

Monitoring diseases with empirical and model generated histories

Enrico W. Coiera

Hewlett-Packard Laboratories, Filton Rd., Stoke-Gifford, Bristol BS12 6QZ
United Kingdom¹

Abstract

Diagnostic monitoring systems track disease hypotheses over time, symbolically interpreting the time-varying patient data produced by medical instrumentation. The need to track multiple interacting diseases recommends a hypothesize, test and refine reasoning architecture which incorporates a robust knowledge representation. Rule-based systems are often inadequate for this task, and deep or model based representations capable of first principles reasoning are currently favoured. However the model based approach may be too low level for many monitoring tasks. While disease interactions may present novel patterns to a monitor, usually the diseases themselves will be familiar. It is proposed that disease histories generated from pathophysiological models are at an appropriate level of abstraction for many monitoring tasks. Histories lie between disease models and rules in depth. Using the QSIM representation, results are presented for model-generated histories that define some limits of their utility in reasoning systems. In particular if the underlying system is non-linear then restrictions exist on the predictions possible with histories alone. As a consequence, we can specify when a monitoring system must switch from history to model-based representations. These results are extended to poorly modelled domains to allow reasoning about disease interactions using only empirically derived histories.

1 INTRODUCTION

The use of instrumentation to monitor a patient's clinical state has become routine in many situations. With the advent of computational techniques drawn from Artificial Intelligence, researchers are now beginning to explore the possibilities of *diagnostic patient monitoring* systems.

Diagnostic monitoring systems track disease hypotheses over time, symbolically

1. This work was carried out while the author was at the Department of Computer Science, University of New South Wales, Sydney, Australia. The Research was assisted by a New South Wales Government Medical Research Scholarship and a research award from the Medical Engineering Research Organisation. This is a revised version of a paper that originally appeared in *Artificial Intelligence in Medicine*, 2, (1990),135-147

interpreting the time-varying patient data produced by medical instrumentation. Such systems may be applied in numerous areas of medical practice, ranging from real time analysis of patient data in Intensive Care Units, to reviews of patient laboratory tests (Coiera, 1989; Dvorak, 1989; Fagan et al.,1984; Weiss, 1981). However there are wider applications beyond the medical domain. Diagnostic monitoring systems will find use in process control and other real-time diagnosis applications e.g. nuclear power plants and space stations (Doyle et al., 1989; Dvorak, 1987).

The reasoning architecture for a diagnostic monitoring system is well suited to the hypothesize and test paradigm. Observations are made of the patient, hypotheses are generated to account for the observations, and their expected behaviour over time is then compared to the observed behaviour of the patient. The result of this comparison forms the test by which the monitoring system can select amongst competing hypotheses and refine its diagnoses.

The four key elements of the hypothesize-test-refine cycle are shown in Figure 1 (Coiera, 1989):

- evoking the initial hypotheses
- matching hypotheses to subsequent observations
- revision of hypotheses based upon the match
- generating new behavioural predictions for the revised hypotheses.

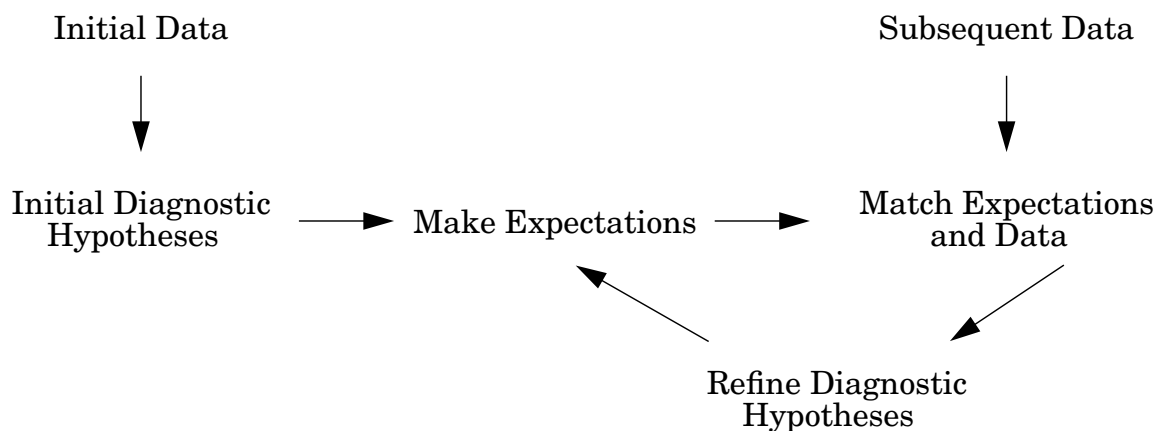


Figure 1. The Diagnostic Monitoring Cycle.

1.1 Monitoring Acid-Base Disturbances

An example drawn from simple acid-base physiology will help demonstrate how a diagnostic monitor might work. We start with a set of blood gases - pH, serum bicarbonate concentration and partial pressure of CO₂ - drawn from a patient and presented in Column 1 of Table 1. These results suggest two possible diagnoses - either an acute or a chronic metabolic alkalosis. The chronic alkalosis would be relatively long standing, while the acute alkalosis would be recent.

Table 1
Example Blood-Gas Data

	Data 1	Data 2	Difference
pH	7.49	7.43	-
$P_a\text{CO}_2$	43	48	+
HCO_3^-	35	31	-

In an attempt to differentiate between these two diagnostic hypotheses, the diagnostic monitor would make predictions outlining the expected behaviour of each over future observations. We would expect that the blood gas values would remain relatively stable for the chronic alkalosis. Compensatory changes would be expected in response to the acute alkalosis. In particular, respiratory compensation should cause the $P_a\text{CO}_2$ to rise and the pH to fall.

These predictions are now matched by the monitor to a second set of gases taken about 12 hours later (Column 2, Table 1). Qualitatively the new gases match the evolving acute alkalosis hypothesis. To retain the chronic alkalosis as a viable alternative, the monitor system must postulate that some new events have occurred over the intervening 12 hours to cause the unexpected shifts. For example, the unlikely occurrence of both a new metabolic and respiratory acidosis in tandem with the chronic alkalosis would explain the data.

