[SKFO89] **Skorstad, G. and K. Forbus**: "Qualitative and Quantitative Reasoning about Thermodynamics," In Collected Papers, Qualitative Physics Workshop, 1989

[WEDE90] **Weld, D. and J. de Kleer, eds.**: *Qualitative Reasoning about Physical Systems*, Morgan Kaufmann, California, 1990

[WILL86] **Williams, B.**: "Doing Time: Putting Qualitative Reasoning on Firmer Ground,"*AAAI-86*, 1986

[DAVI87] **Davis, E.**: "Constraint Propagation with Interval Labels," *Artificial Intelligence* **32**: 281-331, 1987

[DEKL75] **de Kleer, J.**: "Qualitative and Quantitative Knowledge in Classical Mechanics," Technical report AI-TR-352. MIT Artificial Intelligence Laboratory, 1975

[DEKL77] **de Kleer, J.**: "Multiple Representations of Knowledge in a Mechanics Problem-Solver," *IJCAI-77*, 1977

[FORB84] **Forbus, K.**: "Qualitative Process Theory," *Artificial Intelligence* **24**: 85-168, 1984

[FORB86] **Forbus, K.**: "Interpreting Measurements of Physical Systems," *AAAI-86*, 1986

[FOFA90] **Forbus, K. and B. Falkenhainer**: "Self-Explanatory Simulations: An integration of qualitative and quantitative knowledge," *AAAI-90*, 1990

[GEAR71] **Gear, C.**: *Numerical Initial Value Problems in Ordinary Differential Equations*, Prentice-Hall, Englewood Cliffs, New Jersey, 1971

[KUIP84] **Kuipers, B.**: "Commonsense Reasoning about Causality: Deriving Behavior from Structure," *Artificial Intelligence* **24**: 169-204, 1984

[KUIP86] **Kuipers, B.**: "Qualitative Simulation," *Artificial Intelligence* **29**: 289-338, 1986

[KUBE88] **Kuipers, B. and D. Berleant**: "Using Incomplete Quantitative Knowledge in Qualitative Reasoning," *AAAI-88*, 1988

[MOOR79] **Moore, R.**: *Methods and Applications of Interval Arithmetic*, SIAM Studies in Applied Mathematics. SIAM, Philadelphia, 1979

[PFTV86] **Press, W., B. Flannery, S. Teukolsky and W. Vetterling**: *Numerical Recipes*, Cambridge University Press, 1986

[SACK89] **Sacks, E.**: "Qualitative Analysis of Continuous Dynamic Systems by Intelligent Numeric Experimentation," Collected Papers, Qualitative Physics Workshop, 1989

[SIMM83] **Simmons, R.**: "The Use of Qualitative and Quantitative Simulations," *AAAI-83*, 1983

[SIMM83] **Simmons, R.**: "A Theory of Debugging Plans and Interpretations," *AAAI-89*, 1983

task by allowing intervals to be given as values of the input parameters, and generating the plausible behaviors.

A similar application in engineering relates to necessary tolerances of various model parameters. Simply simulating many times with different combinations of plausible model parameter values, as in Monte Carlo simulation, is likely to be tedious, perhaps impractical, and could also miss possible and potentially serious (mis)behaviors.

This sort of tolerance analysis could also be done automatically. Given a model with precisely specified parameter values, those values could be "blurred" into intervals by a preprocessor, then *Q3* could simulate to find the plausible behaviors.

Thus we expect an exciting potential for practical application of qualitative-numeric simulation.

Acknowledgements. We would like to acknowledge useful comments on the work reported here by Paul Fishwick, Pierre Fouché, David Dalle Molle, Raman Rajagopalan, Claire Rollwage, and Dan Weld.

References

[ALHE83] **Alefeld, G. and J. Herzberger**: *Introduction to Interval Computations*, Academic Press, 1983

[BERL89] **Berleant, D.**: "A Unification of Numerical and Qualitative Model Simulation," *Proceedings of the 1989 Workshop on Model Based Reasoning*, sponsored by AAAI-89, 1989

[BERL90] **Berleant, D.**: "Constraint Propagation by Waltz Filtering is Sound," *Workshop Notes, Constraint Directed Reasoning Workshop* (p.2), sponsored by AAAI-90, 1990

[BEKU90] **Berleant, D. and B. Kuipers**: "Qualitative and Numerical Simulation: Bridging the Gap," Submitted to *Artificial Intelligence* (special issue on qualitative simulation), 1990

[BOBR85] **Bobrow, D., ed.**: *Qualitative Reasoning about Physical Systems*, MIT Press, Cambridge, Massachusetts, (reprint of *Artificial Intelligence*, **24**, 1984

