

ARTINT 995

Qualitative simulation: then and now *

Benjamin J. Kuipers

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

Qualitative reasoning about physical systems has become one of the most active and productive areas in AI in recent years. While there are many different kinds of qualitative reasoning, the central role is played by *qualitative simulation*: prediction of the possible behaviors consistent with incomplete knowledge of the structure of physical system.

In the retrospective [8] on my 1984 paper, “Commonsense reasoning about causality: deriving behavior from structure”, I describe the framework for qualitative reasoning that has motivated this work, and the applications that have come out of that framework. That paper [5] includes the conjecture that the structural and behavioral representations for qualitative simulation could be rigorously shown to be abstractions of ordinary differential equations and their solutions. My 1986 paper, “Qualitative simulation”, established that conjecture and legitimized the term *qualitative differential equation* or QDE. It also presented the clear and efficient QSIM algorithm. In this retrospective, I describe aspects of the body of work on qualitative simulation that has developed from there.

1. Background

Three motivating insights led to the development of the QSIM algorithm. First, the design for the QDE representation for qualitative models, presented

Correspondence to: Computer Sciences Department, University of Texas at Austin, Austin TX 78712, USA. E-mail: kuipers@cs.utexas.edu.

*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-8905494 and IRI-8904454, by NASA grant NAG 2-507, and by the Texas Advanced Research Program under grant no. 003658-175.




Fig. 1. All models are abstractions of the world. Qualitative models are related to ordinary differential equations, but are more expressive of incomplete knowledge.

in the 1984 *Artificial Intelligence Journal* paper, had been inspired both by observations of human experts (Kuipers and Kassirer [13]) and by the language of ordinary differential equations, so it was natural to ask how the mathematical similarity could be proved to be a true abstraction relation (Fig. 1).

Second, my previous ENV algorithm and its implementation had become unwieldy both in theory and in practice, so it was time to redesign the algorithm and reimplement the simulator. My research assistant at the time, Christopher Eliot, suggested the approach of proposing transitions and filtering inconsistent combinations. We were both inspired by David Waltz' compelling animation of his constraint filtering algorithm in the classic MIT AI Lab film, "The Eye of the Robot". Rather later, it became clear that the QSIM algorithm was almost a textbook application of Mackworth's node, arc, and path consistency algorithms for constraint satisfaction, but our inspiration came from the film, not the theory.

The third insight came from attempting to do a qualitative simulation, by hand, of the simple undamped oscillator: $x'' = -M^+(x)$. At the end of the first complete cycle, the simulation branches three ways according to whether the oscillation was increasing, steady, or decreasing, although only the steady case is consistent with this equation. After some confusion, it became clear that this apparent bug in the algorithm reflected a fundamental and revealing limitation in the mathematics of qualitative simulation.

2. Abstraction, soundness, and incompleteness

Once the abstraction relations from ODEs to QDEs, and from continuously differentiable functions to qualitative behaviors, are carefully defined,¹ the mathematical results are relatively straightforward.

We can view an ordinary differential equation solver as a theorem-prover for theorems of a special form:

$$\text{DiffEqs} \vdash \text{ODE} \wedge \text{State}(t_0) \rightarrow \text{Beh}. \quad (1)$$

A qualitative simulation algorithm can also be viewed as a special-purpose theorem-prover:

$$\text{QSIM} \vdash \text{QDE} \wedge \text{QState}(t_0) \rightarrow \text{or}(\text{QBeh}_1, \dots, \text{QBeh}_n). \quad (2)$$

The soundness theorem says that when QSIM proves a theorem of form (2), it is true: that is, for any ODE described by the QDE, and $\text{State}(t_0)$ described by $\text{QState}(t_0)$, the solution Beh to the ODE is described by one of the qualitative behaviors, $\text{QBeh}_1, \dots, \text{QBeh}_n$. The constraint filtering algorithm makes the proof very simple: all possible real transitions from one qualitative state to the next are proposed, and only impossible ones are filtered out, so all the real ones must remain.

The incompleteness theorem says that some qualitative behaviors in the disjunct may be *spurious*: that is, not abstracting any real solution to an ODE corresponding to the QDE. In the simple oscillator example, the increasing and decreasing behaviors are spurious. This situation is properly considered incompleteness, since QSIM has failed to prove the stronger theorem with fewer disjuncts.

3. Progress in qualitative simulation

The constraint filtering architecture of the QSIM algorithm lends itself to natural extension with a set of global filters² on complete qualitative states or behaviors. The goal of each filter is to make certain consequences of the qualitative description explicit, and to detect inconsistencies so the behavior

¹A QDE is a description of a set of ODEs, with two essential abstractions. First, a *quantity space* is an abstraction of the real number line to an ordered set of *landmark values*, symbolic names for qualitatively significant values. Second, the arithmetic and differential constraints in the ODE are augmented by a *monotonic function* constraint describing a fixed but unknown function in terms of its direction of change.