Now the diagnostic monitoring system is faced with a new problem. If it wishes to keep both hypotheses active it must generate a new set of behavioural predictions for the revised disease complex of chronic metabolic alkalosis - acute respiratory acidosis - acute metabolic acidosis (Figure 2). What began as a relatively straightforward diagnostic monitoring exercise has rapidly become a difficult prediction task that must cope with interacting diseases. The problem is compounded if one needs to decompose observed disease interactions into their components for therapeutic purposes. For example, it may be necessary to track the progress of each individual disease within an interaction to determine its response to a particular therapeutic intervention.

1.2 Selecting a Knowledge Representation

A critical task in the diagnostic monitoring cycle is thus predicting the future behaviour of disease hypotheses. This is knowledge intensive and the choice of an appropriate knowledge representation becomes a key issue in system design. The representation must be sufficiently rich to capture the intricacies of knowledge in the domain. In this case it must capture the time varying behaviour of disease. Secondly the representation must be sufficiently deep to cope with the complexities of the problem solving task. With patient monitoring we not only need to track individual diseases, but must reason about complex time varying disease interactions.

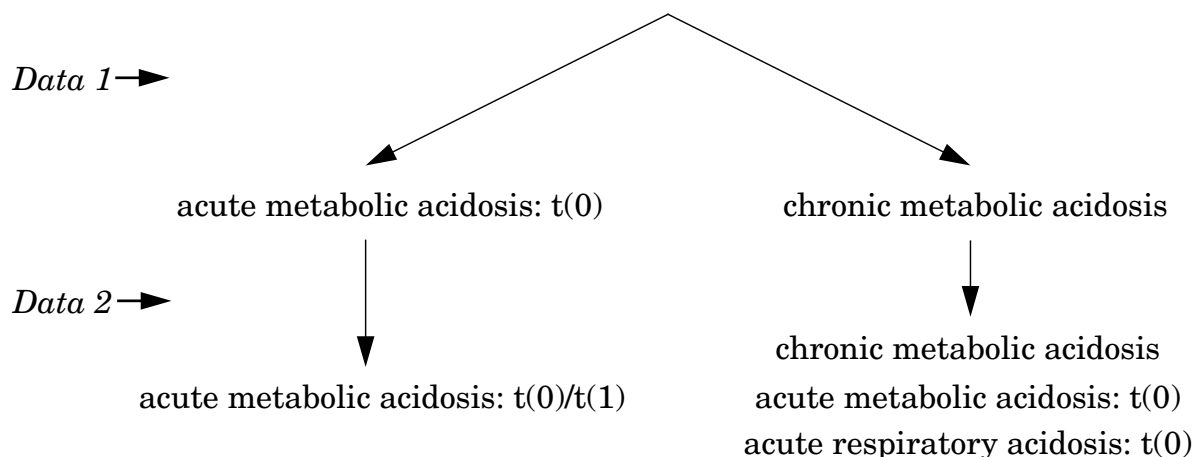


Figure 2. Evolving hypotheses from example data.

The knowledge representation we select must thus capture enough detail of the pathophysiology of disease to allow reasoning about such disease interactions.

Clearly the combinatorial demands of creating a rule for every possible disease interaction in most medical domains is enormous. Allowing for the complexity that medical intervention adds to the presentation of disease, this approach becomes highly limited. Rules cannot cope with the complexities of unpredicted events or previously unencountered combinations of events (Davis, 1987). These limitations have motivated the development of second-generation expert systems relying on some form of 'deep' knowledge (Keravnou and Washbrook, 1989). One approach is to use a model of the underlying structure and function of a domain. With such models, a reasoner may be able to resolve novel situations from first principles.

However, in complex medical domains it seems inefficient to rely completely on model-based systems. Most novel data patterns encountered by a patient monitoring system will not be the result of previously unencountered diseases, but will result from the interactions of known diseases. There is thus scope to use methods and representations that lie in between rule and model based systems, optimising between computational cost and depth of knowledge represented (Coiera, 1991). Further, pathophysiological models are not always available and the use of shallower representations can extend the functionality of monitoring systems in such poorly modelled domains

1.3 Disease Histories

One traditionally successful approach in clinical practice is to discuss observed disease states in relation to an idealised *natural history*. The natural history describes the usual temporal course of a disease in an average patient, unimpeded by other factors. The history may have been obtained empirically, being refined over successive case studies. If the disease process is well understood, the history might be derivable from a pathophysiological disease model.

The notion that the history of an event can be an effective knowledge representation has found currency within AI. Hayes (1979) first suggested that histories might be a useful representation for reasoning with physical domains at a naive level. Since then the notion under one name or another has become central to much work in Qualitative Physics (Weld and de Kleer, 1990).

Histories are an attractive representation because while they seem superficially to be a temporal pattern that could be encoded in a rule, they can be linked to deeper models of a domain. It is this ability to formally link histories with model-based representations that will prove to be advantageous. While reasoning with templates of temporal disease patterns alone has value, there are limitations to the types of reasoning that may be performed with them compared to deeper physiological models (Cohn et al., 1990). In particular, one cannot make inferences about the interaction of events recorded on separate templates. As will be shown below, provided histories are derived from a deeper model, it can be proven that for many cases they are an adequate representation for reasoning with interactions. An important corollary will be to define situations in which a monitor cannot rely on histories alone but must invoke deeper knowledge representations.

We may also wish to reason with histories because first principles reasoning is not possible. When no deep model is available, descriptive disease histories might still allow a diagnostic monitor to draw some conclusions that shallower representations cannot. This is attractive in the medical domain where knowledge about the physiological basis of many diseases is still lacking. At present, most prototypic monitoring systems have selected domains that are well understood and susceptible to modelling (e.g. Ironi, 1990, Uckun, 1992). Histories offer the opportunity to extend the range of diagnostic monitoring systems into less well mapped territory. Reasoning with histories also allows a diagnostic monitoring system to decompose complex disease interactions into their components, and thus track the behaviour of each component within an observed interaction (Coiera, 1989). Such decomposition is not possible with model-based simulations, which will generate only composite interaction histories.

History generation has been variously described as *envisioning* or *qualitative simulation* in the AI literature (de Kleer and Brown, 1984; Kuipers 1986). Programs that generate histories rely on a model of the physical world. There are numerous competing representational ontologies (Coiera, 1992) - for example the physical world can be modelled as a collection of processes in Qualitative Process Theory (Forbus, 1984), as a set of structural or functional device components (Davis, 1984; de Kleer and Brown, 1984), or a set of qualitative mathematical constraints between continuous functions (Kuipers, 1986).

The intuition we will use here is that by their very nature, histories must contain implicit information about the structure and function of the system that generated them. Thus histories should be useful in making inferences that require some deep understanding of a system. In the process of generating the history from the model however, some of that knowledge is lost, since the history captures just one of many potential behaviours derivable from the model.

