NASA-CR-193812

GRANT
IN-61-CR
181315
17P

The University of Texas at Austin
Austin, Texas 78712 USA

System Monitoring and Diagnosis With Qualitative Models.

NAG 2-507.

Final Report

covering the period
1 March 1988
to
30 September 1991.

Prof. Benjamin Kuipers, Principal Investigator
Department of Computer Sciences

NASA Technical Officer: Brian Glass
NASA Ames Research Center
Moffett Field, CA 94035

(NASA-CR-193812) SYSTEM MONITORING AND DIAGNOSIS WITH QUALITATIVE MODELS Final Report, 1 Mar. 1988 - 30 Sep. 1991 (Texas Univ.) 17 p

N94-10711

Unclas

G3/61 0181315

CASI

1 System Monitoring and Diagnosis With Qualitative Models

The world is infinite, continuous, and continually changing over time. Human knowledge and human inference abilities are finite, apparently symbolic, and therefore incomplete. Nonetheless, people normally reason quite effectively about the physical world.

Models of particular systems or mechanisms play an important role in this capability. In service of a task such as diagnosis or design, simulation predicts the behaviors that follow from a particular model. In diagnosis or explanation, these predictions include testable consequences of a diagnostic hypothesis. In design, these predictions make explicit the consequences of a set of design choices.

A qualitative differential equation (QDE) model is a symbolic description expressing a state of incomplete knowledge of the continuous world, and is thus an abstraction of an infinite set of ordinary differential equations models. Qualitative simulation predicts the set of possible behaviors consistent with a QDE model and an initial state.

We have developed a substantial foundation of tools for model-based reasoning with incomplete knowledge: QSIM and its extensions for qualitative simulation; Q2, Q3 and their successors for quantitative reasoning on a qualitative framework; and the CC and QPC model compilers for building QSIM QDE models starting from different ontological assumptions.

The QSIM representation for qualitative differential equations (QDEs) and qualitative behaviors was originally motivated by protocol analysis studies of expert explanations [Kuipers & Kassirer, 1984]. A QDE represents a set of ODEs consistent with natural states of human incomplete knowledge of a physical mechanism [Kuipers, 1984]. Qualitative simulation can be guaranteed to produce a set of qualitative behavior descriptions covering all possible behaviors of all ODEs covered by the QDE [Kuipers, 1986, 1988b, 1989a].

The subsequent evolution of QSIM has been dominated by the mathematical problems of retaining this guarantee while producing a tractable set of predictions. A variety of methods now exist for applying a deeper analysis, changing the level of description, or appealing to carefully chosen additional assumptions, to obtain tractable predictions from a wide range of useful models [Kuipers 1987, 1988a, 1989b; Kuipers & Chiu, 1987; Lee & Kuipers, 1988; Fouché & Kuipers, 1990, 1992; Kuipers, et al, 1991].

Quantitative information can be used to annotate qualitative behaviors, preserving the coverage guarantee while providing stronger predictions. Quantitative information may be expressed as bounds on landmarks and other symbolic elements of the qualitative description [Kuipers & Berleant, 1988], by adaptively inserting new time-points to improve the resolution of the description and converge to a numerical function [Berleant & Kuipers, 1990], and by deriving envelopes bounding the possible trajectories of the system [Kay & Kuipers, 1991]. Observations are interpreted by unifying quantitative measurements against the qualitative behavior prediction, yielding either a stronger prediction or a contradiction. As quantitative uncertainty in the QDE and initial state decrease to zero, the resulting behavioral description converges to the true quantitative behavior, though computational costs can still be high with current methods.

We have developed two model-compilers for QDE models: CC, which takes the component-connection view of a mechanism [Franke & Dvorak, 1989], and QPC, which implements an extended version of Qualitative Process Theory [Crawford, et al, 1990]. Other model-compilers for QDEs, e.g. using bond graphs or compartmental models, have been developed elsewhere. These model-buiding tools will support automatic construction of qualitative models from physical specifications, and further research into selection of appropriate modeling viewpoints.

There are several inference schemes built on the set of all possible behaviors that are particularly well-suited to reliable model-based reasoning for diagnosis and design. For design, desirable and undesirable behaviors can be identified, and additional constraints inferred to guarantee or prevent those behaviors [Franke, 1989, 1991]. This capability supports the design, analysis, and validation of heterogeneous, non-linear controllers even under incomplete knowledge [Kuipers & Åström, 1991].

