y)) \\
range(f(x)) &= [f(\underline{x}), \bar{f}(\bar{x})] \\
range(f^{-1}(y)) &= [\bar{f}^{-1}(\underline{y}), \underline{f}^{-1}(\bar{y})]
\end{aligned}$$

The entries for $x \cdot y$ and $1/x$ describe only the cases where $x, y \geq 0$. It is straight-forward to extend to the full case split on possible combinations of signs, and to handle $\pm\infty$ as bounds.

Table 2: Interval arithmetic

$d(p, q)$ represents the distance between landmarks p and q ; $span(p, q)$ is the interval enclosing both p and q ; and $slope(f)$ is the range bound on $f'(x)$ for a monotonic function f .

The initially given interval values associated with the landmarks $FULL$, if^* and t_0 are propagated across the Q2 equations, following the rules in Table 2. Newly derived values are intersected with old values until a fixed point is reached or until an empty interval is derived. The result of propagation for the equilibrium-partly-full behavior is the following set of bounds for landmark values. In the other two behaviors, an empty interval is derived for some landmark, so the behaviors are refuted.

<i>amount</i>	<i>FULL</i>	[80, 100]
	a_0	[5.15, 56.7]
<i>outflow</i>	o_1	[4, 8]
<i>inflow</i>	if^*	[4, 8]
<i>netflow</i>	n_0	[4, 8]
<i>time</i>	t_0	[0, 0]
	t_1	[0.644, $+\infty$]

Semi-quantitative inference is implemented in QSIM as a filter applied to each partial behavior whenever a successor state is added to the behavior graph. When a partial behavior is refuted, its extensions need not be computed, reducing the branching factor of the behavior graph.

3.2 Order-of-Magnitude Constraints

A different form of partial quantitative knowledge is order-of-magnitude constraints on landmark values [6, 26]. These relations can also be propagated across the Q2 equations derived from a behavior. The result is additional predicted order-of-magnitude relations, or contradictions that refute behaviors from the behavior graph, just as in Q2.

3.3 State Interpolation

The temporal granularity of the qualitative behavior description predicted by QSIM, and hence of the Q2 equations, is determined by the qualitative value changes that take place in the behavior. Thus, semi-quantitative inference takes place over time intervals that are quite large and sometimes infinite, making it difficult to draw strong conclusions.

Q3 [1] addresses this problem by interpolating new landmarks into intervals in quantity spaces, including new time-points into large time-intervals. This provides smaller intervals of change, so the derived error bounds are tighter. The effect is essentially the same as Euler integration, approximating a continuous curve above and below by rectangles. It is possible to show that as the uncertainty in the SQDE approaches zero, and as the size of largest time-interval in the behavior approaches zero, the resulting semi-quantitative prediction converges to the real-valued solution to the corresponding ODE [1].

3.4 Dynamic Envelopes

The rectangular bounds on a variable's behavior derived for time-interval states by Q2 and Q3 are consequences of the Mean Value Theorem and the bounds on the rate of change of the variable over that time-interval. In many cases, we can derive stronger bounds.

Just as static envelopes define real-valued functions providing upper and lower bounds to partially known monotonic functions, it is possible to infer real-valued functions providing upper and lower bounds on the values of variables over time. These "dynamic envelope" functions are the solutions to a real-valued ODE model that can be derived from the bounds and static envelopes in the SQDE and simulated numerically [19].

For example, consider a model of an open-ended U-tube, with constant flow q into tank A and a pressure-driven flow out of tank B:

$$\begin{aligned} F_{ab} &= f(p_1(a) - p_2(b)) \\ \dot{a} &= q - F_{ab} \\ \dot{b} &= F_{ab} - g(b) \end{aligned}$$

The SQDE includes interval bounds $[q, \bar{q}]$ on the landmark value q and static envelopes $\underline{f}, \bar{f}, g, \bar{g}, \underline{p}_1, \bar{p}_1, \underline{p}_2, \bar{p}_2$ on the unknown functions $f, g, p_1, p_2 \in M^+$. The bounding ODE system, which must have double the order of the original QDE, is:

$$\begin{aligned} \dot{\underline{a}} &= \underline{q} - \bar{f}(\bar{p}_1(\underline{a}) - \underline{p}_2(\underline{b})) \\ \dot{\bar{a}} &= \bar{q} - \underline{f}(\underline{p}_1(\bar{a}) - \bar{p}_2(\bar{b})) \\ \dot{\underline{b}} &= \underline{f}(\underline{p}_1(\underline{a}) - \bar{p}_2(\underline{b})) - \bar{g}(\underline{b}) \end{aligned}$$