The QSIM formalism will be used here to represent histories and their underlying models, and a brief introduction to it follows. A formal exposition of QSIM is presented in (Kuipers, 1986). It is not appropriate at this stage in our knowledge to

select out one history generating methodology as being the best. However QSIM is strongly based on mathematical methods and lends itself to rigorous definition and analysis. Further, the relationship between QSIM and competing representational theories like Qualitative Process Theory (Forbus, 1984) are now being explored (Coiera, 1992). This suggests the possibility of providing formal mechanisms for translating models between these representations. While QSIM's mathematical flavour can be construed as too limiting for some domains (Bratko, 1989) the qualitative constraint language has been successfully used to model quite complex physiological systems (Kuipers, 1985). These considerations were uppermost in choosing QSIM for the present work.

2 AN INTRODUCTION TO QSIM

QSIM is an algorithm for performing qualitative simulation. Its inputs are qualitative models and their initial conditions. The models are composed of qualitative differential equations (QDEs) between continuous real valued parameters. The initial conditions for these parameters are specified in qualitative terms. QSIM is thus analogous to numeric simulation methods that take a set of ordinary differential equations (ODEs) and numeric values as their input.

Several qualitative constraint types are used in QSIM to capture mathematical relationships. For some arbitrary continuous and real valued functions a , b , c we can represent simple relationships like $add(a,b,c)$ if $a + b = c$, or $deriv(a,b)$ if $a'(t) = b(t)$. More complex functional relationships can be represented with the monotonic constraint. If two functions are strictly monotonically increasing over an interval of interest, QSIM represents that relationship as $M^+(a,b)$. Many functional relationships such as polynomial and exponential relationships can be mapped onto the monotonic constraint. In medicine we often cannot specify complete functional relationship between parameters, and constraints such as M^+ or its inverse M^- can capture the incomplete qualitative specifications of such relationships. The formal definition of each constraint type is presented in (Kuipers, 1986).

2.1 An Example QSIM Model

A schematic model of the acid-base regulatory mechanism is presented in Figure 3. Briefly pulmonary regulation controls pH by controlling the excretion rate of carbon dioxide, and the kidney exerts its control on pH through manufacture and excretion rates for bicarbonate. The third regulatory component of the system is provided by acid buffers which act as reservoirs for hydrogen ions. We can model this system as a series of equilibrium relationships between carbon dioxide, hydrogen ions and bicarbonate (Narins, 1980):



A QSIM constraint model for the system is presented in Figure 4. For simplicity, we model the three control parameters as their change from normal. Thus ΔpH is positive for any pH value greater than normal.

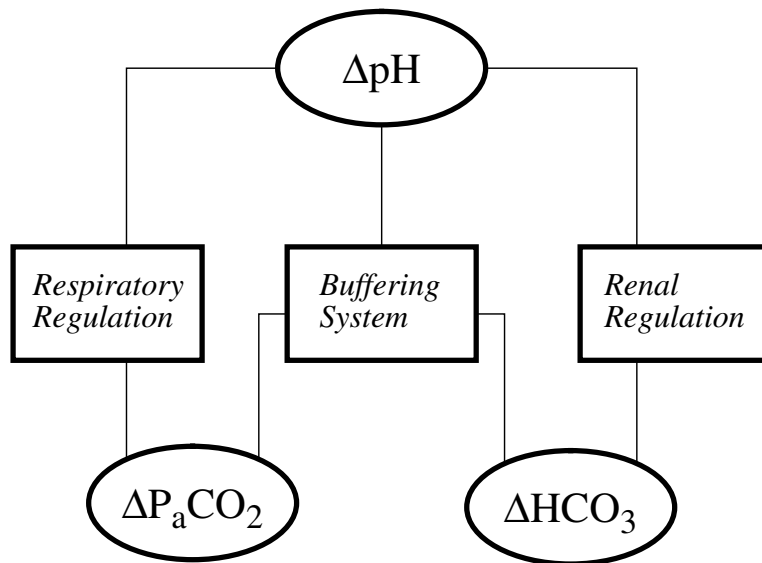


Figure 3. Schematic Model of Acid-Base Regulation.

The action of the major buffers in the equilibrium system can be modelled by the constraint $add(\Delta pH, \Delta P_a CO_2, \Delta HCO_3)$. For example, if $\Delta P_a CO_2$ is zero and ΔpH is positive, then the constraint stipulates that ΔHCO_3 must also be positive. The two derivative relationships are simple models of the renal and respiratory arms of the regulatory mechanism. Thus with $deriv(\Delta P_a CO_2, \Delta pH)$, we can capture the information that when pH is positive (a net alkalosis) then the value for $\Delta P_a CO_2$ will be increasing in an effort to return the pH towards normal.

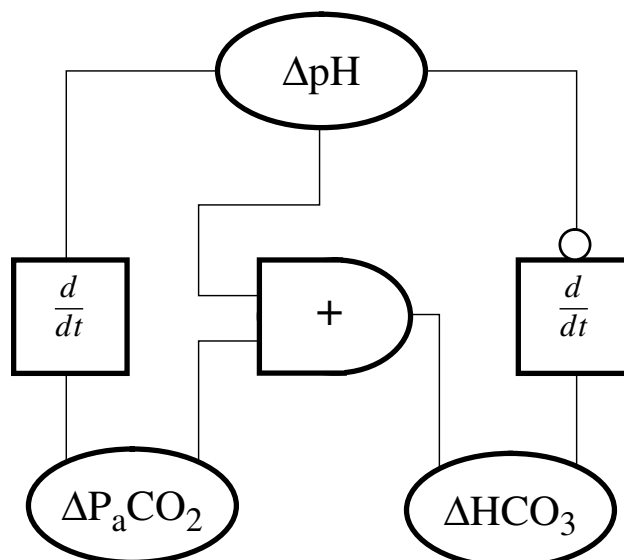


Figure 4. QSIM Model of Acid-Base Regulation.

2.2 Histories are solutions to QDEs in QSIM

When presented with a model and initial conditions, QSIM simulates the resulting time-varying behaviour which we have called a history. Just as QSIM's constraints are qualitative analogues of ordinary differential equations, the histories QSIM generates are direct but qualitative analogues of the solution for a set of ODEs.