For monitoring and diagnosis, plausible hypotheses are unified against observations to strengthen or refute the predicted behaviors. In MIMIC [Dvorak & Kuipers, 1989, 1991], multiple hypothesized models of the system are tracked in parallel in order to reduce the "missing model" problem. Each model begins as a qualitative model, and is unified with *a priori* quantitative knowledge and with the stream of incoming observational data. When the model/data unification yields a contradiction, the model is refuted. When there is no contradiction, the predictions of the model are progressively strengthened, for use in procedure planning and differential diagnosis. Only under a qualitative level of description can a finite set of models guarantee the complete coverage necessary for this performance. The MIMIC approach to monitoring and diagnosis has become very influential, and we are continuing research on it.

2 Publications on the Topic of the NASA Grant

These papers present the results of the research program supported by NASA grant NAG 2-507. A few papers are included that date from before the grant, to show the context of the work, and some of the papers cited below were supported by other funding but represent work that was synergistic with the NASA grant.

1. D. Berleant & B. Kuipers. Combined qualitative and numerical simulation with Q3. *Papers of the Fourth International Workshop on Qualitative Physics*, Lugano, Switzerland, 9-12 July 1990. To appear in Boi Faltings and Peter Struss (Eds.), *Recent Advances in Qualitative Physics*, MIT Press, 1991.

2. C. Chiu & B. J. Kuipers. 1991. Comparative analysis and qualitative integral representations. *Papers of the Third International Workshop on Qualitative Physics*, Stanford, California, July 1989. To appear in Boi Faltings and Peter Struss (Eds.), *Recent Advances in Qualitative Physics*, MIT Press, 1991.

3. J. M. Crawford, A. Farquhar, B. J. Kuipers. 1990. QPC: a compiler from physical models into qualitative differential equations. *Proceedings of the National Conference on Artificial Intelligence (AAAI-90)*, AAAI/MIT Press, 1990.

4. D. T. Dalle Molle, B. J. Kuipers, and T. F. Edgar. 1988. Qualitative modeling and simulation of dynamic systems. *Computers and Chemical Engineering* **12**: 853-866, 1988.

5. D. Dvorak and B. Kuipers. 1989. Model-based monitoring of dynamic systems. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI-89).* Los Altos, CA: Morgan Kaufman.

6. D. Dvorak & B. Kuipers. 1991. Process monitoring and diagnosis: a model-based approach. *IEEE EXPERT* **6**(3): 67-74, June 1991.

7. D. L. Dvorak, D. T. Dalle Molle, B. J. Kuipers, and T. F. Edgar. 1990. Qualitative simulation for expert systems. 1990 Congress, International Federation of Automatic Control (IFAC), Tallin, Estonia, USSR.

8. P. Fouché & B. Kuipers. 1991. An assessment of current qualitative simulation techniques. *Papers of the Fourth International Workshop on Qualitative Physics*, Lugano, Switzerland, 9-12 July 1990. To appear in Boi Faltings and Peter Struss (Eds.), *Recent Advances in Qualitative Physics*, MIT Press, 1991.

9. P. Fouché & B. Kuipers. 1992. Reasoning about energy in qualitative simulation. To appear *IEEE Transactions on Systems, Man, and Cybernetics* **22**(1), 1992.

10. D. W. Franke. Representing and acquiring teleological descriptions. Model-Based Reasoning Workshop, IJCAI-89, Detroit, Michigan, August 1989.

11. D. W. Franke. 1991. Deriving and using descriptions of purpose. *IEEE Expert*, April 1991, pp. 41-47.

12. D. W. Franke and D. Dvorak. 1989. Component-connection models. Model-Based Reasoning Workshop, IJCAI-89, Detroit, Michigan, August 1989.

13. H. Kay & B. Kuipers. 1991. Numerical behavior envelopes for qualitative models. Manuscript.

14. B. J. Kuipers. 1984. Commonsense reasoning about causality: deriving behavior from structure. *Artificial Intelligence* **24**: 169 - 204.

15. B. J. Kuipers. 1986. Qualitative simulation. *Artificial Intelligence* **29**: 289 - 338.

16. B. Kuipers. 1987. Abstraction by time-scale in qualitative simulation. *Proceedings of the National Conference on Artificial Intelligence (AAAI-87).* Los Altos, CA: Morgan Kaufman.

17. B. J. Kuipers. 1988a. Qualitative simulation using time-scale abstraction. *Int. J. Artificial Intelligence in Engineering* **3**(4): 185-191, 1988.