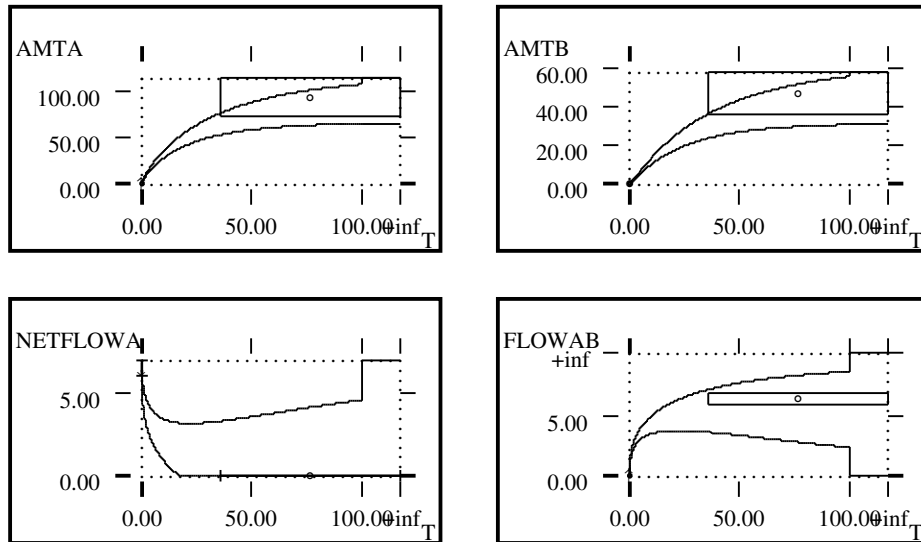


Figure 4: Dynamic envelopes are trajectories guaranteed to bound the true trajectory. They are often tighter than the Q2 bounds, which are represented by rectangular boxes.

$$\dot{\bar{b}} = \bar{f}(\bar{p}_1(\bar{a}) - \bar{p}_2(\bar{b})) - \bar{g}(\bar{b})$$

Dynamic envelopes give improved bounds on the behavior over an interval starting at the initial state, but eventually diverge and provide no constraint farther away (Figure 4). Thus, dynamic envelopes should be combined with inference using the symbolic Q2 and Q3 methods.

3.5 Research Problem: Soft Bounds

Semi-quantitative inference based on intervals and envelopes preserves the QSIM Guaranteed Coverage Theorem: only behaviors that are provably inconsistent are deleted. An important research direction is extending semi-quantitative inference to partial knowledge of quantity in the form of probability distribution functions: Gaussian distributions and more general pdfs. These representations seldom support inference of direct contradictions, making it difficult to refute a behavior entirely. Rather, the goal must be to infer a degree of belief in a behavior, and a distribution of belief over a set of behaviors. In the monitoring context (next section), it will be useful to distribute belief over a set of alternate hypothesized models as well.

4 Monitoring and System Identification

Monitoring is the process of comparing an observation stream with predictions from a model of the system being observed. Monitoring is typically used to detect failures by detecting differences between the observation stream and predictions from a model of the healthy system [24]. System identification is the process of combining a partially-specified model with observations from a system to converge on a more accurate and precisely specified model [23]. Traditional approaches to

monitoring and system identification deal with incomplete knowledge of the system being observed by attempting to select precise models that are close approximations to the unknown true system.

In the qualitative framework, by contrast, the attempt is to select SQDE models that cover sets of precisely-specified models and behaviors. SQUID [17] uses semi-quantitative simulation to unify the quantitative observation stream with a SQDE model to derive a more precisely specified model, still guaranteed to cover all ODE models consistent with the SQDE and the observations; or to derive a contradiction, refuting an entire family of ODE models. MIMIC [10] is an approach to monitoring that tracks multiple SQDE models in parallel, proposing and doing system identification with potential fault models even before the nominal model is refuted.

4.1 SQUID: Semi-Quantitative System Identification

An SQDE model represents a hypothesis about the qualitative structure of the system being observed. The quantitative uncertainty in a given SQDE model is represented by bounds on landmark values, static envelopes around monotonic functions, and dynamic envelopes around predicted behaviors. Information from the observation stream can be used to shrink each of these types of uncertainty.