QSIM histories are a sequence of qualitative state descriptions. Each state in the sequence is distinct and each system parameter is assigned a qualitative value for each state. Parameter values are either distinguished points in a range of interest such as end points or stationary points, or intervals bounded by such points. For example, if the acid base model in Figure 4 is initialised with an above normal carbon dioxide partial pressure, that $\Delta P_a\text{CO}_2$ value is assigned the qualitative state $\text{QS}(0/\infty, \text{dec})$. The first value states that the carbon dioxide value is positive, lying somewhere between zero and infinity.

Time derivatives in QSIM can either be increasing, steady or decreasing - *inc*, *std*, *dec*. The second component of the $P_a\text{CO}_2$ value *dec* thus states that the first time derivative is negative. Representing function values in this way allows QSIM to make predictions when only approximate measures are available.

The simulation algorithm at the heart of QSIM is detailed in (Kuipers, 1986) and will not be presented here. Figure 5 shows the history QSIM generates with a high initial carbon dioxide value, clinically equivalent to a respiratory acidosis. A sequence of three distinct qualitative time states is produced. At time point $t(0)$, the initial increase in $\Delta P_a\text{CO}_2$ pushed the ΔpH value lower than normal. The time interval $t(0)/t(1)$ corresponds to the beginning of renal compensation with a rise in serum bicarbonate levels. At $t(1)$ a new homeostatic equilibrium is reached, with pH now returned to normal and bicarbonate and CO_2 achieving new steady states.

The bicarbonate steady state is labelled $\text{QS}(\text{dHCO}_3(1), \text{std})$ where $\text{dHCO}_3(1)$ is the label generated by QSIM for the unspecified but positive equilibrium value. Similarly, $\text{dCO}_2(1)$ is the label QSIM attaches to the new positive equilibrium value for $\Delta P_a\text{CO}_2$ at its final state. It should be noted that this is a very simplified disease history and many subtleties of the true disease behaviour are not captured here. More complex behaviours, that for example include finer details of the buffer response, can be captured with a more detailed model. A history for metabolic acidosis that was also produced from the acid-base model is presented alongside the first in Figure 5.

3 PREDICTING DISEASE INTERACTIONS

If histories are to be an acceptable knowledge representation for a patient monitoring system, they must predict the complex behaviours that result from disease interactions. We can use the deeper model-based representation of disease behaviour from QSIM as a 'gold standard'. We will need to prove that any interaction behaviour that can be produced by QSIM can also be produced by manipulating single disease histories. Recall that we wish to use histories in situations when a monitor knows all the likely diseases it will encounter and will only have to reason about novel disease combinations.

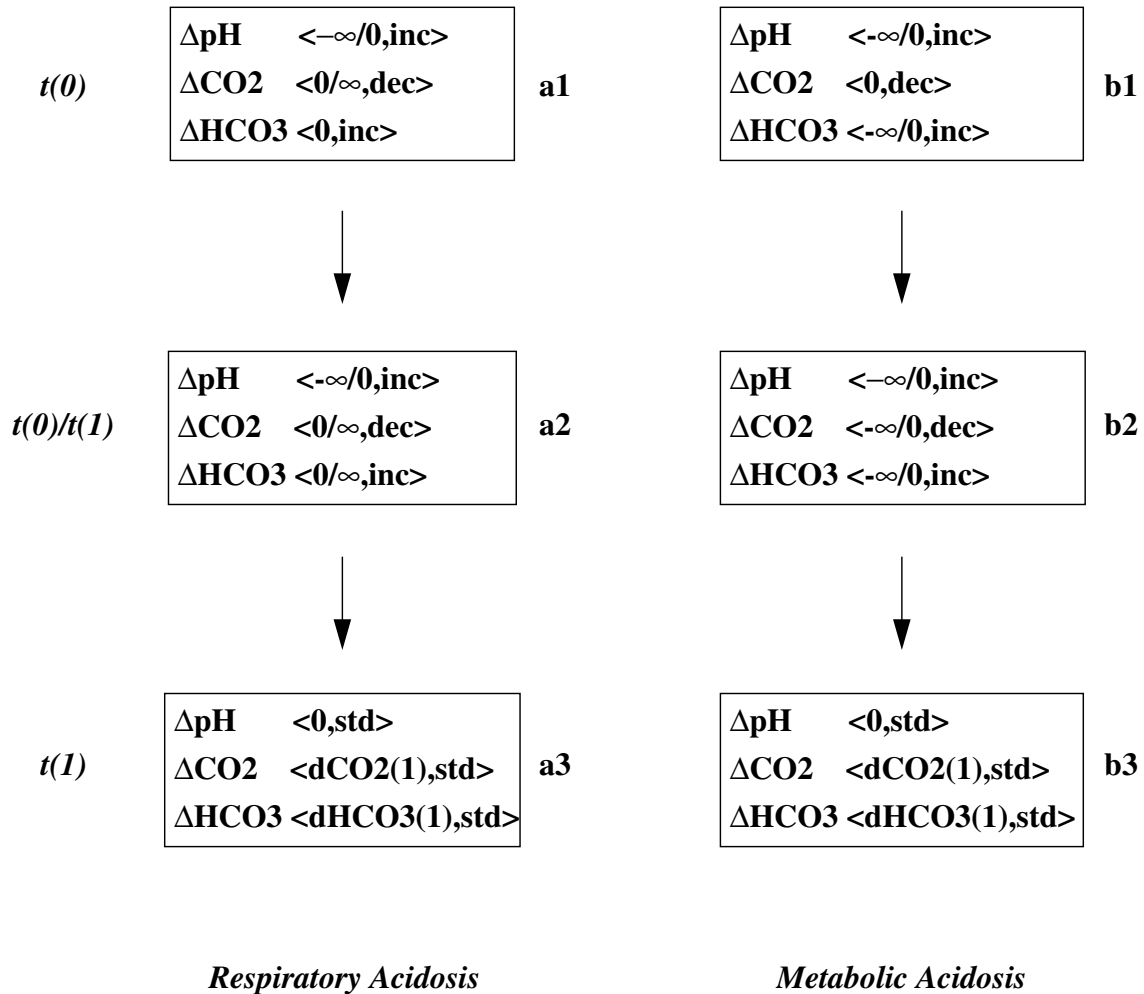


Figure 5. Simulated histories for respiratory and metabolic acidoses.