18. B. J. Kuipers. 1988b. The qualitative calculus is sound but incomplete: a reply to Peter Struss. *Int. J. Artificial Intelligence in Engineering* **3**(3): 170-173, 1988.

19. B. Kuipers. 1989a. Qualitative reasoning: modeling and simulation with incomplete knowledge. *Automatica* **25**: 571-585.

20. B. J. Kuipers. 1989b. Qualitative reasoning with causal models in diagnosis of complex systems. In L. Widman, K. Loparo, & N. Nielson (Eds.), *Artificial Intelligence, Simulation and Modeling.* New York: John Wiley & Sons, 1989, pp. 257-274.

21. B. J. Kuipers. 1990. Simulation, Qualitative. In M. G. Singh (Ed.), *Systems & Control Encyclopedia, Supplementary Volume 1.* NY: Pergamon Press.

22. Benjamin Kuipers & Karl Åström. 1991. The composition of heterogeneous control laws. In *Proceedings of the American Control Conference*, 1991.

23. B. Kuipers and D. Berleant. 1988. Using incomplete quantitative knowledge in qualitative reasoning. In *Proceedings of the National Conference on Artificial Intelligence (AAAI-88).* Los Altos, CA: Morgan Kaufman.

24. B. Kuipers and D. Berleant. 1990. A smooth integration of incomplete quantitative knowledge into qualitative simulation. UT AI TR 90-122.

25. B. Kuipers and C. Chiu. 1987. Taming intractible branching in qualitative simulation. *Proceedings of the Tenth International Joint Conference on Artificial Intelligence (IJCAI-87).* Los Altos, CA: Morgan Kaufman.

26. B. J. Kuipers, C. Chiu, D. T. Dalle Molle & D. R. Throop. 1991. Higher-order derivative constraints in qualitative simulation. *Artificial Intelligence* **51**: 343-379.

27. B. J. Kuipers and J. P. Kassirer. 1984. Causal reasoning in medicine: analysis of a protocol. *Cognitive Science* **8**: 363 - 385.

28. W. W. Lee and B. Kuipers. 1988. Non-intersection of trajectories in qualitative phase space: a global constraint for qualitative simulation. In *Proceedings of the National Conference on Artificial Intelligence (AAAI-88).* Los Altos, CA: Morgan Kaufman.

3 Abstracts of Relevant Papers

1. D. Berleant & B. Kuipers. Combined qualitative and numerical simulation with Q3. *Papers of the Fourth International Workshop on Qualitative Physics*, Lugano, Switzerland, 9-12 July 1990. To appear in Boi Faltings and Peter Struss (Eds.), *Recent Advances in Qualitative Physics*, MIT Press, 1991.

Abstract

A simulation is a sequence of predicted states of a modeled system. A *qualitative-quantitative* simulation is a simulation containing both qualitative, and quantitative, state information such that the qualitat6ive information alone would be a qualitative simulation, and the quantitative information alone would be a numberical simulation. In this paper, each state is described with both qualitative and numerical data. Qualitative-quantitative simulation is a generalization of both qualitative simulation and numerical simulation, providing a framework for viewing historically disparate *genres* of simulation.

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

2. C. Chiu & B. J. Kuipers. 1991. Comparative analysis and qualitative integral representations. *Papers of the Third International Workshop on Qualitative Physics*, Stanford, California, July 1989. To appear in Boi Faltings and Peter Struss (Eds.), *Recent Advances in Qualitative Physics*, MIT Press, 1991.

Abstract

Comparative analysis is applied to a qualitative behavior of an incompletely known mechanism, to determine the effect of a given perturbation on the behavior as a whole. This class of inference is useful in diagnosis, design, planning, and generally for understanding the relations among a set of alternate qualitative behaviors.

Comparative analysis depends on information which is implicit, and relatively difficult to extract, from qualitative differential equations. By introducing the definite integral as a descriptive term linking qualitative variables and their landmarks, we show that the qualitative *integral* representation (QIR) makes the required information easily accessible.

Inspired by observations of expert physicists, we have adopted an approach to inference that allows global algebraic manipulation of the QIR. Within this approach, comparative analysis can be decomposed into a search and algebraic manipulation problems. Several detailed examples are presented to clarify our method.