SQUID [17] first segments the observation stream into qualitatively distinct regions of monotonic change, called trends, separated by critical points (Figure 5(a)). Then it uses MSQUID [18], a specialized neural-net-based method for estimating monotonic functions and bounding envelope functions covering the observed data points out to a specified confidence level (Figure 5(b)).


Since the observations and the predictions are now in the same bounds-and-envelopes representation, they can be combined to either refine (Figure 5(c)) or refute (Figure 5(d)) the current hypothesized model. Refinements to the dynamic envelopes predicted by the SQDE model can then be propagated back to the landmark bounds and the static envelopes around monotonic functions, so the SQDE model will be able to make more precise predictions for future cases. More precise predictions are useful for many purposes, of course, but in particular they make the model easier to refute, so that more subtle contradictions between observation and prediction can be detected in the future.

4.2 MIMIC: Monitoring with Semi-Quantitative Models


Starting with an SQDE representing the nominal (“healthy”) state of a system, SQUID can be used to monitor the system by using the information in the observation stream to progressively refine the uncertainty in the model (Figure 6). If the observation stream refutes the nominal model, then fault diagnosis is required.

However, if the system is complex and there is significant uncertainty in the SQDE, then indications of possible faults may have appeared in the observation stream well before the nominal model could actually be refuted. It is well known that operator failure in complex dynamic systems (e.g. the Three Mile Island nuclear plant failure) often occurs due to operator fixation on a single hypothesized model of the system that is only refuted after it is too late to fix a developing problem [25].


The MIMIC approach to monitoring [10, 27] tracks multiple hypotheses in parallel, each expressed as an SQDE model. Any desired features in the observation stream can be used to trigger fault hypotheses, launching additional trackers to run in parallel, even before the nominal model is refuted. Multiple active trackers and their predictions can be analyzed to select observations or




(a)



(b)



(c)



(d)

Figure 5: (a) Classify observation points in monotonic or steady regions. (b) MSQUID infers monotonic envelopes from data. (c) Refine the model by intersecting with the data. (d) Refute the model when the data fails to match.

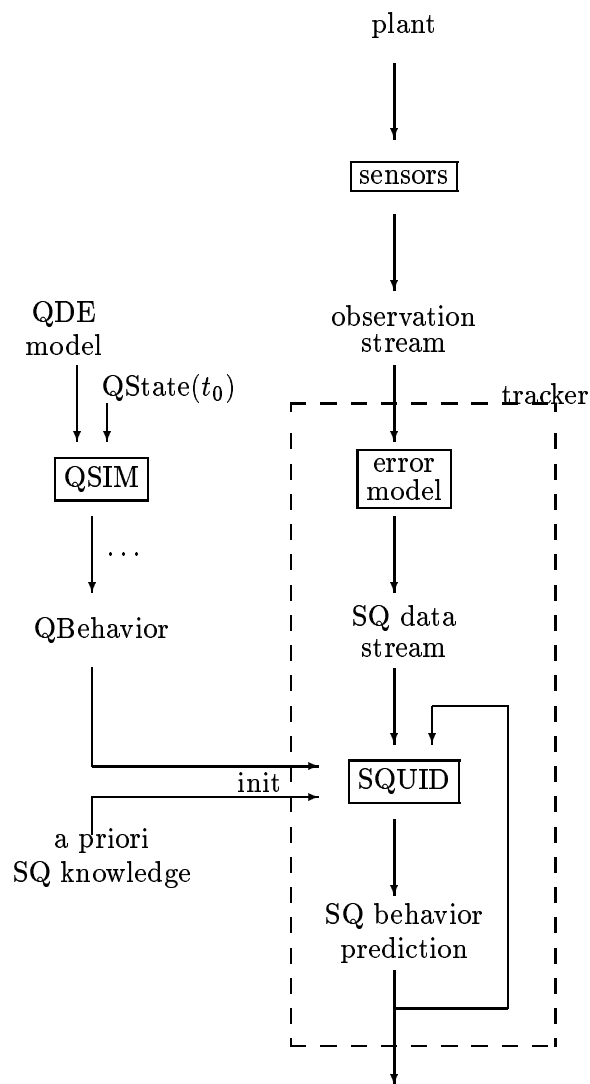


Figure 6: A *tracker* embodies a hypothesis, expressed as an SQDE model, and uses SQUID to refine (or refute) the model using information in the observation stream.

plan experiments for differential diagnosis. Since each SQDE model predicts the future behavior of the system based on its hypothesis, the information is available for a cost-benefit analysis of the costs and risks of bad outcomes versus the costs of diagnosis or repair.

