In *Proceedings of the Eleventh National Conference on Artificial Intelligence (AAAI-93)*
Cambridge, MA: AAAI/MIT Press, 1993, p. 606-613.

Numerical Behavior Envelopes for Qualitative Models*

Herbert Kay and Benjamin Kuipers
Department of Computer Sciences
University of Texas at Austin
Austin, Texas 78712
bert@cs.utexas.edu and kuipers@cs.utexas.edu

Abstract

Semiquantitative models combine both qualitative and quantitative knowledge within a single semiquantitative qualitative differential equation (SQDE) representation. With current simulation methods, the quantitative knowledge is not exploited as fully as possible. This paper describes *dynamic envelopes* – a method to exploit quantitative knowledge more fully by deriving and numerically simulating an *extremal system* whose solution is guaranteed to bound all solutions of the SQDE. It is shown that such systems can be determined automatically given the SQDE and an initial condition. As model precision increases, the dynamic envelope bounds become more precise than those derived by other semiquantitative inference methods. We demonstrate the utility of our method by showing how it improves the dynamic monitoring and diagnosis of a vacuum pumpdown system.

Introduction

Many models of real systems are incompletely specified either because a precise model of the system does not exist or because the parameters of the model span some range of values. Qualitative simulation methods [de Kleer and Brown, 1984; Forbus, 1984; Kuipers, 1984; Kuipers, 1986] permit such systems to be simulated in the face of this incompleteness by transforming the system into a related system in a more abstract space of qualitative values where model imprecision can be dealt with by the rules of qualitative mathematics. Semiquantitative models [Kuipers and Berleant, 1988; Berleant and Kuipers, 1992] reduce model imprecision by adding numerical knowledge to the purely qualitative representation. Predictions from semiquantitative models are more precise (i.e., more tightly bounded), while still retaining the accuracy (i.e., all possible behaviors are found) provided by purely qualitative methods.

This paper presents a new inference method called *dynamic envelopes* that more fully exploits the semiquantitative representation than existing methods. It works by numerically simulating a set of (typically nonlinear) differential equations whose solutions are guaranteed to bound all behaviors of the semiquantitative QDE. This approach captures the benefits of both qualitative and quantitative reasoning as all possible behaviors of the system are simulated [Kuipers, 1986], and tighter numerical bounds are deduced yielding more precise predictions for each behavior. These benefits are especially important in monitoring tasks where early detection of deviations is vital.

We represent semiquantitative models as QSIM QDEs [Kuipers, 1986] augmented with envelopes for all monotonic functions and numeric ranges for all model variables. We call this representation an SQDE (for semiquantitative QDE). Our technique generates a bounding ordinary differential equation (ODE) system derived from the SQDE that is numerically simulated to yield bounds on all model variables. Note that since the ODE system is in general a non-linear vector function defined over a multidimensional state space, it has no closed-form solution and so the integration must be performed numerically. The resulting bounds on the SQDE as a function of t are called the *dynamic envelopes* for the system.

The strength of this method is apparent when compared to other semiquantitative approaches such as FuSim [Shen and Leitch, 1991] and Q2 [Kuipers and Berleant, 1988]. These simulators also use SQDE models, but produce overly conservative bounds because they use a simulation time-step determined by qualitative distinctions. To better understand this, consider simulating the second order model of the two-tank cascade in Figure 1a using Q2, an extension to QSIM [Kuipers, 1986]. Assume that the partially known monotonic function $f \in M^+$ is bounded by the functions (*static envelopes*) shown in Figure 1b. Q2 pro-

*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, IRI-8904454, and IRI-9017047, by NASA contract NCC 2-760, and by the Jet Propulsion Laboratory.

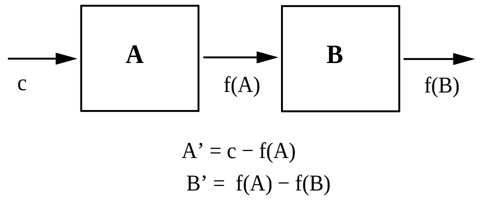

(a) System definition (c is constant).