3. J. M. Crawford, A. Farquhar, B. J. Kuipers. 1990. QPC: a compiler from physical models into qualitative differential equations. *Proceedings of the National Conference on Artificial Intelligence (AAAI-90)*, AAAI/MIT Press, 1990.

Abstract

Qualitative reasoning can, and should, be decomposed into a *model-building* task, which creates a qualitative differential equation (QDE) as a model of a physical situation, and a *qualitative simulation* task, which starts with a QDE, and predicts the possible behaviors following from the model.

In support of this claim, we present QPC, a model builder that takes the general approach of Qualitative Process Theory, describing a scenario in terms of views, processes, and influences. However, QPC builds QDEs for simulation by QSIM, which gives it access to a variety of mathematical advances in qualitative simulation incorporated in QSIM.

We present QPC and its approach to Qualitative Process Theory, provide an example of building and simulating a model of a non-trivial mechanism, and compare the representation and implementation decisions underlying QPC with those of QPE.

4. D. T. Dalle Molle, B. J. Kuipers, and T. F. Edgar. 1988. Qualitative modeling and simulation of dynamic systems. *Computers and Chemical Engineering* **12**: 853-866, 1988.

Abstract

Qualitative simulation is a promising technique for design and analysis, particularly in model-based reasoning systems. The purpose of qualitative simulation is to explain process observations by reasoning from physical descriptions to behavioral descriptions. Qualitative simulation has the ability to yeild partial conclusions from incomplete knowledge of the process. In this work, the qualitative simulation algorithm, QSIM, is used to model qualitatively several systems from chemical engineering. The QSIM algorithm successfully generated qualitative descriptions for the open-loop responses for all of the systems studied including linear, nonlinear and multivariable processes. Some models required the use of redundant constraints to eliminate otherwise ambiguous parameters. The closed-loop behavior of a mixing tank has also been successfully analyzed with qualitative versions of feedback control. The correct dynamic behavior and the qualitative features of proportional control, such as offset, are captured by the QSIM algorithm.

5. D. Dvorak and B. Kuipers. 1989. Model-based monitoring of dynamic systems. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI-89).* Los Altos, CA: Morgan Kaufman.

Abstract

Industrial process plants such as chemical refineries and electric power generation are examples of continuous-variable dynamic systems (CVDS) whose operation is continuously monitored for abnormal behavior. CVDSs pose a challenging disgnostic problem in which values are continuous (not discrete), relatively few parameters are observable, parameter values keep changing, and diagnosis must be performed while the system operates.

We present a novel method for monitoring CVDSs which exploits the system's dynamic behavior for diagnostic clues. The key techniques are: modeling the physical system with dynamic qualitative/quantitative models, inducing diagnostic knowledge from qualitative

simulations, continuously comparing observations against fault-model predictions, and incrementally creating and testing multiple-fault hypothesis. The important result is that the diagnosis is refined as the physical system's *dynamic* behavior is revealed over time.

6. D. Dvorak & B. Kuipers. 1991. Process monitoring and diagnosis: a model-based approach. *IEEE EXPERT* **6**(3): 67-74, June 1991.

Abstract

This paper describes a method for monitoring and diagnosis of process systems based on three foundational technologies: semi-quantitative simulation, measurement interpretation (tracking), and model-based diagnosis. Compared to existing methods based on fixed-threshold alarms, fault dictionaries, decision trees, and expert systems, several advantages accrue:

- imprecise knowledge of parameter values and functional relationships (both linear and non-linear) can be expressed in the semi-quatitative model and used during simulation, producing a valid range for each variable;
- incremental simulation of the model in step with incoming sensor readings, with subsequent comparison of observations to predictions, permits earlier fault detection than with fixed thresholds;
- by using a structural model of the plant and tracing upstream from the site of unexpected readings, model-based diagnosis permits efficient generation of fault candidates without resort to pre-compiled (and often incomplete) symptom–fault patterns;
- by injecting a hypothesized fault into the model and tracking its predictions against observations, the dynamic behavior of the plant is exploited to corroborate or refute hypotheses;
- by simulating ahead in time from the current state, an operator can be forewarned of nearby undesireable states that the plant might enter.

7. D. L. Dvorak, D. T. Dalle Molle, B. J. Kuipers, and T. F. Edgar. 1990. Qualitative simulation for expert systems. 1990 Congress, International Federation of Automatic Control (IFAC), Tallin, Estonia, USSR.