If we present QSIM with the previous acid-base model and initialise it now for two diseases - respiratory acidosis and metabolic acidosis - three possible behaviours are predicted (Figure 6). The leftmost branch of the behaviour tree shown in Figure 6 shows the unlikely scenario where disease and homeostatic forces are exactly balanced, producing blood gas parameters that return to normal values at the equilibrium state $t(1)$. The middle branch shows bicarbonate retention in response to the respiratory acidosis dominating the opposing metabolic acidosis which would keep the bicarbonate low. In the third branch at the right, the initially high $P_a\text{CO}_2$ level due to the respiratory acidosis is dominated by the drive to lower $P_a\text{CO}_2$ in response to the metabolic acidosis. Clinically, a larger variety of behaviours is possible, and again a more detailed model would reveal this.

3.1 Superposition

The principle of superposition for linear ODEs states that the sum of any two solutions for a system is also a solution to that system (Rabenstein, 1972). Recall that qualitative disease histories are direct analogues of such solutions. This result allows histories derived from a linear system model to be directly added to produce an interaction history. Thus solutions produced by the sum of any two histories from a linear QSIM model should also be a legal history for that model.

However, qualitative state addition is an underconstrained process. When two states are added, several possible states may result. Adding a qualitative state $QS_a(0/\infty, inc)$ to $QS_b(0/\infty, dec)$ produces either $QS_{a+b}(0/\infty, inc)$, $QS_{a+b}(0/\infty, dec)$, or $QS_{a+b}(0/\infty, std)$. Further such solutions could be assembled into infinitely long behaviours e.g.

$$\begin{aligned} &QS_{a+b}(0/\infty, inc, t_n), QS_{a+b}(0/\infty, std, t_{n/n+1}), \\ &QS_{a+b}(0/\infty, dec, t_{n+1}), QS_{a+b}(0/\infty, std, t_{n+1/n+2}), \\ &QS_{a+b}(0/\infty, inc, t_{n+2}), QS_{a+b}(0/\infty, std, t_{n+2/n+3})..... \end{aligned}$$

This problem has been called "chatter" (Kuipers, 1987) and results from the inherent ambiguity of the qualitative algebra.

This problem can be demonstrated in the interacting acidoses example. The three interaction disease behaviours in Figure 6 can be generated by adding the two single disease histories in Figure 5. The qualitative addition of two histories involves adding various permutations of the individual states from each history. The addition of the respiratory acidosis history composed of three distinct qualitative states labelled $a1$, $a2$ and $a3$ in Figure 5, to the three metabolic acidosis states $b1$, $b2$ and $b3$ generates 9 possible interaction states. A huge number of behaviours could be assembled from these 9 states, and if we include cycles some will be non-terminating. One assembled behaviour that does reproduce a genuine interaction is labelled down the rightmost branch in figure 6. Thus state $t(0)$ in the tree can be produced by adding states $a1$ and $b1$.

To make the qualitative addition of histories a useful procedure, constraints must be introduced into the addition that eliminate spurious behaviours, retaining only those behaviours that QSIM itself would generate. In general these constraints arise from understanding the form that histories must take, and from assumptions about the nature of the interaction that can be drawn from domain knowledge.

When two histories are added, performing a permutation of legal state combinations, a state space is created. The task of creating an interaction history becomes one of assembling a sequence of states from that space to produce the interaction history. This process can be thought of as a depth first search through the state space, selecting paths that satisfy several specific criteria.

Assume we begin with an initial interaction state produced when we add the initial states of two histories. The process of path traversal involves searching for a successor to our initial state from the state space, and iterating the process until the final interaction state is reached. If such a state is not found, the search backtracks along the assembled path until an alternative state can be found and a new path segment explored. A final state for an interaction history would result from the combination of the final states of the initial histories.