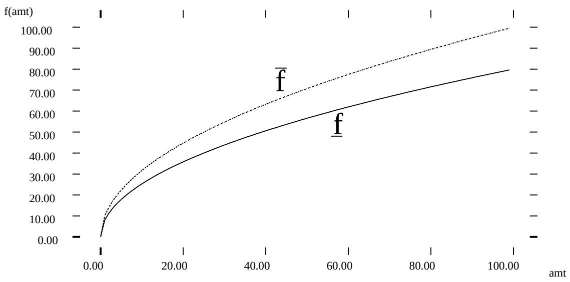

(b) Envelopes for $f(amt) \in M^+$.

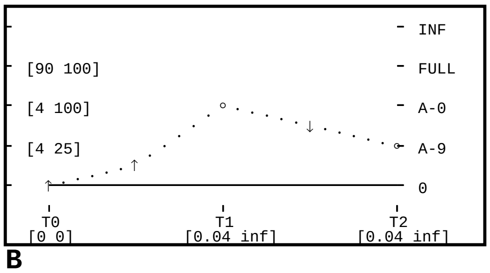

(c) Q2 behavior for $B(t)$ with initial condition A=FULL, B=0.

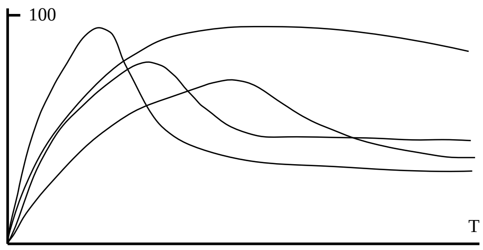

(d) Possible curves corresponding to (c).

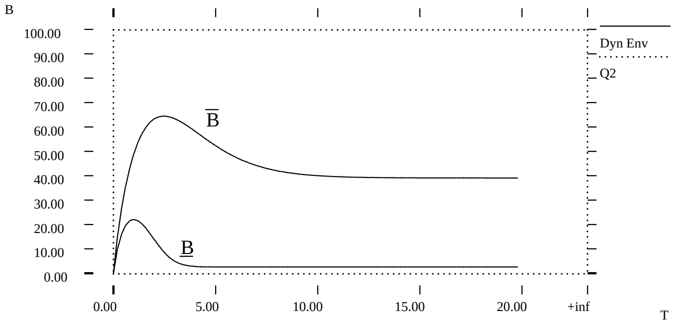

(e) Dynamic envelopes defining the lower bound $\underline{B}(t)$ and the upper bound $\overline{B}(t)$ on $B(t)$. The rectangular Q2 range prediction is superimposed. Note that the dynamic envelopes are much tighter than the Q2 bound.

Figure 1: A second order cascaded tank system and its behaviors.

duces predictions by first constructing a qualitative description of the behavior. This behavior consists of a series of time-points where one or more model variables change qualitative value, and the intervening time-intervals over which the behavior is qualitatively unchanging. It then applies a range propagation algorithm at each time-point which assigns interval values to each model variable at the time-point. These intervals represent the possible range of values that the variable could have at that time-point. Figure 1c shows the Q2 plot of the amount in tank B with the given static functional envelopes and an initial state with tank A full and tank B empty. Note that there are only three time-points in the behavior description and, due to model imprecision, the possible range for each time-point (other than T0) is infinite. There is therefore little than can be said about the possible trajectories described by the prediction. For example, *any* of the behaviors shown in Figure 1d are consistent with the prediction.

The problem is that the precision of a numerical simulation is directly related to the number and density of the time-points in the simulation and a simulator whose time-step is based solely on the qualitative distinctions cannot adequately control these quantities. The dynamic envelope method avoids this problem by using a standard numerical method (such as Runge Kutta) which chooses time-points based on local simulation error estimates. This method results in a much smaller time-step, and hence a more precise simulation. The remaining imprecision – the difference between the dynamic envelopes in Figure 1(e) – more closely reflects the incomplete knowledge in the model itself.

Dynamic Envelopes

To numerically simulate the bounds of an SQDE, bounding equations for each state variable must be generated. Our method attempts to find a set of *extremal equations* for a system. An extremal equation is a bound on the derivative of a state variable (as opposed to a bound on the *value* of the state variable). It may be either *minimal* or *maximal*.