Abstract

Monitoring dynamic chemical processes poses a challenging diagnostic problem when the diagnosis must be performed while the system operates, when multiple faults are common, and when observations are limited to a relatively small set of variables. The monitoring process involves collecting measurements from sensors, combining this data into a picture of the current state of the system, and assessing any departure from expected behavior. We present a method called MIMIC for monitoring continuous-variable dynamic systems. MIMIC relies primarily on knowledge derived from a qualitative or semi-quantative model of the monitored system and exploits the system's temporal behavior for diagnosis. The goal of the disgnostic system is to mimic the condition of the physical system by identifying parameter ranges in a model of the process that are consistent with the observations.

8. P. Fouché & B. Kuipers. 1991. An assessment of current qualitative simulation techniques. *Papers of the Fourth International Workshop on Qualitative Physics*, Lugano, Switzerland, 9-12 July 1990. To appear in Boi Faltings and Peter Struss (Eds.), *Recent Advances in Qualitative Physics*, MIT Press, 1991.

Abstract

QSIM is a powerful Qualitative Simulation algorithm, which now includes many features that have proven to be necessary in Qualitative Simulation. These features are: reasoning with Higher-Order Derivatives, having Multiple Levels of Abstraction, reasoning in the Phase Space representation, and reasoning about Energy. The aim of this paper is to provide a comprehensive view of all these techniques, by explaining their rationale, showing the problems they address and how they interact. Remaining problems in Qualitative Simulation are also discussed.

9. P. Fouché & B. Kuipers. 1992. Reasoning about energy in qualitative simulation. To appear *IEEE Transactions on Systems, Man, and Cybernetics* **22**(1), 1992.

Abstract

Qualitative modeling and simulation make it feasible to predict the possible behaviors of a mechanism consistent with an incomplete state of knowledge. Though qualitative simulation predicts all possible behaviors of a system, it can also produce suprious behaviors, i.e. behaviors which correspond to no solution of any ordinary differential equation consistent with the qualitative model. In this paper we present a method for reasoning about energy, which eliminates an important source of spurious behaviors. We apply this method to an industrially significant mechanism – a non-linear, proportional-integral controller – and show that qualitative simulation captures the main qualitative properties of such a system, such as stability and zero-offset control. We believe that this is a significant step toward the application of qualitative simulation to model-based monitoring, diagnosis, and design of realistic mechanisms.

10. D. W. Franke. Representing and acquiring teleological descriptions. Model-Based Reasoning Workshop, IJCAI-89, Detroit, Michigan, August 1989.

Abstract

Teleological descriptions capture the *purpose* of an entity, mechanism, or activity with which they are associated. These descriptions can be utilized in diagnostic reasoning by providing focus in hypothesis generation and selection. Teleological descriptions can also be utilized in design to index existing designs for reuse and to express design rationale.

While a teleological description of a mechanism is distinct from any structural and behavioral descriptions, it is claimed that a teleological description is constructed with references to elements of structural and behavioral descriptions. In particular, the purpose of a component or activity can be expressed in terms of the behaviors it *prevents* or *guarantees*. While these teleological descriptions reference elements of behavioral descriptions, they are independent of any particular behavioral language or model domain. Higher level operators can be constructed from these primitive operators. A technique for deriving teleological descriptions is described, along with its relationship to design requirements and constraints.

11. D. W. Franke. 1991. Deriving and using descriptions of purpose. *IEEE Expert*, April 1991, pp. 41-47.

Abstract

When one examines human-generated descriptions of systems or mechanisms, one finds that they are rich with descriptions of purpose. Given representation and acquisition schemes, such descriptions can be utilized in explanation, diagnostic, and design systems. We describe a language, TeD, for representing descriptions of purpose, along with a design method in which descriptions of purpose can be captured and subsequently utilized for design reuse. This language is independent of any particular structure or behavior description languages, but builds upon generalizations of such languages. In particular, the purpose of a component or activity is expressed in terms of behaviors prevented, guaranteed, or introduced by the component or activity. The detailed relationship between TeD and structure and behavior descriptions is described, and a design method for acquiring and utilizing teleology descriptions is given for an example design.

12. D. W. Franke and D. Dvorak. 1989. Component-connection models. Model-Based Reasoning Workshop, IJCAI-89, Detroit, Michigan, August 1989.

Abstract

The relation between *part* and *whole* is the key to describing the structure of a mechanism. Different modeling methods have different concepts of what should count as a "part" of a system, and how the parts should relate to each other. The mathematical, differential-equation-based approach to modeling taken in QSIM essentially says that the "parts" of a mechanism are the continuous variables that characterize its state, and their relations are mathematical constraints inherited from the physical structure of the system.