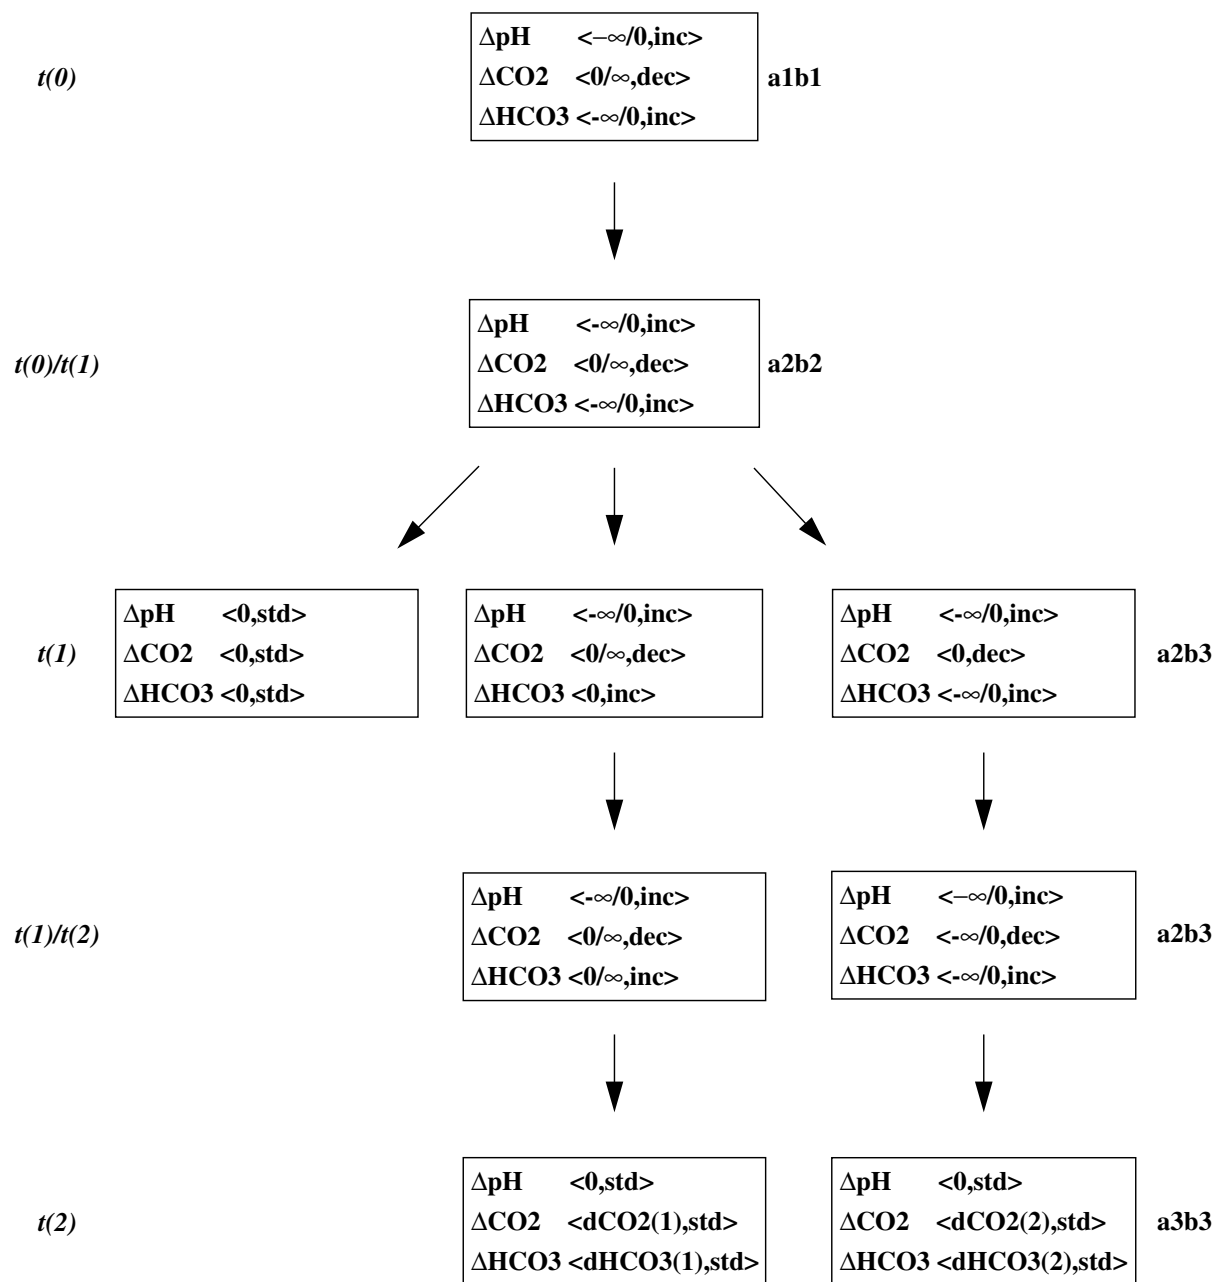


Figure 6. History tree for concurrent respiratory and metabolic acidosis.

3.1.1 Constraining the size of the State Space

The first step in the process is to generate the search space of interaction states. From the principle of superposition it is known that the qualitative addition of states from linear systems will produce legal interaction states. Solutions other than these may also be generated because of the ambiguity of the qualitative algebra. Such spurious states are a feature of QSIM and qualitative reasoning in general. Several quite powerful methods have been developed to eliminate them in

QSIM, in particular methods reliant on global filters (Kuipers and Chiu, 1987; Lee and Kuipers, 1988). A number of local methods also exist to reduce these spurious solutions. Williams (1988) in particular developed a strengthened qualitative algebra to deal with this problem.

In addition to these, several other methods can be used to reduce the number of spurious states generated in the initial addition of histories. These include:

1/ *Limiting Temporal Extent*. Rather than performing a full permutation of state additions, only those additions that preserve temporal and functional continuity need be performed. In particular, transient states taken from an individual history cannot exist more than momentarily within a sum state. Such transients, when they occur, are usually present in the initial state of a history. The addition of initial states containing transients to non-initial states from other histories is thus prohibited. Hence states $a1$ and $b1$ in Figure 5 can only be added to each other, producing the composite state $a1b1$ in Figure 6. A state $a1b2$ would imply that $a1$ could exist for more than a moment, spanning the history segment $a1b1$ to $a1b2$. However the value for ΔHCO_3 in $a1$ is $\langle 0, \text{inc} \rangle$ and this value cannot exist across an interval of time.

2/ *Relative Magnitude Assumptions*. If it is known that in absolute terms one history dominates another during an interaction, such that the sum values are largely determined by it, we can use that knowledge to restrict the initial search space. If $\text{QS}_a(0/\infty, \text{dec})$ dominates $\text{QS}_b(-\infty/0, \text{inc})$ then the result will always be $\text{QS}_{a+b}(0/\infty, \text{dec})$. A more detailed exploration of order of magnitude reasoning can be found in (Raiman, 1988; Mavrovountiotis and Stephanopoulos, 1988).

3/ *Model consistency*. The process of history addition can clearly also be constrained by the QSIM system model itself. Each sum state generated in the addition process is compared to the constraint model, and states that would violate that model are disallowed. Thus a sum state must be legal for the system that produced its component states.

3.2 Path Creation

The technique of envisioning (Forbus, 1988) generates all possible state descriptions for a system model, based upon a set of background assumptions. The state space created by history addition is analogous to that created by envisionment, and recreates part or the whole of the envisionment space.

In his Logic of Occurrence, Forbus (1987) presents a formal relationship between histories and envisionments. In particular, he describes methods for selecting legal paths from the envisionment state space, and includes methods for avoiding impossible cyclic behaviours. These methods need not be reproduced here, but are directly applicable to the task at hand. Some additional points will need to be made however. In particular, several local and global constraints can be applied during path creation.

1/ *Functional progression*. For a candidate state to be considered as a potential successor to an interaction state, it must be continuous with that state. Recall

that a characteristic of QSIM functions is that they are real valued and continuous. Thus we would disallow a path segment that introduced a discontinuity in the values of parameters in the path e.g. jumping from $QS(0, \text{std})$ to $QS(1, \text{std})$ would be disallowed. Further, continuity determines whether adjacent states in a history can have the same description. In particular, a transition from an interval to a point in time should not have the same state descriptions since a distinguished time point represents a critical point in a function's behaviour i.e. by definition a behavioural change is expected. The reverse transition, from point to interval, does however allow similar state descriptions (Kuipers, 1986).

2/ Temporal progression To prevent local cycles and also to enforce the intuition that for an interaction history to advance to a new state, at least one of the individual histories must advance to a new state, a strict progression is enforced. This means that if we added state *a1* (Figure 5) time stamped $t(0)$ and *b1* time stamped $t(0)$, then the successor chosen must result from the sum of states time stamped with either $t(0)$ or $t(0)/t(1)$. A jump to a sum state produced by a $t(1)$ state would be disallowed. Retrograde jumps are equally prohibited.