Let $A : \mathbf{x}' = \mathbf{f}(\mathbf{x})$ be an ODE system with state vector $\mathbf{x}$. For each $x_i \in \mathbf{x}$, let $x_i' = f_i(\mathbf{x}_i)$ be the equation for the derivative of x_i where $\mathbf{x}_i \subset \mathbf{x}$ is the set of state variables that f_i depends upon. For each x_i, let $\underline{x}_i$ and $\overline{x}_i$ denote the lower and upper bounds on x_i. We will use the term y_i to refer to either $\underline{x}_i$ or $\overline{x}_i$. We say that $y_i' = g_i(\mathbf{x}_i)$ is a minimal equation for x_i if $y_i = x_i$ implies $y_i' < x_i'$ and maximal if $y_i = x_i$ implies $y_i' > x_i'$. The function g_i is called an *extremal expression* for f_i.

A set of equations is an *extremal system* for the system A if it consists of a minimal and a maximal equation for each $x_i \in \mathbf{x}$.

We can generate a set of extremal equations for any SQDE that is written as a system of equations of the form $x_i' = f(\mathbf{x}_i)$ where f is an expression composed of addition, subtraction, multiplication, division, unary minus, and arbitrary monotonic functions. The algorithm uses the functions $L(e)$ and $U(e)$ which take an expression and return the corresponding minimal or maximal expression as defined in Table 1.

e	$L(e)$	$U(e)$
c	$\underline{c}$	$\overline{c}$
x_j	$L(x_j)$	$U(x_j)$
x_i	$\beta(x_i)$	$\beta(x_i)^\dagger$
$A+B$	$L(A)+L(B)$	$U(A)+U(B)$
$A \times B$	$L(A) \times L(B)$	$U(A) \times U(B)^\ddagger$
$A-B$	$L(A)-U(B)$	$U(A)-L(B)$
$A \div B$	$L(A) \div U(B)$	$U(A) \div L(B)$
$-A$	$-U(A)$	$-L(A)$
$M^+(A)$	$\underline{M}^+(L(A))$	$\overline{M}^+(U(A))$
$M^-(A)$	$\underline{M}^-(U(A))$	$\overline{M}^-(L(A))$

Table 1: Translation table for extremal expressions of the equation $x_i' = f_i(\mathbf{x}_i)$. Let $\beta(f_i)$ be the desired bound on x_i' ($\beta = L$ or $\beta = U$). The table is applied recursively to the subexpressions of f_i. The symbol x_j is any state variable other that x_i, c is a constant, M^+ and M^- are monotonic functions, $\underline{c}$ and $\overline{c}$ return the lower or upper range values of c, $\underline{M}^*$ and $\overline{M}^*$ return the lower or upper functional envelope of the monotonic function. For state variables, $L(x)$ returns the variable $\underline{x}$ and $U(x)$ returns the variable $\overline{x}$.

The extremal equations are generated by computing for each x_i the expressions $L(f_i)$ and $U(f_i)$ using Table 1. This yields a set of $2n$ equations which represent an ODE of order $2n$ which is the extremal system for the SQDE.

Let the relation $\mathbf{R_i}$ be $\leq$ when $y_i \equiv \underline{x}_i$ and $\geq$ when $y_i \equiv \overline{x}_i$. In [Kay, 1991], the following theorem is proved:

> Let $A : \mathbf{x}' = \mathbf{f}(\mathbf{x})$ be an ODE system. Let $\alpha : \mathbf{y}' = \mathbf{g}(\mathbf{y})$ be an extremal system for A. Assume that for all i y_i $\mathbf{R_i}$ x_i at $t = 0$. Then for all t, $y_i(t)$ $\mathbf{R_i}$ $x_i(t)$.

This states that if the state of the extremal system starts on the "correct side" of the SQDE, then it will remain on that side and hence bound the solution for all time.

Once the extremal system has been found, it can be simulated by a standard numerical simulation technique such as Runge-Kutta. The complete simulation method is thus :

† The bound on x_i is the same as that for f_i regardless of whether an $L(x_i)$ or $U(x_i)$ is desired.

‡ The expressions for multiplication and division are for the case where A and B are positive. For other cases, the expressions in the table for $L(e)$ and $U(e)$ are computed differently, using information about the signs of A and B.

