

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

June 25, 1992

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.

*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, by the Texas Advanced Research Program under grant no. 003658175, and by the Jet Propulsion Laboratory.

1 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 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 inferential 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 will 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] which also use SQDEs as models, but produce behaviors that have fixed bounds (i.e., are not time varying) over time intervals. To help understand this, consider the Q2 system, an extension to QSIM [Kuipers, 1986] that also operates on SQDEs. It produces predictions by applying an interval propagation algorithm at each qualitative timepoint of the simulation. The propagation method uses the Mean Value Theorem to constrain the ranges of a variable at adjacent time points.

This means that predictions over the interval between timepoints is simply the maximal range difference between the values at the timepoints.

As an example, consider the second order model of a two-tank cascade in Figure 1a. Assume that the partially known monotonic function $f \in M^+$ is bounded by the functions as shown in Figure 1b. 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. Over the time interval $[t_0, t_2]$, B falls in the range $[0, 100]$. Figure 1d shows some possible trajectories that fall in $[0, 100]$ and obey the qualitative first derivative information provided by QSIM. Ranges are propagated across time intervals by solving the equation

$$B(t_{n+1}) = B(t_n) + (t_{n+1} - t_n)B'(t^*) \quad (1)$$

where $B'(t^*)$ is in the range $[B'(t_n), B'(t_{n+1})]$. Since this range is formed from the values of B' at the two surrounding timepoints, Q2 predicts that over the interval B falls somewhere between the minimum and maximum range values of $B(t_n)$ and $B(t_{n+1})$. This forms a box in the $B-t$ plane. The dynamic envelope method shrinks this box by replacing the use of the mean value theorem with explicit integration using the equation

$$x(t_{n+1}) = x(t_n) + \int_{t_n}^{t_{n+1}} x' dt. \quad (2)$$

2 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 \bar{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 \bar{x}_i . We say that $y'_i = g_i$ 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

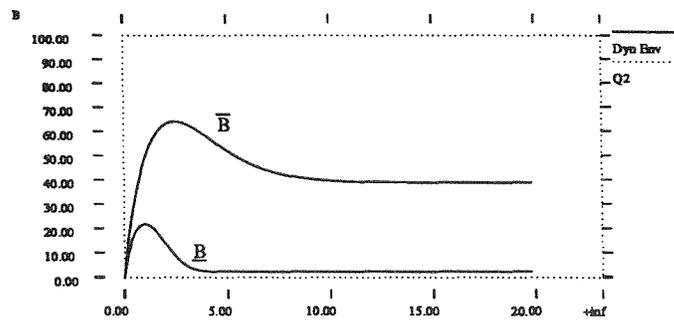
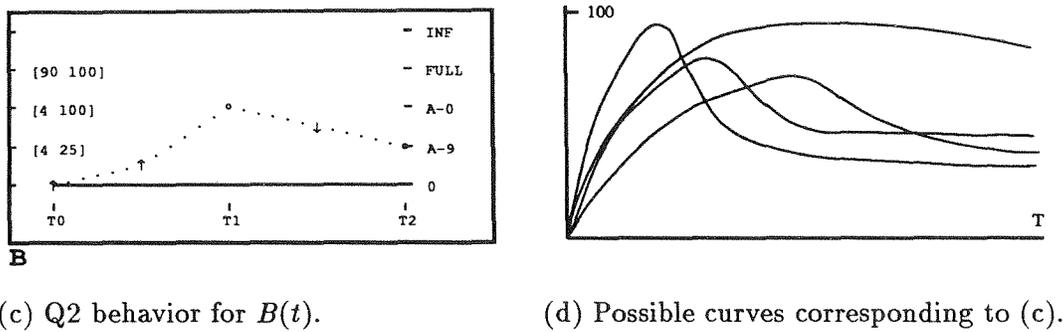
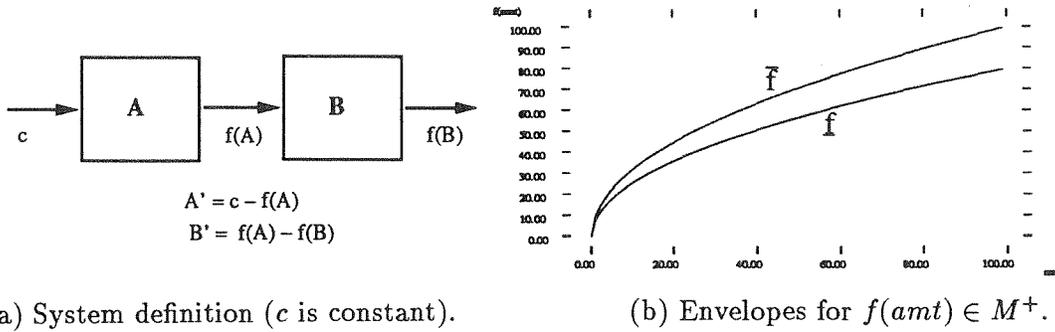