The temporal progression constraint may be further specialised to take advantage of contextual information when it is available:

i/ Synchronous progression in time. If it is known that the two histories being added are of equal path length and that during the interaction both histories progress in a synchronous manner, only states with the same time stamp can be added. Thus in Figure 6 the leftmost history is the product of a synchronous progression of component histories $a1b1 \rightarrow a2b2 \rightarrow a3b3$.

ii/ Asynchronous progression in time. The complement of the previous rule is that the histories being added progress asynchronously. Here only one of the component states can progress at a time.

While the first two constraints are always applied during path creation, the specialised temporal constraints are applied only when additional information is available to the reasoning system. For example, a monitoring system will not always know that two diseases are evolving in a synchronous manner.

3.3 Qualitative Superposition of Non-linear Systems

The process of history addition described so far has been applied to histories arising from linear systems. Unfortunately many systems of interest are non-linear, and the principle of superposition for ODEs does not extend to non-linear systems. However, relaxing the requirements for superposition to produce qualitative solutions may allow it to be applied to a range of non-linear systems. One thus seeks situations in which solutions to non-linear QDEs behave qualitatively as if they came from a linear system.

The problem can be approached in the following way. All QSIM systems are first order and any n th order non-linear system can be reduced to such a first order system. Each first order QDE that could appear in a system can be mapped onto the QSIM constraint types identified previously. Thus the problem reduces to identifying situations in which qualitative superposition is legal for each non-linear constraint type.

A constraint that satisfies superposition for non-linear systems will be called a

reasonable constraint, and we can define such a constraint formally:

Let C be a QDE of the form $C(a,b)$ where a and b are continuous real valued functions. Let $a_1, a_2, .. a_n$ be solutions for a and $b_1, b_2, .. b_n$ be the solutions for b , such that $C(a_1, b_1), C(a_2, b_2), .. C(a_n, b_n)$ are solutions for C .

If $C(a_x + a_y, b_x + b_y)$ is true for any solutions x, y then C is a reasonable constraint.

Starting with the formal constraint definitions in Kuiper's (1986) paper, it can be shown without difficulty that *add* and *deriv* will be reasonable under all circumstances. Non-linear behaviour however resides in QSIM's M^+ , M^- and *mult* constraints.

We can prove formally (Coiera,1992b) that monotonic constraints can support superposition for most combinations of qualitative values, except when the derivatives being added are of opposite sign. In this case, as with the multiplication constraint, when one of the two solutions being added dominates in the sum solution, then the constraint will display reasonable behaviour.

Dominance of one history over another means that when there is ambiguity in the sum of their values, we choose the value which favours the dominant history. For example, take the CO_2 values from the respiratory acidosis and metabolic acidosis histories in Figure 5 at $t(0)/t(1)$. The value is positive in the respiratory acidosis QS($0/\infty, dec.$) and negative in the metabolic acidosis QS($-\infty/0, dec.$). Using qualitative algebra, the result of the addition of these two values could be positive, negative or zero. If the positive respiratory acidosis value dominates then the result will also be positive.

While history dominance is the most useful condition identified so far with which non-reasonable system histories can be made to exhibit reasonableness, there may potentially be many such conditions. Their identification should be a fruitful research area.

3.3.1 Reasoning with Empirically Derived Histories

One advantage of reasoning at the level of histories is that with some poorly understood domains, we can still perform relatively 'deep' level reasoning about history interactions in the absence of a system model. For a non-model derived history to be useful in the framework presented above it must be well-formed. By this we mean that it should display all the features of a true model-derived history such as continuity of function behaviour amongst real-valued parameters.

The results presented above indicate both that it is possible to reason with empirically derived histories, and caution that in some circumstances the conclusions so derived could be in error. In particular, if the underlying system is non-linear, and the constraints are not reasonable, then the process of history addition is not legal. Since a reasoning system by definition does not know whether an empirically derived history comes from a linear or non-linear system, any conclusions about history interactions must be considered defeasible.

Implicit in the exposition of history addition was the assumption that the histories being added came from the same system model. It makes no mathematical sense to add solutions to different sets of QDEs. Thus when dealing with empirically derived histories, they should be the products of the same underlying system. However, this is not universally possible to ensure. For example, assume a

patient currently has a disease A. When a new disease event B occurs that can be modelled as a perturbation within the existing system model (as in the previous example), we can reason about that interaction with histories. If however, a new disease event B occurs that alters the underlying physiological mechanisms influencing A, then we can no longer reason with the old A history. In effect a new model has come into being, and a new A history would need to be obtained if we wanted to reason about A and B at the history level. In situations like this, a reasoning system would need to switch to a model-based reasoning modality, hypothesising new mechanism faults in its system model and then projecting their effects.

4 CONCLUSION

The premise at the beginning of this paper was that while a diagnostic monitoring system might need to manipulate knowledge at a deep level because of the type of problem presented to it, that level need not be as deep as a full system model. Histories have the advantage that they are in effect compiled model-based simulations and can be stored within a reasoning system. Further, their manipulation is computationally less intensive than a full model-based simulation.

The mechanisms by which QSIM derived histories can be used to make predictions about disease interaction are based on the mathematical definition of the QSIM system. As such, they are general rather than domain specific. However, one would caution against making general conclusions about the suitability of this representation until work with competing paradigms has been explored.

The major benefit of using such a formal knowledge representation is that it both allows us to define where history based reasoning is appropriate, and also when it is inappropriate. We now have some direct measure of the trade off in predictive power as we move from one level of abstraction to a deeper one. Further, we can accurately define when switching to the deeper model-based representations is appropriate. This is of critical importance as it is becoming clear that diagnostic monitoring systems must employ a layered approach to problem solving (Hayes-Roth et al., 1989). Several different knowledge representations and inference methods may need to be invoked to deal with the problems presented in a single patient.

Reasoning with empirically derived system histories is one of the most exciting prospects of the history based approach, despite its inherent uncertainties. The two major limitations to history based reasoning were raised in relation to empirically derived histories, but are generally true. Firstly, we may not be able to use histories when no reasonableness conditions apply and the system being examined is non-linear. Secondly, when a change in the system model occurs because of a new event, then any current histories are no longer valid. In effect, the change in system model requires that each current history be recomputed. It is clear from these limitations that not all monitoring tasks will be amenable to history based reasoning. Equally, however, not all domains are amenable to formal modelling, and history based prediction will offer more powerful reasoning than that offered by rule-based systems.