1. For each initial state of the SQDE, generate its extremal system.
2. Using a numerical simulator, simulate the extremal system for all initial states.

A simple example

To demonstrate the method, we apply it to the second-order model in Figure 1a. The qualitative equations of the system are

$$\begin{aligned} A' &= c - f(A) \\ B' &= f(A) - f(B) \end{aligned}$$

where $c \in (0, \infty)$ and $f \in M^+$. The semiquantitative model also includes numerical bounds on c such that $\underline{c} \leq c \leq \overline{c}$ and static envelope functions $\underline{f}$ and $\overline{f}$ such that $\underline{f} < f < \overline{f}$. The corresponding extremal system is :

$$\begin{aligned} \underline{A}' &= \underline{c} - \overline{f}(\underline{A}) \\ \underline{B}' &= \underline{f}(\underline{A}) - \overline{f}(\underline{B}) \\ \overline{A}' &= \overline{c} - \underline{f}(\overline{A}) \\ \overline{B}' &= \overline{f}(\overline{A}) - \underline{f}(\overline{B}) \end{aligned}$$

Note that in this case, the extremal system partitions into two separate systems, one for $\underline{A}$ and $\underline{B}$, the other for $\overline{A}$ and $\overline{B}$. This is *not* the case in general. Figure 1e shows the behavior produced by the dynamic envelope method that corresponds to the Q2-produced behavior shown in Figure 1c. Note that the numerical bounds are much tighter than those of Q2.

Using dynamic envelopes to infer behavior characteristics

The dynamic envelope method bases its prediction on the ability to bound the first derivatives of the system. As a result, the extremal systems are not generally members of the class of ODEs represented by the SQDE. Therefore, the dynamic envelopes do not necessarily have the same shape as the behaviors of the SQDE. This means that only "0th order" bounds are predicted. The width of the bounds will increase with increasing imprecision in the SQDE. The prediction may also become weak because the extremal system may be numerically unstable, which results in the envelopes diverging from each other with time. In such a case, the dynamic envelope bounds can eventually be worse than those from Q2.

To combat this effect, we combine the Q2 and dynamic envelope simulation methods and thus gain the benefits of both. Q2 describes behaviors as a series of time-points with intervening time intervals. It places ranges on the location of each time-point and the values of the model variables at the time-point. The time-intervals are defined simply by their adjacent time-points. We gain predictive precision by intersecting dynamic envelopes with Q2 in two ways:

- By intersecting over time-intervals, we improve the precision over the interval. The Q2 time-interval prediction is simply that each model variable is somewhere between the values that it has at the adjacent qualitative time-points. Because it uses a smaller time-step, dynamic envelopes can be more precise over such intervals (as seen in Figure 1).
- By intersecting at time-points, we potentially reduce the ranges of the model variables at the time-points. This not only improves the precision of the prediction, but may also open up gaps into which semi-quantitative time-point interpolation methods such as Q3 [Berleant and Kuipers, 1992] can insert time-points.

Note that since both the dynamic envelope method and Q2 bound all real behaviors of the model, if an intersection is empty, the behavior is refuted. Hence, dynamic envelopes can be used as a behavior filter.

The Vacuum Chamber

In this section we model a complex system, the vacuum chamber, and use the dynamic envelope simulation method to improve the response time of a monitoring system based on the MIMIC system [Dvorak and Kuipers, 1989; Dvorak, 1992].

The production of high vacuum is of great importance to semiconductor fabrication as many of the steps (such as sputtering and molecular beam epitaxy) cannot be performed if there are foreign particles in the process chamber.

Unfortunately, creating such ultra-high vacua can be expensive and time-consuming. To reach ultimate pressures of 10^{-9} Torr can take several hours[3] and something as innocuous as a fingerprint left on the chamber during servicing can cause a huge performance loss.

Because of this risk, it is important to service vacuum equipment only when there is a problem. This suggests a need for a monitoring system that can detect when the system goes out of tolerance. The normal approach to monitoring is to run the pumpdown process until the chamber reaches a steady-state pressure and then to compare this pressure to the expected value. Unfortunately, it can take several hours to reach a steady-state pressure. If the monitoring method could detect failures before the chamber reaches a steady-state pressure, the time and expense of unnecessarily running the pumpdown procedure could be avoided.