However, a physical system frequently consists of a set of components that relate through explicit connections (a form of description that is frequently more meaningful to a domain expert than the differential equations). This paper describes CC, a model-building program that accepts a component-connection description of a physical system and translates it to the qualitative differential equations of QSIM. CC provides facilities for component abstraction and hierarchical component definition, raising the level of abstraction for modeling via QSIM. Component *modes* can be specified, and are translated into QSIM operating regions. CC uses the general variable types of bond graphs (a technique for dynamic physical system modeling). Finally, this component-connection paradigm provides the framework for information utilized in other model-based reasoning tasks such as diagnosis.

13. H. Kay & B. Kuipers. 1991. Numerical behavior envelopes for qualitative models. Manuscript.

Abstract

We describe a method for improving the bounds on the behaviors of a qualitative differential equation (QDE) model augmented with numerical information, by numerically simulating systems whose solutions are guaranteed to bound the solutions of any system that corresponds to the QDE. It is shown that when such systems exist, they can be determined automatically given the QDE and an initial condition. We explain our method and compare it with other approaches on a simple first-order model. Finally, we show how the method improves the dynamic monitoring and diagnosis of a vacuum pump-down system.

14. B. J. Kuipers. 1984. Commonsense reasoning about causality: deriving behavior from structure. *Artificial Intelligence* **24**: 169 - 204.

Abstract

This paper presents a qualitative reasoning method for predicting the behavior of mechanisms characterized by continuous, time-varying parameters. The structure of a mechanism is described in terms of a set of parameters and the constraints that hold among them: essentially a "qualitative differential equation." The qualitative behavior description consists of a discrete set of time-points, at which the values of the parameters are described in terms of ordinal relations and directions of change. The behavioral description, or envisionment, is derived by two sets of rules: propagation rules which elaborate the description of the current time-point, and prediction rules which determine what is known about the next qualitatively distinct state of the mechanism. A detailed example shows how the envisionment method can detect a previously unsuspected landmark point at which the system is in stable equilibrium

15. B. J. Kuipers. 1986. Qualitative simulation. *Artificial Intelligence* **29**: 289 - 338.

Abstract

Qualitative simulation is a key inference process in qualitative causal reasoning. However, the precise meaning of the different proposals and their relation with differential equations is often unclear. In this paper, we present a precise definition of qualitative structure and behavior descriptions as abstractions of differential equations and continuously differentiable functions. We present a new algorithm for qualitative simulation that generalizes the best features of existing algorithms, and allows direct comparasons among alternate approaches. Starting with a set of constraints abstracted from a differential equation, we prove that the QSIM algorithm is guaranteed to produce a qualitative behavior corresponding to any solution to the original equation. We also show that any qualitative simulation algorithm will sometimes produce spurious qualitative behaviors: ones which do not correspond to any mechanism satisfying the given constraints. These observations suggest specific types of care that must be taken in designing applications of qualitative causal reasoning systems, and in constructing and validating a knowledge base of mechanism descriptions.

16. B. Kuipers. 1987. Abstraction by time-scale in qualitative simulation. *Proceedings of the National Conference on Artificial Intelligence (AAAI-87).* Los Altos, CA: Morgan Kaufman.

Abstract

Qualitative simulation faces an intrinsic problem of scale: the number of limit hypotheses grows exponentially with the number of parameters approaching limits. We present a method called *Time-Scale Abstraction* for structuring a complex system as a hierarchy of smaller, interacting equilibrium mechanisms. Within this hierarchy, a given mechanism views a slower one as being constant, and a faster one as being instantaneous. A perturbation to a fast mechanism may be seen by a slower mechanism as a displacement of a monotonic function constraint. We demonstrate the time-scale abstraction hierarchy using the interaction between the water and sodium balance mechanisms in medical physiology, an example drawn from a larger, fully implemented, program. Where the structure of a large system permits decomposition by time-scale, this abstraction method permits qualitative simulation of otherwise intractibly complex systems.

17. B. J. Kuipers. 1988a. Qualitative simulation using time-scale abstraction. *Int. J. Artificial Intelligence in Engineering* **3**(4): 185-191, 1988.