Figure 1: A second order cascaded tank system and its behaviors.

e	$L(e)$	$U(e)$
c	\underline{c}	\bar{c}
x_j	$L(x_j)$	$U(x_j)$
x_i	$\beta(x_i)$	$\beta(x_i)$
$A + B$	$L(A) + L(B)$	$U(A) + U(B)$
$A \times B$	$L(A) \times L(B)$	$U(A) \times U(B)$ ¹
$A - B$	$L(A) - U(B)$	$U(A) - L(B)$
$A \div B$	$L(A) \div U(B)$	$U(A) \div L(B)$
$-A$	$-U(A)$	$-L(A)$
$M^+(A)$	$\underline{M}^+(L(A))$	$\overline{M}^+(U(A))$
$M^-(A)$	$\underline{M}^-(U(A))$	$\overline{M}^-(L(A))$

Table 1: Translation table for extremal expressions. 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 . x_i is the state variable whose derivative is f_i , x_j is any other state variable, c is a constant, M^+ and M^- are monotonic functions, \underline{c} and \bar{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 \bar{x} .

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.

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 $y_i \equiv \bar{x}_i$. 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.

¹The expressions for multiplication and division are for the case where A and B are positive. For other cases, the expressions for $L(e)$ and $U(e)$ are computed differently, using information about the signs of A and B .

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

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.

2.1 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 $\underline{c} \leq c \leq \bar{c}$ and static envelope functions \underline{f} and \bar{f} such that $\underline{f} < f < \bar{f}$. The corresponding extremal system is :

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

Note that in this case, the extremal system partitions into two separate systems, one for \underline{A} and \underline{B} , the other for \bar{A} and \bar{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.

2.2 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. A weakly constrained SQDE can

in some cases generate an extremal system that is unstable, which limits the length of simulation time over which results can be obtained. With a very weakly constrained SQDE, the envelopes may even fall partially outside the bounds predicted by Q2.

This suggests that combining the inferences produced by dynamic envelopes and Q2 can lead to better overall predictions. With this in mind, there are three ways that dynamic envelope predictions can be used to infer SQDE behavior bounds :

- Intersect Q2 behavior predictions with the dynamic envelopes. This guarantees that the dynamic envelopes will never increase the bound width predicted by Q2.
- Combine the dynamic envelopes predictions with qualitative derivative predictions. Consider the case where the lower bound \underline{x} of the dynamic envelope for a variable x reaches a maximum \underline{x}_{max} at t_n and then turns down. If QSIM has predicted that x' is always positive then we can infer that for all $t > t_n$ the lower bound on x should be \underline{x}_{max} .
- Use the envelopes to provide new timepoint estimates to Q2. The dynamic envelope prediction at some timepoint is used to create a new state where the quantitative ranges for all model variables are taken from the simulation. This state is then spliced into the behavior and Q2 is run on the augmented behavior. Since the dynamic envelope method has provided new information about the behavior (namely, its bounds at some specific timepoint), Q2 should be able to further constrain its behavioral prediction. This technique of state splicing is also a key component to the Q3 [Berleant and Kuipers, 1991] system.

At the moment, we have implemented the first of these techniques.

3 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 hours² and something as

²Atmospheric pressure is 760 Torr.

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 there is no practical theory for the sorption³ 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 :

$$A' = -flowAB(A, B) - ptp(A) + leakrate(A) \quad (3)$$

$$B' = flowAB(A, B) \quad (4)$$

$$flowAB(A, B) = area \cdot adsorbrate(A, B) - desorbrate(B) \quad (5)$$

$$adsorbrate(A, B) = mi(pressure(A)) \cdot sf(B) \quad (6)$$

$$ptp(A) = pressure(A) \cdot speed(pressure(A)) \quad (7)$$

$$leakrate(A) = C_{leak} \cdot (p_{atm} - pressure(A)) \cdot c_i \quad (8)$$

³ *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.