A model-based method that can track the state of the system while it is changing is one way to solve this problem. In order to construct such a system, a model of the pumpdown process must be constructed. The difficulty in modeling this process numerically is that

[3] Atmospheric pressure is 760 Torr.

there is no practical theory for the sorption [4] of gases. Therefore, any useful model must deal with uncertainties in the underlying modeling assumptions. Qualitative modeling permits reasoning with these types of uncertainties.

The pumpdown process is intuitively very simple. A chamber at atmospheric pressure initially contains some amount of gas. A pump, which can displace a certain amount of gas per unit time and pressure, removes gas from the chamber, hence lowering the pressure. For a simple vacuum pump, this process will continue until the pump reaches its cutoff pressure at which point the minimum pressure within the pump is the same as the pressure within the chamber.

For pumps that operate in the high vacuum range (between 10^{-3} and 10^{-5} Torr), there are additional effects to consider. The most significant of these is that of "outgassing" – a process where gas initially present in the walls of the chamber desorbs and thereby increases the chamber pressure.

Our model takes into account both the effects of the pump and outgassing. The system is described by the following equations :

$$
\begin{aligned}
A' &= -fl(A,B) - ptp(A) + leak(A) && (1)\\
B' &= fl(A,B) && (2)\\
fl(A,B) &= area \cdot ads(A,B) - des(B) && (3)\\
ads(A,B) &= mi(pr(A)) \cdot sf(B) && (4)\\
ptp(A) &= pr(A) \cdot speed(pr(A)) && (5)\\
leak(A) &= C_{leak} \cdot (p_{atm} - pr(A)) \cdot c_l && (6)
\end{aligned}
$$

where A is the amount of gas in the chamber and B is the amount of gas adsorbed in the chamber walls (all other terms are defined in Table 2 in the Appendix).

For a working vacuum chamber, the leak rate is zero and hence $C_{leak} = 0$. For model-based diagnosis, however, fault models of the system must also be created. By setting C_{leak} to a positive value, the above system models a chamber with a leak. The behavior of both the working and leaking models is for A to decrease until it reaches a steady state. With $C_{leak} > 0$ the steady state value of A will be higher than when $C_{leak} = 0$.

Simulation results

The two systems were augmented with envelopes for the functions $speed(p)$, $des(B)$, $mi(A)$, and $sf(B)$ and then simulated with both Q2 and the dynamic envelope method using the values described in Table 3. The resulting envelopes are shown in Figure 2 together with the corresponding Q2 range predictions. First, notice that Q2 predicts identical ranges for the normal and faulty model whereas the dynamic envelope method predicts no overlap between the two models after $t = 4$ minutes. Second, notice that the dynamic envelope prediction for the lower envelope of the normal system is less precise than Q2 prediction. This situation is not a problem since the diagnostic algorithm uses the intersection of the Q2 and dynamic envelope predictions.

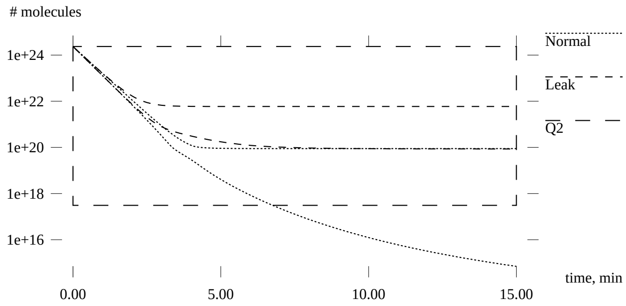

Figure 2: The predicted behaviors of the vacuum chamber A variable as a function of time for both a normal and a leaking model are shown using dynamic envelopes (the dotted and short-dashed envelopes, respectively). The behaviors of the two hypotheses are clearly distinguished after $t = 4$ minutes. For comparison, the Q2 predictions for both hypotheses are also displayed, although since Q2 is unable to disambiguate the behaviors quantitatively, the Q2 results are represented by the same prediction (long-dashed box).