²These filters are “global” in the sense that they apply to complete qualitative state descriptions, not just to individual assignments of values to variables, or tuples of assignments. The filters also vary according to whether their scope is an individual state or an entire behavior.

can be filtered out. The creation of a suitable set of global filters has been an ongoing and productive line of research.

3.1. State-based filters

- *Infinite values and infinite times.* The QSIM abstraction is defined over the extended number line, so $+\infty$ and $-\infty$ are represented by landmark values in each quantity space. There are useful constraints on the possible combinations of finite and infinite values, times, and rates of change (Kuipers [6]).
- *Higher-order derivatives.* Certain unconstrained or “chattering” sets of qualitative behaviors can be pruned by deriving and applying expressions for higher-order derivatives of key variables in the QDE (Kuipers and Chiu [11], Kuipers et al. [12], building on earlier work by Williams [19] and by de Kleer and Bobrow [2]). The derivation may require additional assumptions about the behavior of unspecified monotonic functions.
- *Ignoring direction of change.* Chattering behaviors can also be collapsed into a single description without an additional assumption by ignoring certain qualitative features, at the cost of additional possible spurious behaviors (Kuipers and Chiu [11], Kuipers et al. [12]).

3.2. History-based filters

- *Non-intersection of trajectories in qualitative phase space.* The solution to a differential equation can be viewed as a trajectory in phase space. These trajectories cannot intersect themselves or each other at finite times. Methods for testing for self-intersection, applicable even under the qualitative behavior description, were developed independently by Lee and Kuipers [14] and by Struss [18].
- *Kinetic energy theorem.* Under very general circumstances, a QDE can be viewed as representing motion in response to a force, which in turn can be decomposed into a conservative and a non-conservative component. Then, over any segment of behavior, the change in kinetic energy of the system must be equal to the sum of conservative and non-conservative work. This equation can often be evaluated qualitatively, and eliminates an important source of spurious behaviors (Fouché and Kuipers [3]).

3.3. Quantitative constraints

Methods for adding quantitative information to qualitative behaviors can be used both to exploit additional *a priori* knowledge, and to interpret quantitative observations by unifying them with a qualitative behavior.

- *Q2: bounds on landmarks and monotonic functions.* A qualitative behavior predicted by QSIM can serve as a framework for representing quantitative information by annotating landmark values with real intervals and monotonic function constraints with real-valued functions serving as bounding envelopes. The quantitative bounds can be propagated across constraints to derive tighter bounds, or to detect a contradiction and filter out the behavior (Kuipers and Berleant [9,10]).
- *Q3: adaptive discretization.* The quantitative precision of the prediction from Q2 is drastically limited by the coarse grain-size of the qualitative behavior. The grain-size can be adaptively refined by inserting additional qualitative states, to converge to a real-valued function as uncertainty goes to zero (Berleant and Kuipers [1]).

3.4. Operating region transitions

A given QDE model has a region of applicability. When a behavior is about to cross the boundary of that region, simulation stops within the current region. If a model exists for the region on the other side of the boundary, a transition is created to a new state defined with respect to that model. In QSIM this is done by an explicit transition function that specifies which values are inherited, asserted, or inferred in the new state. The two states linked by a region transition are both considered time-points, and refer to the “same” point in time. This has two different interpretations:

- The two regions may have different constraints, but have identical descriptions of the state on their shared boundary. Therefore, the two transition states are alternate descriptions of the same physical state in time.
- The transition may represent the two sides of a “discontinuous” change: really a continuous but fast process whose extent is abstracted to zero for the purposes of the current model (Nishida and Doshita [16]).

This distinction can be illustrated with two models of a bouncing ball (provided as examples with the distributed version of QSIM): one models the bounce as a continuous transition between a gravity model and a spring model, and the other models the bounce as a discontinuous reflection of the velocity when the ball strikes the floor.

3.5. Time-scale abstraction

Time-scale abstraction allows us to decompose a model of a complex system into a hierarchy of simpler models of the system operating at different time-scales. A process in the midst of a time-scale hierarchy can view slower processes as constant and faster processes as acting instantaneously. That is, it can take a quasi-equilibrium view of the faster process, and abstract its

behavior to a monotonic function (Kuipers [7]). There is much more to be done in this area, particularly drawing on traditional mathematical work on time-scales.

4. Open problems

There are many important open problems that naturally arise from the QDE representation and the QSIM algorithm. I list three interesting ones.

- *Qualitative phase portrait analysis.* Derive the set of all possible qualitative phase portraits of a given second-order QDE.

A phase portrait captures the set of all possible behaviors of a dynamical system, for all initial states. It thus fills the same role as the “total envisionment”, but with a more expressive language for qualitative features. Sacks [17] and Yip [20] have demonstrated important results in the intelligent control of numerical experiments to map phase portraits of dynamical systems, given numerically specific equations.