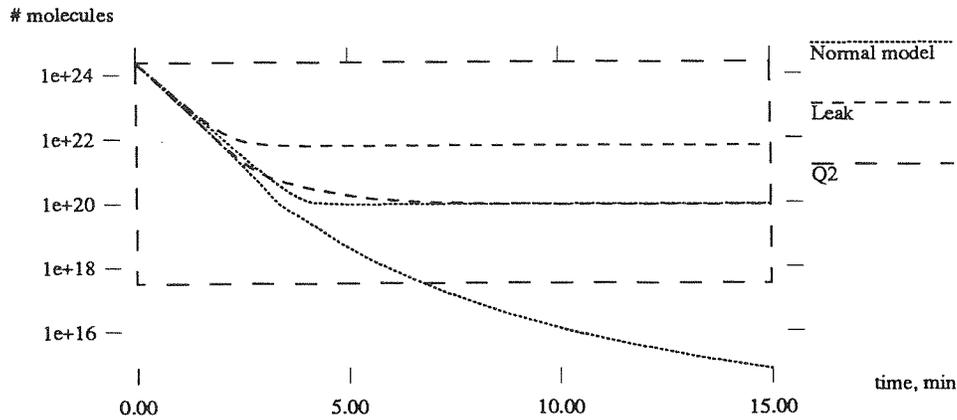


Figure 2: The predicted behaviors of the vacuum chamber A variable as a function of time for both a normal and a leaking model are shown using dynamic envelopes (the dotted and short-dashed envelopes, respectively). The behaviors of the two hypotheses are clearly distinguished after $t = 4$ minutes.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

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 $C_{leak} = 0$.

3.1 Simulation results

The two systems were augmented with envelopes for the functions $speed(p)$, $desorbate(B)$, $mi(A)$, 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.

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 fault⁴. 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 Related Work

There has been considerable interest in the combination of qualitative and quantitative reasoning. This work includes the development of combined qualitative representations (see [Williams, 1988, Kuipers and Berleant, 1988, Cheng and Stephanopoulos, 1988, Karp and Friedland, 1989]) 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, 1991, Berleant, 1989]. In their method, called Q3, the range of a qualitative parameter is narrowed through an adaptive discretization technique that subdivides qualitative intervals. As the intervals are further subdivided, the ranges shrink and the accuracy of the prediction increases. In contrast, dynamic envelopes use direct numerical

⁴Q2 detected the fault because of a difference in the qualitative behavior of the two models that is detectable after the chamber pressure becomes constant.

simulation. In particular, for an accurate set of envelope functions, Q3 can provide an arbitrary level of precision by subdividing to a fixed level of granularity. Reaching a high level of precision may require many subdivisions. Dynamic envelopes will provide a high level of precision with a fixed simulation cost, but only over fixed period of time.

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.

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. As such, it is not particularly suited to tasks such as process monitoring in which an exact numerical model cannot be found.

5 Conclusions

The dynamic envelope method combines qualitative and quantitative simulation so that both representations can be used in problem solving. It produces all behaviors

associated with a particular model, and it provides 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 model the model is very imprecise, combining dynamic envelopes with other QSIM prediction techniques leads to improved prediction 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 this research is to develop a method whereby the predictive capability of a simulation improves as more information about the underlying process is gained. The dynamic envelope method is a step in that direction.

6 Acknowledgments

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

A Vacuum system terms and SQDE quantitative knowledge

Term	Definition (units)
A	amount of gas in the chamber (# molec)
B	amount of gas dissolved in chamber walls (# molec)
$area$	surface area of the chamber (cm^2)
$pressure(A)$	the pressure corresponding to A molecules (Torr) [Assuming a fixed volume v for the chamber]
$ptp(A)$	pump throughput (Torr-liters/min)
$speed(p)$	pump speed (liters/min)
$adsorbrate(A, B)$	rate : chamber gas \rightarrow walls (# molec/ cm^2 -min)
$desorbrate(B)$	rate : chamber gas \leftarrow walls (# molec/ cm^2 -min)
$flow(A, B)$	net flow of gas out of the walls (# molec/min)
$mi(A)$	# molecules incident on chamber walls (# molec/ cm^2 -min)
$sf(B)$	sticking factor : fraction of $mi(A)$ that “stick” to walls
$leakrate(A)$	rate : room air \rightarrow chamber (# molec/min)
C_{leak}	leak conductance (liters/min)
P_{atm}	atmospheric pressure (760 Torr)
c_l	conversion constant : Torr-liters \rightarrow #. molecs

Table 2: Definition of terms used in equations 3 through 8.