Our diagnostic program is based on a simplified version of the MIMIC system [Dvorak and Kuipers, 1989]. We provided our own predefined fault models and used dynamic envelopes rather than Q2 to predict variable ranges. We then ran our system against a stream of pressure measurements (taken every minute) that simulated a gasket leak in our vacuum system. Our diagnostic system was able to detect the leak after four measurements, whereas the diagnostic system using only Q2 required nine measurements to detect the fault[5]. Further improvements are possible by recomputing the envelopes of both models after every new measurement is taken. Note that leak model envelopes are predicted based on an assumed leak size range. Because of this, when MIMIC refutes the leak model, it is really partitioning the space of possible leak sizes into three regions (those within the range, those bigger, and those smaller) with the first two regions refuted. This provides a method for converging on the precise leak size through successive partitions based on refining the leak size hypothesis.

[4] *Desorption* is the process by which gases trapped on a substance are released. The reverse process is called *adsorption*. Adsorption is different from absorption in that the gases do not dissolve into the substance; they simply "stick" to its surface.

[5] Q2 detected the fault because of a difference in the qualitative behavior of the two models that is detectable after the chamber pressure becomes constant.

Related Work

There has been considerable interest in the combination of qualitative and quantitative reasoning. This work includes the development of combined qualitative/quantitative representations (see [Williams, 1988; Kuipers and Berleant, 1988; Cheng and Stephanopoulos, 1988; Karp and Friedland, 1989; Simmons, 1986]) and the use of numerical and qualitative knowledge for process monitoring [Dvorak and Kuipers, 1989] and process planning [Fusillo and Powers, 1988; Lakshmanan and Stephanopoulos, 1988; LeClair and Abrams, 1988]. The methods and software described in [Kuipers and Berleant, 1988] and [Dvorak and Kuipers, 1989] (Q2 and MIMIC) are integral parts of this research. Recently Berleant and Kuipers have extended Q2 to provide a single representation for both qualitative and quantitative simulation [Berleant and Kuipers, 1992; Berleant, 1989]. In their method, called Q3, the range of a qualitative parameter is narrowed through an adaptive discretization technique that subdivides qualitative intervals. Q3 and the dynamic envelope method take different approaches to improving the precision of semiquantitative inference. Our experience suggests that the two methods may have complimentary strengths, so we are exploring methods for coordinating their application.

The problem of predicting behavioral bounds on uncertain systems is also addressed in control theory and ecological system simulation. Sensitivity analysis [Deif, 1986] is used to investigate the effect of small-scale perturbations to a model. Tolerance banding [Ashworth, 1982], [Lunze, 1989] is used to predict the effect of larger-scale model uncertainties. Both methods are normally restricted to linear models and hence permit uncertainty in parameter values or initial conditions only. The dynamic envelope method is not restricted by linearity assumptions and so it can also handle models with uncertain (and possibly nonlinear) functional relations. Bounding techniques that do not rely on linearity assumptions have been developed for VLSI simulation [Zukowski, 1986], however they rely on domain-specific assumptions about MOS VLSI circuits.

Interval analysis [Moore, 1979] also provides methods for simulating SQDEs by recasting standard numerical ODE solvers to work with interval arithmetic [Markov and Angelov, 1986]. In contrast, the dynamic envelope method recasts the SQDE into an ODE of higher order and uses a standard numerical ODE solver directly. An advantage of this approach is that model imprecision is separated from the error introduced by the simulator. Another benefit is that we can directly take advantage of advances in the field of numerical analysis by switching to more powerful simulators as they are developed.

This research also relates to the measurement interpretation theories ATMI [Forbus, 1986] and DATMI [DeCoste, 1990]. Both of these methods abstract a measurement stream into qualitative values and then select possible behaviors by comparing measurement segments to states in the total envisionment graph. By hypothesizing measurement errors, DATMI also manages to interpret noisy sensor data. By contrast, the dynamic envelope method augments the qualitative behavior with numerical envelopes that are guaranteed to bound any solution of the system and then compares the measurement data directly. This approach has the advantage that distinctions between models can be detected over intervals where their qualitative behaviors are identical. Furthermore, by recomputing the envelopes as new measurements are received, the bounding solutions can be further tightened. Measurement faults can also be modeled by assuming that the measurement data itself represents a range rather than a precise point.