4.3 Research Problems: Tractability

The benefit of the MIMIC approach to monitoring is that the qualitative representation makes it possible for a finite set of SQDE models to cover an infinite set of ODE models and initial conditions. The tractability of MIMIC monitoring, however, does depend on the tractability of qualitative simulation (section 2.1) to avoid running multiple trackers for functionally identical models; and on the ability to do semi-quantitative reasoning with probability distributions (section 3.5), to make it possible to identify low-probability models even when they are not absolutely refuted. The tractability of MIMIC also depends on the balance between the process that triggers new hypotheses based on features in the observation stream and the process that dismisses hypotheses based on inconsistency or low probability.

5 Building Qualitative Models

When a scientist or engineer builds an ODE model in the traditional way, he or she does a considerable amount of qualitative reasoning about the relevant constraints to include in the model, and the types of behaviors that might result from the interactions among those constraints. Since numerical simulation of an ODE model predicts only a single behavior, the scientist may experiment with a variety of parameter values, looking for qualitative changes in the resulting behavior.

Qualitative simulation can be regarded as a computational tool for performing an important subset of the qualitative reasoning the scientist or engineer does on the way to a numerically precise ODE model. It predicts the set of all possible behaviors of an incompletely specified SQDE model, automatically exploring a space of possible precise values for symbols denoting constants (i.e., landmark values) and functions (i.e., monotonic function constraints).

The other aspect of qualitative reasoning done by the human model-builder is to decide which constraints to include in a model. Several methods have been explored to automate this aspect of model-building [28]. These methods are as applicable to the creation of quantitative ODE models as they are to qualitative models.

A model embodies the assumption that the objects and relationships included within the model are the only ones that need to be considered for the purpose for which the model was created. That is, by its very nature, a model embodies a *Closed World Assumption*: anything not explicitly included in the model is excluded.

The component-connection approach builds models of systems that divide into modules that interact through well-defined interfaces [21, ch. 13]. Electrical circuits and fluid-flow systems are physical domains where component-connection models are often appropriate, but the same framework can be useful in more abstract settings such as compartmental modeling in biology [16, 15] and system dynamics in economics [13]. Component-connection models are also useful for model-based diagnosis [7, 9], where the goal is to account for the misbehavior of a device by identifying the smallest (or most probable) set of components whose failure can explain the observations.

Compositional modeling [12, 21, ch. 14] selects a relevant set of *model fragments* from a knowledge base of mechanisms, decides when the Closed World Assumption is appropriate, and transforms

a set of model fragments into a model. An important research task is to reason effectively about the appropriate selection of modeling assumptions to focus the compositional model-building process.

6 Conclusion

Qualitative simulation of partially specified models complements numerical simulation of completely specified models. Qualitative and semi-quantitative differential equation models make it possible to express natural types of incomplete knowledge, and qualitative and semi-quantitative simulation make it possible to predict all possible behaviors consistent with the available knowledge. Monitoring, system identification, design and verification can all benefit from the ability of a finite set of qualitative models to cover the predictions of an infinite set of precise models. Automated model building methods can utilize libraries of domain-specific model fragments to create ODE, QDE and SQDE models as needed.