Abstract

Qualitative simulation faces an intrinsic problem of scale: the number of limit hypotheses grows exponentially with the number of parameters approaching limits. We present a method called *Time-Scale Abstraction* for structuring a complex system as a hierarchy of smaller, interacting equilibrium mechanisms. Within this hierarchy, a given mechanism views a slower one as being constant, and a faster one as being instantaneous. A perturbation to a fast mechanism may be seen by a slower mechanism as a displacement of a monotonic function constraint. We demonstrate the time-scale abstraction hierarchy using the interaction between the water and sodium balance mechanisms in medical physiology, an example drawn from a larger, fully implemented, program. Where the structure of a large system permits decomposition by time-scale, this abstraction method permits qualitative simulation of otherwise intractibly complex systems.

18. B. J. Kuipers. 1988b. The qualitative calculus is sound but incomplete: a reply to Peter Struss. *Int. J. Artificial Intelligence in Engineering* **3**(3): 170-173, 1988.

Abstract

Peter Struss has made a valuable contribution to the mathematics of qualitative reasoning through his careful analysis of qualitative algebras. In particular, he has firmly demonstrated that the *varying granularity* property of qualitative representations is incompatible with familiar algebraic properties such as the associative and distributive laws.

However, two points of clarification are required:

(a) Qualitative methods for solving algebraic (and differential) equations are correctly regarded as *sound* but *incomplete*. Struss' assertion, that these methods are complete but unsound, is incorrect.

(b) There are several useful techniques that ameliorate the impact of the incompleteness: meta-level reasoning with solutions to simple quantitative instances of qualitative equations; inclusion of constraints that are quantitatively redundant, but qualitatively independent; and choice of landmarks to provide corresponding values across quantity spaces.

19. B. Kuipers. 1989a. Qualitative reasoning: modeling and simulation with incomplete knowledge. *Automatica* **25**: 571-585.

Abstract

Recently developed methods for qualitative reasoning may fill an important gap in the modeling and control toolkit. Qualitative reasoning methods provide greater expressive power for states of incomplete knowledge than differential or difference equations, and thus make it possible to build models without incorporating assumptions of linearity or specific values for incompletely known constants. Even with incomplete knowledge, there is enough information in a qualitative description to support qualitative simulation, predicting the possible behaviors of an incompletely described system. We survey results from several approaches to qualitative reasoning, and provide a detailed example of the application of these methods to a simple problem. The mathematical validity of qualitative simulation is also assessed. Initial results have been encouraging, and steps are now being taken to develop additional mathematical power, hierarchical decomposition methods, and incremental quantitative constraints, to make qualitative reasoning into a formal reasoning method useful on realistic problems.

20. B. J. Kuipers. 1989b. Qualitative reasoning with causal models in diagnosis of complex systems. In L. Widman, K. Loparo, & N. Nielson (Eds.), *Artificial Intelligence, Simulation and Modeling.* New York: John Wiley & Sons, 1989, pp. 257-274.

Abstract

This chapter describes research that we have been doing in qualitative reasoning. The goal of this work is to understand the role that qualitative reasoning about the structure and behavior of mechanisms might play in medical diagnosis. Although the motivation for this work is medical diagnosis, some of the examples discussed are of simple physical systems, and we anticipate that our results will be applicable to a variety of nonmedical domains.

This chapter discusses the motivation for using qualitative causal models as a part of a diagnostic process. The nature of qualitative models is described and an example given of a qualitative model of a relatively simple medical mechanish, the water balance mechanism of the kidney. Finally, in somewhat more detail a recent development that shows promise of solving certain previously open problems in qualitative reasoning is discussed.

21. B. J. Kuipers. 1990. Simulation, Qualitative. In M. G. Singh (Ed.), *Systems & Control Encyclopedia, Supplementary Volume 1.* NY: Pergamon Press.

Abstract

Qualitative simulation is a method for predicting the possible qualitatively distinct behaviors of a system from an incomplete qualitative description of its structure. Where numerically-based simulation methods frequently require added assumptions, such as numerical parameter values and linear approximations to unknown or intractable functional relations, qualitative simulation can be applied to a qualitative description of values and relations provides a correspondingly weaker result. The result of a qualitative simulation is frequently a branching tree of possible behaviors. These methods are particularly valuable in situations characterized by incomplete knowledge such as prediction in biology, medicine, or economics, or in model-based diagnosis of unknown faults.

22. Benjamin Kuipers & Karl Åström. 1991. The composition of heterogeneous control laws. In *Proceedings of the American Control Conference*, 1991.