[BRAN77] **Brandt, A.**: "Multi-Level Adaptive Solutions to Boundary-Value Problems," *Mathematics of Computation* **31**: 333-390, 1977

[COOP89] **Cooper, M.**: "An optimal k-consistency algorithm," *Artificial Intelligence* **41**(1) (November 1989) 89-95

5 Conclusions

Previous work on combining qualitative simulation with quantitative information has produced ways of mediating between the qualitative and the quantitative domains. Qualitative-numeric simulations are conceptually elegant, unifying the concepts of numerical simulation and qualitative simulation. *Q3*, an implemented qualitative-numeric simulator language, has the following properties.

•Qualitative-numeric simulation provides a conceptual bridge between two seemingly distinct approaches to simulation: numerical simulation and qualitative simulation.

•Full-spectrum data expressibility makes explicit a conceptual link between qualitative and numerical representations of value. Intervals, of course, are only one full-spectrum representation.

Qualitative-numeric, interval-based simulations also have some derivative properties, including:

•Partial quantitative information is sufficient to run the simulation.

•Different qualitative behaviors are retained if consistent with the initial conditions, and represented explicitly (unlike numerical simulators).

Finally, *Q3* has other useful properties:

•Transitions among the continuous segments of piece-wise continuous simulations are generated automatically.

•Step sizes need not be constant, and time points need not be precise numbers.

•Qualitatively landmarks are named and given numerical bounds.

•Inferred bounds are guaranteed correct (even with e.g. non-linear models) because Waltz filtering is correct and *Q3* utilizes Waltz filtering of interval labels.

6 Future Work: Applications

One intent of qualitative-numeric simulation is to help move qualitative simulation toward real world problems.

Take an example from the field of finance. A common problem is determining the feasibility of a goal given certain initial assumptions such as interest rates. Unfortunately the exact values of such input parameters often cannot be known ahead of time with desired precision. Thus, there is the necessity of simply trying numerical simulations with various combinations of plausible input condition values, to find out what different behaviors are possible or to get an idea of how feasible the desired system behavior (such as making a profit) might be. This is not only tedious, it may be impossible from a practical viewpoint to try all desired combinations of plausible values for several or many input parameters. Qualitative-numeric, full-spectrum simulation will help in such a

Figure 5: **a)** In the left hand column, an object is dropped from about 8 feet. In **b)**, in the right column, two extra time points were inserted into the simulation of **a**, leading to much better inferences about the values of `V5` and `T1`.

instantiations ([Berleant [BERL89]).

A theoretical convergence result (Berleant and Kuipers [BEKU90]) shows that for many simulations, intervals can be narrowed all the way to numbers in the limit as step size approaches 0.

Figure 3: A qualitative-numeric simulation of the on/off air conditioning system. Initial conditions are specified as individual numbers, rather than intervals as in figure 1. Each initial condition that was specified there as an interval is initialized here by the number at the middle of that interval, thus initializing `TI9, TI13, ON` and some constants not shown such as Resistance to Heat Flow.

Figure 4: A model variable of figure 3, plotted to scale.

the simulation. Therefore when there are more states there are more constraint

lead to worse conclusions, follows from the *implausibility* condition[6] on constraints used to Waltz filter interval labels: numbers that are implausible as values of a constraint argument must not become plausible when other intervals are narrowed, where the constraint determines whether a number is plausible or not. The implausibility condition is intuitive, easy to satisfy and all the usual algebraic operations satisfy it.

Let C be any constraint and Q be an argument of C. Define *adjacent* such that Q is *adjacent* to A if A is an argument of C and $A \neq Q$. C satisfies the *implausibility* condition if numbers which C disallows Q from having as a quantitative value, and therefore are outside the current interval associated with Q, cannot become allowed when intervals of adjacent arguments are narrowed. In other words, a constraint used for Waltz filtering must never be defined to require widening the interval of any argument due to narrowing any other argument. Thus narrowing any initializing intervals[7] can lead to narrowing their adjacent intervals, but not to widening them.

Any intervals that are narrowed by the action of some constraint due to initial conditions can likewise cause narrowing, but not widening, of their own adjacent intervals. This narrowing ripples outward, and can eventually reach all intervals in the constraint network. Since intervals are only narrowed, never widened, computational stability is assured. Unfortunately, cycles occur when intervals that were narrowed once before become narrowed again and again. This can lead to non-termination. Luckily non-termination is not hard to prevent, although halting computation means not finding the narrowest possible intervals (Davis [DAVI87]).