It is known that the phase portraits of all second-order systems can be described in terms of a simple qualitative language (Hirsch and Smale [4]). Preliminary experiments suggest that these terms can be inferred from intelligently guided *qualitative* simulation of a QDE model. This project would require automated algebraic analysis of the QDE to search for Lyapunov functions and other derived qualitative properties of the QDE. The resulting qualitative phase portrait would depend on fewer assumptions and thus have wider applicability than the corresponding numerically-based description.

- *Automatic formulation of numerical problems.* Use the tree of qualitative behaviors to formulate problems for a numerical equation-solver, for example an optimizer.

Each predicted qualitative behavior represents a qualitatively equivalent set of continuous behaviors. The qualitative behavior description can be mapped naturally onto a set of equations over landmark values and other symbolic terms (Kuipers and Berleant [10]). It should be possible to transform that set of equations into the appropriate forms for input to a variety of numerical equation-solving algorithms. For example, an optimizer could be used to find the numerical values for certain landmarks that optimize the value of some objective function.

The set of continuous behaviors corresponding to a single qualitative behavior provides useful assumptions to the equation-solver. Where there are several qualitative behaviors, the numerical solutions found along each branch can be combined, in the case of an optimizer by searching for the maximum value.

- *Completeness*. Is the problem of spurious behaviors a fundamental limitation of qualitative reasoning, or is the QDE language sufficiently limited that sound *and* complete qualitative simulation is possible?

On the one hand, recently developed methods are capable of detecting and filtering out many of the previously-troublesome sources of spurious behaviors. On the other hand, algebraic equivalence to zero is recursively unsolvable for a language rich enough to include the transcendental functions (Moses [15]). Either outcome to this question would be of considerable interest to the QR community.

References

- [1] D. Berleant and B.J. Kuipers, Qualitative-numeric simulation with Q3, in: B. Faltings and P. Struss, eds., *Recent Advances in Qualitative Physics* (MIT Press, Cambridge, MA, 1992) 3–16.
- [2] J. de Kleer and D.G. Bobrow, Qualitative reasoning with higher-order derivatives, in: *Proceedings AAAI-84*, Austin, TX (1984).
- [3] P. Fouché and B.J. Kuipers, Reasoning about energy in qualitative simulation, *IEEE Trans. Syst. Man Cybern.* **22** (1) (1992) 47–63.
- [4] M.W. Hirsch and S. Smale, *Differential Equations, Dynamical Systems, and Linear Algebra* (Academic Press, New York, 1974).
- [5] B.J. Kuipers, Commonsense reasoning about causality: deriving behavior from structure, *Artif. Intell.* **24** (1984) 169–203.
- [6] B.J. Kuipers, Qualitative simulation, *Artif. Intell.* **29** (1986) 289–338.
- [7] B.J. Kuipers, Abstraction by time-scale in qualitative simulation, in: *Proceedings AAAI-87*, Seattle, WA (1987).
- [8] B.J. Kuipers, Reasoning with qualitative models, *Artif. Intell.* **59** (1993) 125–132 (this volume).
- [9] B.J. Kuipers and D. Berleant, Using incomplete quantitative knowledge in qualitative reasoning, in: *Proceedings AAAI-88*, St. Paul, MN (1988).
- [10] B.J. Kuipers and D. Berleant, A smooth integration of incomplete quantitative knowledge into qualitative simulation, Artificial Intelligence Lab TR AI 90-122, University of Texas at Austin (1990).
- [11] B.J. Kuipers and C. Chiu, Taming intractable branching in qualitative simulation, in: *Proceedings IJCAI-87*, Milan, Italy (1987).
- [12] B.J. Kuipers, C. Chiu, D.T. Dalle Molle and D.R. Throop, Higher-order derivative constraints in qualitative simulation, *Artif. Intell.* **51** (1991) 343–379.
- [13] B.J. Kuipers and J.P. Kassirer, Causal reasoning in medicine: analysis of a protocol, *Cogn. Sci.* **8** (1984) 363–385.
- [14] W.W. Lee and B.J. Kuipers, Non-intersection of trajectories in qualitative phase space: a global constraint for qualitative simulation, in: *Proceedings AAAI-88*, St. Paul, MN (1988).
- [15] J. Moses, Algebraic simplification: a guide for the perplexed, *Commun. ACM* **14** (8) (1971) 527–537.
- [16] T. Nishida and S. Doshita, Reasoning about discontinuous change, in: *Proceedings AAAI-87*, Seattle, WA (1987).
- [17] E.P. Sacks, Automatic qualitative analysis of dynamic systems using piecewise linear approximations, *Artif. Intell.* **41** (1990) 313–364.
- [18] P. Struss, Global filters for qualitative behaviors, in: *Proceedings AAAI-88*, St. Paul, MN (1988).

- [19] B.C. Williams, The use of continuity in a qualitative physics, in: *Proceedings AAAI-84*, Austin, TX (1984) 350–354.
- [20] K.M.-K. Yip, *KAM: A System for Intelligently Guiding Numerical Experimentation by Computer* (MIT Press, Cambridge, MA, 1991).