References

- [1] Daniel Berleant and Benjamin Kuipers. Qualitative and quantitative simulation: bridging the gap. *Artificial Intelligence*, 95(2):215–255, 1997.
- [2] Girish Bhat, Rance Cleaveland, and Orna Grumberg. Efficient on-the-fly model checking for CTL*. In *Proc. Conf. on Logic in Computer Science (LICS-95)*, 1995.
- [3] Giorgio Brajnik and Daniel J. Clancy. Focusing qualitative simulation using temporal logic: theoretical foundations. *Annals of Mathematics and Artificial Intelligence*, 22:59–86, 1998.
- [4] D. J. Clancy and B. J. Kuipers. Static and dynamic abstraction solves the problem of chatter in qualitative simulation. In *Proc. 14th National Conf. on Artificial Intelligence (AAAI-97)*. AAAI/MIT Press, 1997.
- [5] D. J. Clancy and B. J. Kuipers. Qualitative simulation as a temporally-extended constraint satisfaction problem. In *Proc. 15th National Conf. on Artificial Intelligence (AAAI-98)*. AAAI/MIT Press, 1998.
- [6] P. Dague, P. Deves, and O. Raiman. Troubleshooting: When modeling is the trouble. In *Proc. 6th National Conf. on Artificial Intelligence (AAAI-87)*, pages 600–605, San Mateo, CA, 1987. Morgan Kaufmann.
- [7] R. Davis. Diagnostic reasoning based on structure and behavior. *Artificial Intelligence*, 24:347–410, 1984.
- [8] J. de Kleer and J.S. Brown. A qualitative physics based on confluences. *Artificial Intelligence*, 24:7–83, 1984.
- [9] Johan de Kleer and Brian C. Williams. Diagnosing multiple faults. *Artificial Intelligence*, 32:97–130, 1987.
- [10] Daniel Dvorak and Benjamin Kuipers. Process monitoring and diagnosis: a model-based approach. *IEEE Expert*, 6(3):67–74, June 1991.
- [11] E. Allen Emerson. Temporal and modal logic. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, pages 995–1072. Elsevier Science Pub. B. V./MIT Press, 1990.
- [12] Kenneth Forbus. Qualitative process theory. *Artificial Intelligence*, 24:85–168, 1984.
- [13] J. Forrester. *Urban Dynamics*. MIT Press, Cambridge, MA, 1969.
- [14] Patrick Hayes. The second naive physics manifesto. In J.R. Hobbs and R.C. Moore, editors, *Formal Theories of the Commonsense World*, pages 1–36. Ablex Publishing Corporation, Norwood, NJ, 1985.
- [15] L. Ironi and M. Stefanelli. A framework for building and simulating qualitative models of compartmental systems. *Computer Methods and Programs in Biomedicine*, 42:233–254, 1994.
- [16] John A. Jacquez. *Compartmental Analysis in Biology and Medicine*. University of Michigan Press, Ann Arbor, second edition, 1985.
- [17] H. Kay, B. Rinner, and B. J. Kuipers. Semi-quantitative system identification. *Artificial Intelligence*, 119:103–140, 2000.
- [18] H. Kay and L. H. Ungar. Estimating monotonic functions and their bounds. *American Institute of Chemical Engineering (AIChE) Journal*, 46(12):2426–2434, 2000.
- [19] Herbert Kay. SQSIM: a simulator for imprecise ODE models. *Computers and Chemical Engineering*, 23(1):27–46, October 1998.
- [20] B. Kuipers. Qualitative simulation. *Artificial Intelligence*, 29:289–338, 1986.

- [21] B. J. Kuipers. *Qualitative Reasoning: Modeling and Simulation with Incomplete Knowledge*. MIT Press, Cambridge, MA, 1994.
- [22] B. J. Kuipers and K. Åström. The composition and validation of heterogeneous control laws. *Automatica*, 30(2):233–249, 1994.
- [23] Lennart Ljung. *System Identification: Theory for the User*. Prentice-Hall, Englewood Cliffs, NJ, 1987.
- [24] Ron Patton, Paul Frank, and Robert Clark. *Fault Diagnosis in Dynamic Systems: Theory and Applications*. Prentice Hall, New York, 1989.
- [25] Charles Perrow. *Normal Accidents: Living With High-Risk Technologies*. Basic Books, New York, 1984.
- [26] O. Raiman. Order of magnitude reasoning. *Artificial Intelligence*, 51:11–38, 1991.
- [27] Bernhard Rinner and Benjamin Kuipers. Monitoring piecewise continuous behaviors by refining semi-quantitative trackers. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI-99)*, Stockholm, Sweden, 1999.
- [28] Cis Schut and Bert Bredeweg. An overview of approaches to qualitative model construction. *Knowledge Engineering Review*, 11(1):1–25, March 1996.
- [29] Benjamin Shults and Benjamin Kuipers. Proving properties of continuous systems: qualitative simulation and temporal logic. *Artificial Intelligence*, 92:91–129, 1997.
- [30] Peter Struss. Mathematical aspects of qualitative reasoning. *Int. J. Artificial Intelligence in Engineering*, 3(3):156–169, 1988.
- [31] Louise Travé-Massuyès and Robert Milne. Application oriented qualitative reasoning. *Knowledge Engineering Review*, 10(2):181–204, 1995.
- [32] Brian Williams. Doing time: putting qualitative reasoning on firmer ground. In *Proc. 5th National Conf. on Artificial Intelligence (AAAI-86)*, pages 105–112, San Mateo, CA, 1986. Morgan Kaufmann.