The work on SIMGEN [Forbus and Falkenhainer, 1990] is also related to the work described in this paper. It, too, generates a standard numerical simulation by extracting the relevant information from a qualitative model. It differs in that it generates an exact numerical model based on a library of predefined functions rather than generating a bounded model expressing the inexactness of the qualitative model. Since it sacrifices accuracy for precision, it is not particularly suited to tasks such as process monitoring in which an exact numerical model cannot be found.

Conclusions

The dynamic envelope method combines qualitative and quantitative simulation so that both representations can be used in problem solving. QSIM produces all behaviors associated with a particular model, and dynamic envelopes provide detailed numerical ranges for each behavior. Because the generation of extremal systems is guided by the qualitative behaviors, the expense of needless numerical simulation is eliminated. Because the envelope systems are automatically generated from the SQDEs used by Q2, the method can be used with any existing Q2 model.

The precision of the dynamic envelope predictions depends on the precision of the SQDE. As model precision increases, dynamic envelope predictions become more precise than Q2 predictions. Even when the model is very imprecise, combining dynamic envelopes with other QSIM prediction techniques leads to improved precision.

In monitoring tasks, the dynamic envelope method improves the predictive power of SQDEs both in accuracy (meaning that fault hypotheses can be more easily eliminated) and failure detection time (meaning that there is more time to recover from failures). In cases where measurement acquisition is expensive, the increased accuracy of the predictions may allow fewer measurements to be made and errors to be detected sooner.

The ultimate goal of our research is to construct a self-calibrating monitoring and diagnosis system that can learn models directly from observations of a physical system. Part of this task involves developing a semiquantitative simulation method that produces precise predictions without excessive computation. The dynamic envelope method helps address this need by providing a new form of inference especially suited for high precision semiquantitative models.

Acknowledgments

The authors would like to thank Adam Farquhar for his comments on an earlier version of this paper.

Vacuum system terms and SQDE quantitative knowledge

Term	Definition (units)
A	amount of chamber gas (molec)
B	amount of gas in chamber walls (molec)
$area$	surface area of chamber (cm^2)
$pr(A)$	pressure exerted by A molecules [assuming fixed chamber volume v] (Torr)
$ptp(A)$	pump throughput (Torr-liters/min)
$speed(p)$	pump speed (liters/min)
$ads(A, B)$	rate : chamber gas $\rightarrow$ walls (molec/cm^2-min)
$des(B)$	rate : chamber gas $\leftarrow$ walls (molec/cm^2-min)
$fl(A, B)$	net flow of gas out of walls (molec/min)
$mi(A)$	# molecules incident on walls (molec/cm^2-min)
$sf(B)$	sticking factor : fraction of $mi(A)$ that "stick" to walls
$leak(A)$	rate : room air $\rightarrow$ chamber (molec/min)
C_{leak}	leak conductance (liters/min)
p_{atm}	atmospheric pressure (760 Torr)
c_l	constant : Torr-liters $\rightarrow$ molecs

Table 2: Definition of terms used in equations 1 through 6.

Term	Value or envelope description
A	$[2.34 \times 10^{24}, 2.34 \times 10^{24}]$ molec
B	$[1.36 \times 10^{21}, 1.50 \times 10^{21}]$ molec
$area$	$[13100, 14500]$ cm^2
C_{leak}	$[0.01, 0.001]$ liters/min
v	90 liters
$speed(p)$	M^+ piecewise linear with narrowing envelope
$des(B)$	M^+ linear with both envelopes equal
$mi(p)$	M^+ linear with both envelopes equal
$sf(B)$	M^- exponential envelope narrowing from $[2, 0.5]$ to $[0, 0]$ at $B \approx 0.3$

Table 3: Initial ranges and functional envelopes for the vacuum chamber model. These values are based on data from Duval [Duval, 1988].

References

Ashworth, M. J. 1982. *Feedback Design of Systems With Significant Uncertainty.* John Wiley and Sons, New York.

Berleant, Daniel and Kuipers, Benjamin 1992. Combined qualitative and numerical simulation with q3. In Faltings, Boi and Struss, Peter, editors 1992, *Recent Advances in Qualitative Physics.* MIT Press.