Figure 3 shows the inferences that can be made when the initial conditions to the qualitative-numeric simulation of figure 2 are replaced by single numbers that are the midpoints of the intervals of the initial conditions of figure 2. Comparing the two figures, it is clear that the inferred intervals tend to be markedly narrower when the initial conditions are single numbers. Most of the inferred values are also numbers, the narrowest possible intervals. Figure 4 shows one of the model variables of figure 3 plotted using a constant scale.

Convergence and time step size. Most numerical simulation techniques, such as the Runge-Kutta method, give more accurate results as the step size decreases. *Q3* behaves likewise, as illustrated in figure 5. It is not surprising that inserting states into a simulation improves the resulting inferences, because the total number of constraints that *Q3* uses to do constraint propagation increases when states are added in: each constraint template that is part of the model definition is instantiated for the values of the model variables in each state in

[6] The implausibility condition is actually required of constraints used to do *Waltz filtering* with *any* kind of label (Berleant [BERL90]). *Q3* uses Waltz filtering with interval labels to make its inferences about the ranges of values of model variables.

[7] Placing any initial condition on the model constitutes narrowing some interval from $(-\infty, \infty)$ to some subrange of $(-\infty, \infty)$.

in which to insert some value for it. For example, if time `T1` = $[1, 5]$ and time `T2` = $[2, 6]$, a state inserted between them cannot specify a value for time without adding new information. In that case, another variable will have its variable specified, and the new state will have a time of $[1, 6]$ which might be narrowed as the inference algorithm proceeds.[4]

4.2 Correctness of the inferred ranges.

By *correct* we mean that any real mechanism satisfying the structure and initial conditions of the model will exhibit values within the intervals inferred by *Q3*. *Q3* is correct because Waltz filtering with interval labels is correct (Davis [DAVI87]; Cooper [COOP89]). [5]

4.3 Convergence

By *convergence* we mean the property that a simulated curve becomes a better and better approximation to the actual curve, the more work we do, until at some limit (e.g. infinitesimal time step size) the prediction coincides exactly with the actual curve.

Commonly used numerical simulation methods, such as the Runge-Kutta technique, have been proven to converge as the time step size approaches zero [GEAR71]. Naturally we want to know if *Q3* behaves similarly. There are two aspects to showing convergence in *Q3*. One is showing that, like the Runge-Kutta technique, convergence occurs as the step size approaches zero, assuming we are given precise initial conditions. The other is showing that inferred intervals get narrower as the initial conditions become more precise, so that we can be confident that increased certainty about initial conditions leads to better final results. This property is known as *stability*. We now look at both of these.

Inferences improve as initial conditions become more precise. Showing that narrower initial intervals can lead to better conclusions, and do not

[4] If there are no gaps anywhere in the simulation in which to insert a new state, it would be necessary to use the technique of *case splitting* (Berleant and Kuipers [BEKU90].

[5] The correctness of numerical inferences assumes that arithmetic operations on the machine are correct. In practice, they are not; round-off error is a subject in numerical analysis. *Q3* maintains correctness in practice as well as principle by subtracting a small amount from the low bound of every interval and adding a small amount to the high bound before storing it. As long as this epsilon is greater than the round-off error introduced by the machine, correctness is maintained. Any decent machine is reliable to a few decimal places, and an epsilon of 1 part in 10^6 has been found sufficient.

Since Common Lisp supports rational numbers, *Q3* can also represent interval endpoints as rationals, avoiding the need for using an epsilon. However this slows down computation significantly.

Figure 2: Qualitative-numeric simulation of the air conditioned dwelling, using the initialization of figure 1. Transitions between the `air-conditioner-ON` and the `air-conditioner-OFF` model initializations were generated qualitatively. (Plots are not to scale. Named values have equidistant labels on the axes even when their numerical differences are not equidistant, distorting the plots but minimizing the label overlap of figure 4.)

symbols erased. However the simulator that produced figure 1 does not quite generalize standard numerical simulation because the granularity of the `TIME` axis discretization is too coarse for quality numerical inferences and cannot be made finer, as it is determined solely by qualitative state transitions.

Thus, viewed numerically, figure 1 is too fuzzy as the values are weaker than necessary. We now look at a solution to this problem.