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 unequal envelopes
$desorbrate(B)$	M^+ linear with both envelopes equal
$mi(p)$	M^+ linear with both envelopes equal
$sf(B)$	M^- exponential with unequal envelopes

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

References

- [Ashworth, 1982] M. J. Ashworth. *Feedback Design of Systems With Significant Uncertainty*. John Wiley and Sons, New York, 1982.
- [Berleant and Kuipers, 1991] Daniel Berleant and Benjamin Kuipers. Bridging the gap from qualitative to numerical simulation. Technical Report AI91-158, Artificial Intelligence Laboratory, University of Texas at Austin, Austin, Texas 78712, March 1991.
- [Berleant, 1989] Daniel Berleant. A unification of numerical and qualitative model simulation. In *Proceedings of the Model-based Reasoning Workshop*, 1989.
- [Cheung and Stephanopoulos, 1990] Jarvis Tat-Yin Cheung and George Stephanopoulos. Representation of process trends – part i. a formal representation framework. *Computers and Chemical Engineering*, 14(4/5):495-510, 1990.
- [de Kleer and Brown, 1984] Johan de Kleer and John Seely Brown. A qualitative physics based on confluences. *Artificial Intelligence*, 24:7-83, 1984.
- [DeCoste, 1990] Dennis DeCoste. Dynamic across-time measurement interpretation. In *Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI-90)*, pages 373-379, 1990.
- [Deif, 1986] Assam Deif. *Sensitivity Analysis in Linear Systems*. Springer-Verlag, Berlin, 1986.
- [Duval, 1988] Pierre Duval. *High Vacuum Production in the Microelectronics Industry*. Elsevier Science Publishers, Amsterdam, 1988.
- [Dvorak and Kuipers, 1989] Daniel Louis Dvorak and Benjamin Kuipers. Model-based monitoring of dynamic systems. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 1238-1243, 1989.
- [Dvorak, 1992] Daniel Louis Dvorak. 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, May 1992.
- [Forbus and Falkenhainer, 1990] Kenneth D. Forbus and Brian Falkenhainer. Self-explanatory simulations: An integration of qualitative and quantitative knowledge. In *Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI-90)*, pages 380-387, 1990.

- [Forbus, 1984] Kenneth Forbus. Qualitative process theory. *Artificial Intelligence*, 24:85–168, 1984.
- [Forbus, 1986] Kenneth D. Forbus. Interpreting measurements of physical systems. In *Proceedings of the Fifth National Conference on Artificial Intelligence (AAAI-86)*, pages 113–117, 1986.
- [Fusillo and Powers, 1988] R. H. Fusillo and G. J. Powers. Operating procedure synthesis using local models and distributed goals. *Computer and Chemical Engineering*, 12(9/10):1023–1034, 1988.
- [Karp and Friedland, 1989] Peter D. Karp and Peter Friedland. Coordinating the use of qualitative and quantitative knowledge in declarative device modeling. In Lawrence E. Widman, Kenneth A. Loparo, and Norman R. Nielson, editors, *Artificial Intelligence, Simulation and Modeling*, chapter 7. John Wiley and Sons, 1989.
- [Kay, 1991] Herbert Kay. Monitoring and diagnosis of multi-tank flows using qualitative reasoning. Master's thesis, The University of Texas at Austin, May 1991.
- [Kuipers and Berleant, 1988] Benjamin Kuipers and Daniel Berleant. Using incomplete quantitative knowledge in qualitative reasoning. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, pages 324–329, 1988.
- [Kuipers, 1984] Benjamin Kuipers. Commonsense reasoning about causality : Deriving behavior from structure. *Artificial Intelligence*, 24:169–204, 1984.
- [Kuipers, 1986] Benjamin Kuipers. Qualitative simulation. *Artificial Intelligence*, 29:289–338, September 1986.
- [Lakshmanan and Stephanopoulos, 1988] R. Lakshmanan and G. Stephanopoulos. Synthesis of operating procedures for complete chemical plants – i. hierarchical, structured modelling for nonlinear planning. *Computer and Chemical Engineering*, 12(9/10):985–1002, 1988.
- [LeClair and Abrams, 1988] Steven R. LeClair and Frances L. Abrams. Qualitative process automation. In *Proceedings of the 27th National Conference on Decision and Control*, pages 558–563, 1988.
- [Lunze, 1989] Jan Lunze. *Robust Multivariable Feedback Control*. Prentice Hall, 1989.

- [Shen and Leitch, 1991] Qiang Shen and Roy Leitch. Synchronized qualitative simulation in diagnosis. In *Working Papers from the Fifth International Workshop on Qualitative Reasoning about Physical Systems*, pages 171–185, 1991.
- [Williams, 1988] Brian C. Williams. Minima : A symbolic approach to qualitative algebraic reasoning. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, pages 264–269, 1988.