Berleant, Daniel 1989. A unification of numerical and qualitative model simulation. In *Proceedings of the Model-based Reasoning Workshop.*

Cheung, Jarvis Tat-Yin and Stephanopoulos, George 1990. Representation of process trends – part I. a formal representation framework. *Computers and Chemical Engineering* 14(4/5):495–510.

de Kleer, Johan and Brown, John Seely 1984. A qualitative physics based on confluences. *Artificial Intelligence* 24:7–83.

DeCoste, Dennis 1990. Dynamic across-time measurement interpretation. In *Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI-90).* 373–379.

Deif, Assam 1986. *Sensitivity Analysis in Linear Systems.* Springer-Verlag, Berlin.

Duval, Pierre 1988. *High Vacuum Production in the Microelectronics Industry.* Elsevier Science Publishers, Amsterdam.

Dvorak, Daniel Louis and Kuipers, Benjamin 1989. Model-based monitoring of dynamic systems. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence.* 1238–1243.

Dvorak, Daniel Louis 1992. Monitoring and diagnosis of continuous dynamic systems using semiquantitative simulation. Technical Report AI92–170, Artificial Intelligence Laboratory, University of Texas at Austin, Austin, Texas 78712.

Forbus, Kenneth D. and Falkenhainer, Brian 1990. Self-explanatory simulations: An integration of qualitative and quantitative knowledge. In *Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI-90).* 380–387.

Forbus, Kenneth 1984. Qualitative process theory. *Artificial Intelligence* 24:85–168.

Forbus, Kenneth D. 1986. Interpretating meaurements of physical systems. In *Proceedings of the Fifth National Conference on Artificial Intelligence (AAAI-86).* 113–117.

Fusillo, R. H. and Powers, G. J. 1988. Operating procedure synthesis using local models and distributed goals. *Computer and Chemical Engineering* 12(9/10):1023–1034.

Karp, Peter D. and Friedland, Peter 1989. Coordinating the use of qualitative and quantitative knowledge in declarative device modeling. In Widman, Lawrance E.; Loparo, Kenneth A.; and Nielson, Norman R., editors 1989, *Artificial Intelligence, Simulation and Modeling.* John Wiley and Sons. chapter 7.

Kay, Herbert 1991. Monitoring and diagnosis of multi-tank flows using qualitative reasoning. Master's thesis, The University of Texas at Austin.

Kuipers, Benjamin and Berleant, Daniel 1988. Using incomplete quantitative knowledge in qualitative reasoning. In *Proceedings of the Seventh National Conference on Artificial Intelligence.* 324–329.

Kuipers, Benjamin 1984. Commonsense reasoning about causality : Deriving behavior from structure. *Artificial Intelligence* 24:169–204.

Kuipers, Benjamin 1986. Qualitative simulation. *Artificial Intelligence* 29:289–338.

Lakshmanan, R. and Stephanopoulos, G. 1988. Synthesis of operating procedures for complete chemical plants – i. hierarchical, structured modelling for nonlinear planning. *Computer and Chemical Engineering* 12(9/10):985–1002.

LeClair, Steven R. and Abrams, Frances L. 1988. Qualitative process automation. In *Proceedings of the 27th National Conference on Decision and Control.* 558–563.

Lunze, Jan 1989. *Robust Multivariable Feedback Control.* Prentice Hall.

Markov, S. and Angelov, R. 1986. An interval method for systems of ode. In *Lecture Notes in Computer Science #212 – Interval Mathematics 1985.* Springer-Verlag, Berlin. 103–108.

Moore, Ramon E. 1979. *Methods and Applications of Interval Analysis.* SIAM, Philadelphia.

Shen, Qiang and Leitch, Roy 1991. Synchronized qualitative simulation in diagnosis. In *Working Papers from the Fifth International Workshop on Qualitative Reasoning about Physical Systems.* 171–185.

Simmons, Reid 1986. "commonsense" arithmetic reasoning. In *Proceedings of the Fifth National Conference on Artificial Intelligence (AAAI-86).* 118–124.

Williams, Brian C. 1988. Minima : A symbolic approach to qualitative algebraic reasoning. In *Proceedings of the Sixth National Conference on Artificial Intelligence.* 264–269.

Zukowski, Charles A. 1986. *The Bounding Approach to VLSI Circuit Simulation.* Kluwer Academic Publishers, Boston.