4.1 Inserting new states between existing states.

In standard numerical simulation techniques, reducing the step size tends to produce better results. The same is true for qualitative-numeric simulation. For an existing simulation, e.g. that of figure 1, the steps may be made smaller by inserting new states between existing ones. The results of this are shown in figures 2 and 5. In figure 5b, new states with specific time values were inserted. Comparing 5a and 5b, the result is dramatically improved quantitative inferences about model behavior. The improvement was not so dramatic in figure 2 compared to figure 1, but that figure illustrates another benefit of state insertion: quantitative and qualitative information about the system's intermediate states is inferred and represented explicitly. This insertion process is a type of adaptive discretization which we call *step-size refinement.* Other forms of adaptive discretization include adaptive step-size control (Press et al. [PFTV86]; [GEAR71]) and multigrid methods (Brandt [BRAN77]).

In *Q3*, inserted states explicate information that was implicitly already present. In figure 5a, $T1 > 0.4$, so inserting states for times 0.2 and 0.4 is justified. An inserted state, however, can only specify the numerical value of one model variable; specifying the values of more than one variable is making an assertion about the model state that wasn't already implicitly there.[3] However, numerical information about the other model variables in the inserted state is usually inferable from model constraints and surrounding states, using the inference method of Kuipers and Berleant [KUBE88]. When inference is completed, the intervals bounding the values of many of the model variables in the various states of the simulation will have been narrowed. This constitutes improved knowledge about the behavior of the model. *Ability to infer more about behavior is the reason for inserting new states.* For a very simple example described in much greater detail see Berleant [BERL89]. The improvement can be dramatic, as in figure 5b compared with 5a.

Although an inserted state can only numerically describe a value for one model variable when it is inserted, any variable is potentially eligible, not just time. Figure 2 contains several states where another variable was used; these states specified `Inside Temperature` values of $79.5, 80.01 (printed as 80.0), 80.5, 81.0$ and 81.5. The only requirement for a variable to be used is that there be a gap

[3] Specifying more than one value in a newly inserted state is valid if the values form a *measurement.* Then the co-occuring values are justified externally. *Q3* can insert states expressing measurements.

Figure 1: The air conditione dwelling. Constants are not plotted (e.g., `Heat Flow Resistance` remains constant at $R6 = [17.1, 20.9]$). With partial quantitative information, six behaviors can be ruled out (marked with an `X`), and quantitative conclusions, in the form of intervals, are made about values of model variables in the remaining behavior.

also depends on the air conditioner cooling capacity. Normally the temperature inside declines, until the thermostat transitions to `OFF` and the cycle repeats.

Figure 1 shows the normally operating system in which the inside temperature gets warmer until the air conditioner turns on, then the air gets cooler until it turns off, then the cycle repeats. The simulation also found several other behaviors. One example is, the indoor temperature rises asymptotically toward the outside temperature, never getting high enough for the air conditioner to turn on.

By adding quantitative information, even incomplete quantitative information, quantitative inferences may be made about various qualitative values in a simulation (figure 1). These inferences are done by constraint propagation among model variable value((Kuipers and Berleant [KUBE88]). For example, one of the air conditioner constraints is `((mult R HFin dTemp))`, which declares that `mult`iplying the `R`esistance of the dwelling to heat flow in or out, by the `H`eat `F`low `in` rate occurring despite that resistance, equals `dtemp`, the temperature difference between inside and outside. One of its instantiations in the behavior shown in figure 1 is $R6 * IF29 = IO29$. Such constraints, determined by the model and behavior, are used to filter intervals containing the possible quantitative values of the qualitative values constrained by the instantiation. Constraint propagation with interval labels was examined by Davis [DAVI87]. When doing the propagation, each value is expressed as an *interval arithmetic*[2] expression in terms of the other values of the constraint. This expression is evaluated using the rules of interval arithmetic (e.g. Moore [MOOR79; Alefeld and Herzberger [ALHE83]). The result is intersected with the current label, and the intersection assigned as the new label. Davis [DAVI87] argues that constraint propagation with interval labels is correct. For more rigorous proofs see e.g. Cooper [COOP89] and Berleant [BERL90].

Quantitative inferences may lead to higher level conclusions such as ruling out behaviors and making predictions about values (Kuipers and Berleant [KUBE88], also illustrated in figure 1.

4 Qualitative-numeric Simulation with State Insertion.

Recall that a qualitative-numeric simulation generalizes qualitative and numerical simulation. Figure 1 illustrates a precursor of qualitative-numeric simulation. Removing the numerical annotations leaves a qualitative simulation of a behavior, and removing all qualitative information produces a rather weak numerical simulation-like trace, which would look like figure 1 with the non-numerical

[2]Intervals are a *full spectrum* representation, because they can represent along a continuous spectrum from qualitative values to numbers. For example, a qualitative value of "+" corresponds to $(0, \infty)$. Gradually improving the specificity about the value might for example move us through $[5, \infty]$, $[6, \infty]$, $[6, 9]$ and $[7, 9]$, finally ending at the fully specified value $[8, 8]$.