The formal limits to history based reasoning raise other interesting questions

about the ways physicians utilise disease histories for prediction tasks. The limits presented here may not be an everyday problem because the disease systems reasoned with may inherently approximate linear systems, or because the physician intuitively turns to a full first principles model based approach when difficult interactions are encountered. Another possibility is that the reasoning mechanism used by physicians is robust to prediction errors, the physician quickly revising predictions when they obviously are not born out clinically. Perhaps the likeliest explanation is that clinicians have access to contextual information that will allow calculations of the relative magnitude of effects from individual histories to be taken into account.

In conclusion, diagnostic monitoring systems are becoming an increasingly important research area, and several competing architectures and knowledge representations are in evolution. History based systems have their place within diagnostic monitoring systems, and offer to be efficient reasoning tools. Further, they offer a generality that should find them used in a variety of systems complementing pure model-based techniques.

REFERENCES

- I.Bratko, Qualitative modelling and learning in KARDIO, *Proceedings Fifth Australian Conference on Applications for Expert Systems*, New South Wales Institute of Technology, (1989), 1 - 22.
- A.Cohn, S.Rosenbaum, M.Factor, P.Miller, DYNASCENE: An approach to computer-based intelligent cardiovascular monitoring using sequential clinical scenes, *Meth. Inform. Med.*, 29, (1990), 122-131.
- E.Coiera, Reasoning with Qualitative Disease Histories for Diagnostic Patient Monitoring, Ph.D Thesis, School of Electrical Engineering and Computer Science, University of New South Wales, Sydney, Australia, (1989).
- E.Coiera, Qualitative Disease Histories - Optimising Representational Depth in a Multilevel System, *Workshop on Representing Knowledge in Medical Decision Support Systems, 12th International Joint Conference on Artificial Intelligence*, Sydney, Australia, (1991).
- E.Coiera, The Qualitative Representation of Physical Systems, *The Knowledge Engineering Review*, 7,1,(1992).
- E.Coiera, Qualitative Superposition, *Artificial Intelligence*, (1992), to appear.
- R.Davis, Diagnostic Reasoning Based on Structure and Behaviour, *Artificial Intelligence*, 24, (1984), 347 - 410.
- R.Davis, Robustness and Transparency in Intelligent Systems, *Proceedings Third Australian Conference on the Applications of Expert Systems*, Sydney, (1987), 143 - 164.
- R.Doyle, S.Sellers, D.Atkinson, A Focused, Context-Sensitive Approach to Monitoring, *Proceedings 11th International Joint Conference on Artificial Intelligence*, (1989), 1231-1237.
- J.de Kleer, J.S Brown, A Qualitative Physics Based on Confluences, *Artificial Intelligence*, 24, (1984), 7 - 83.

- D.Dvorak, Expert Systems for Monitoring and Control, Technical Report AI87-55, Artificial Intelligence Laboratory, The University of Texas at Austin, (1987).
- D.Dvorak, B.Kuipers, Model-Based Monitoring of Dynamic Systems, *Proceedings 11th International Joint Conference on Artificial Intelligence*, (1989), 1238 - 1243.
- L.M.Fagan, J.C.Kunz, E.A.Feigenbaum, J.Osborn, Extensions to the Rule-Based formalism for a Monitoring Task, in B.G.Buchanan, E.H.Shortliffe, (eds.), *Rule Based Expert Systems*, Addison Wesley, (1984).
- K.D.Forbus, Qualitative Process Theory, *Artificial Intelligence*, 24, (1984), 85 - 168.
- K.D.Forbus, The Logic of Occurrence, *Proceedings 10th International Joint Conference on Artificial Intelligence*, (1987), 409-415.
- K.D.Forbus, Commonsense Physics, *Ann. Rev. Comput. Sci.*, 3, (1988), 197-232.
- P.J.Hayes, The Naive Physics Manifesto, in D.Michie (ed.), *Expert Systems in the Microelectronic Age*, Edinburgh University Press, (1979).
- B.Hayes-Roth, R.Washington, R.Hewett, et al., Intelligent Monitoring and Control, *Proceedings 11th International Joint Conference on Artificial Intelligence*, (1989), 243-249.
- L.Ironi, M.Stefannelli, G.Lanzola, Qualitative Models in Medical Diagnosis, *Artificial Intelligence in Medicine*, 2,(1990), 85-101.
- E.Keravnou, J.Washbrook, Deep and shallow models in medical expert systems, *Artificial Intelligence in Medicine*, 1, (1989), 11-28.
- B.Kuipers, Qualitative Simulation in Medical Physiology: A Progress Report, Technical Report MIT/LCS/TM-280, Laboratory for Computer Science, MIT, (1985).
- B.Kuipers, Qualitative Simulation, *Artificial Intelligence*, 29, (1986), 289-338.
- B.Kuipers, C.Chiu, Taming Intractable Branching in Qualitative Simulation, *Proceedings 10th International Joint Conference on Artificial Intelligence*, (1987), 1079-1085.
- W.Lee, B.Kuipers, Non-Intersection of Trajectories in Qualitative Phase Space: A Global Constraint for Qualitative Simulation, *AAAI-88 - Proceedings Seventh National Conference on Artificial Intelligence*, (1988), 286-290.
- M.Mavrovountiotis, G.Stephanopoulos, Formal Order of Magnitude Reasoning in Process Engineering, *Computer Chemical Engineering*, 12, (1988), 867-880.
- R.G.Narins, M.Emmett, Simple and Mixed Acid-Base Disorders: A Practical Approach, *Medicine*, 59, (1980), 161 - 187.
- A.L.Rabenstein, *Introduction to Ordinary Differential Equations*, Academic Press, New York, (1972).
- O.Raiman, Order of Magnitude Reasoning, *AAAI-88 - Proceedings Seventh National Conference on Artificial Intelligence*, (1988), 100-104.
- S.Uckun, B.Dawant, Qualitative modelling as a paradigm for diagnosis, prediction and therapy management in critical care environments, *Artificial Intelligence in Medicine*, (1992), (to appear).

- S.M.Weiss, C.A.Kulikowsky, R.S.Galen, Developing Microprocessor-Based Expert Models for Instrument Interpretation, *Proceedings Seventh International Joint Conference on Artificial Intelligence*, (1981), 853 - 855.
- B.Williams, MINIMA - A Symbolic Approach to Qualitative Algebraic Reasoning, *AAAI-88 - Proceedings Seventh National Conference on Artificial Intelligence* (1988), 264-269.
- D.Weld, J.de Kleer, (eds.) *Qualitative Reasoning About Physical Systems*, Morgan Kaufman Publishers Inc., San Mateo, (1990).