Qualitative models represent *classes* of systems in contrast to the complete system specification typical of numerical simulation. Another reason for using qualitative modeling is that information precise enough for a numerical model may not be available. However, the predictions of a qualitative model are relatively weak and may lead rapidly to intractably many possible futures. Thus we would like a simulation technique which has the strengths of both qualitative and numerical simulation.

Work on adding numerical information to qualitative models and simulations is reported by de Kleer [DEKL75;DEKL77], Simmons [SIMM83], Williams [WILL86], Forbus [FORB86], Kuipers and Berleant [KUBE88], Sacks [SACK89], Skorstad and Forbus [SKFO89], Forbus and Falkenhainer [FOFA90], and others. These works succeeded in integrating quantitative and qualitative information to a significant degree. In most cases, however, qualitative and quantitative knowledge was maintained in distinct representations.

The work presented here describes a system, *Q3*, that extends previous work, integrating the qualitative and the quantitative in a fine-grained way. *Q3* produces *qualitative-numeric* simulations, i.e., simulations that can be viewed as either qualitative simulations augmented with extra quantitative information, or numerical simulations augmented with extra qualitative information. One result is, *Q3* provides the capability of creating and using intermediate states not present at all in a qualitative simulation. The effect of using new intermediate states in qualitative-numeric simulation is two-fold: to improve the quality of numerical predictions by decreasing the simulation step size, and to explicitly represent predictions about previously unrepresented system states.

Q3 conforms to the philosophy that a system combining qualitative and numeric reasoning should be able to handle varying degrees of partial information. Thus *Q3* uses intervals as a flexible representation for partial information.

Let us look at a physical system we will use to help describe the research, discussing first the system, then its simulation by previously reported methods, and finally its qualitative-numeric simulation with *Q3*.

3 The Cycling Thermostat

The state of an on/off thermostat system changes continuously except for abrupt transitions when the thermostat turns `ON` or `OFF`. Such a system needs a *combined* simulation technique that handles piece-wise continuous systems. *Q3* can do this, so we chose the on/off thermostat system to illustrate the research presented here.

Intuitively, the temperature of a dwelling in hot weather rises at a rate depending on the thermal mass of the air in the dwelling, the amount by which the temperature outside exceeds that inside, and the thermal resistance of the walls. When the inside temperature gets high enough, the thermostat transitions to `ON`, and the air conditioner pumps out heat. The system path now

Qualitative-Numeric Simulation with *Q3*

Daniel Berleant and Benjamin Kuipers

1 Overview

A simulation contains a sequence of predicted states of a modeled system. A *qualitative-numeric* simulation is one containing both qualitative and quantitative state information, such that the qualitative information alone would be a qualitative simulation, and the quantitative information alone would be a numerical simulation. Qualitative-numeric simulation is thus a generalization of both qualitative simulation and numerical simulation, providing a framework for viewing historically disparate genres of simulation.

Qualitative-numeric simulation also holds promise as an applied technique. Since it is a generalization of numerical simulation it has useful properties associated with numerical simulation that qualitative simulation does not have, like numerical predictions. And as a generalization of qualitative simulation, it has useful properties of qualitative simulation not present in numerical simulation, like dealing with weakly defined models, and automatically making qualitative distinctions among device behaviors and among model variable values.

These points are demonstrated in an implemented qualitative-numeric simulator, *Q3*.[1]

2 Introduction

Simulation of models to predict behavior of real systems has a long history and much practical use. *Discrete event* simulation, used e.g. in computer engineering, simulates models whose state variables change discontinuously. Simulation of *continuous* systems also has a history of practical use, usually coming under the heading of *numerical methods*. *Combined* simulation refers to simulations containing both discrete and continuous changes in state. Artificial Intelligence has relatively recently entered the simulation field with *qualitative simulation*, e.g. de Kleer [DEKL77;DEKL84], Forbus [FORB84], Kuipers [KUIP84;KUIP86], and others (see collections in Bobrow [BOBR84] and Weld and de Kleer [WEDE90]). Qualitative simulation derives predictions of behavior from non-numeric models.

[1] This work was done in the Qualitative Reasoning Group at the Artificial Intelligence Laboratory, The University of Texas at Austin. Research in the Qualitative Reasoning Group is supported in part by NSF grants IRI-8602665, IRI-8905494 and IRI-8904454, NASA grants NAG 2-507 and NAG 9-200, and Texas Advanced Research Program grant 003658175.