Abstract

To design a control system to operate over a wide range of conditions, it may be necessary to combine control laws which are appropriate to the different operating regions of the system. The fuzzy control literature, and industrial practice, provide certain non-linear methods for combining heterogeneous control laws, but these methods have been very difficult to analyze theoretically. We provide an alternate formulation and extension of this approach that has several practical and theoretical benefits. First, the elements to be combined are classical control laws, which provide high-resolution control and can be analyzed by classical methods. Second, operating regions are characterized by fuzzy set membership functions. The global heterogeneous control law is defined as the weighted average of the local control laws, where the weights are the values returned by the membership functions, thereby providing smooth transitions between regions. Third, the heterogeneous control system may be described by a qualitative differential equation, which allows it to be analyzed by qualitative simulation, even in the face of incomplete knowledge of the underlying system or the operating region membership functions. Examples of heterogeneous control laws are given for level control of a water tank and for motion control of a mobile robot, and several alternate analysis methods are presented.

23. B. Kuipers and D. Berleant. 1988. Using incomplete quantitative knowledge in qualitative reasoning. In *Proceedings of the National Conference on Artificial Intelligence (AAAI-88)*. Los Altos, CA: Morgan Kaufman.

Abstract

Incomplete knowledge of the structure of mechanisms is an important fact of life in reasoning, commonsense or expert, about the physical world. Qualitative simulation captures an important kind of incomplete, ordinal, knowledge, and predicts the set of qualitatively possible behaviors of a mechanism, given a qualitative description of its structure and initial state. However, one frequently has *quantiative* knowledge as well as qualitative, though seldom enough to specify a numerical simulation.

We present a method for incrementally exploiting incomplete quantitative knowledge, by using it to refine the predictions of a qualitative reasoner. Incomplete quantitative descriptions (currently ranges within which unknown values are assumed to lie) are asserted about some landmark values in the quantity spaces of qualitative parameters. Unknown monotonic function constraints may be bounded by numerically computable envelope functions. Implications are derived by local propagation across the constraints in the model.

When this refinement process produces a contradiction, a qualitatively plausible behavior is shown to conflict with the quantitative knowledge. When all predicted behaviors of a given model are contradicted, the model is refuted. If a behavior is not refuted, propagation of quantitative information results in a mixed quantitative/qualitative description of behavior that can be compared with other surviving predictions for differential diagnosis.

24. B. Kuipers and D. Berleant. 1990. A smooth integration of incomplete quantitative knowledge into qualitative simulation. UT AI TR 90-122.

Abstract

Qualitative and quantitative representations and inference methods provide alternate means for reasoning about the behavior of deterministic systems. The strength of qualitative reasoning is the ability to derive useful, though incomplete, conclusions from incomplete knowledge of the structure of a system. We show how quantitative information, even when very incomplete, can be integrated smoothly into the framework of qualitative reasoning.

Our algorithm, Q2, can draw more powerful conclusions than would be possible for a qualitative simulator alone, without sacrificing the expressive power and graceful degradation capabilities of qualitative simulation. Each qualitative behavior produced by QSIM implies a collection of algebraic equations defined over the terms appearing in the behavior description. In particular, landmark values are names for unknown real numbers, and so serve exactly as algebraic variables. Qualitatively distinct behaviors imply distinct sets of equations. The equations follow from the definitions of the qualitative constraints and fundamental theorems of the differential and integral calculus.

Incomplete knowledge of quantitative values, in the form of bounding intervals, can be propagated across the equations to produce either (a) a contradiction refuting the current qualitative behavior, or (b) a qualitative behavior description in which landmarks and other terms are annotated with quantitative ranges. We sketch the proof of soundness for Q2, discuss the use of mixed qualitative and quantitative reasoning for measurement interpretation, and present examples of model-based reasoning with QSIM and Q2 applied to diagnosis and design.

25. B. Kuipers and C. Chiu. 1987. Taming intractible branching in qualitative simulation. *Proceedings of the Tenth International Joint Conference on Artificial Intelligence (IJCAI-87).* Los Altos, CA: Morgan Kaufman.

Abstract

Qualitative simulation of behavior from structure is a valuable method for reasoning about partially known physical systems. Unfortunately, in many realistic situations, a qualitative description of structure is consistent with an intractibly large number of behavioral

predictions. We present two complementary methods, representing different trade-offs between generality and power, for taming an important case of intractible branching. The first method applies to the most general case of the problem. It changes the level of the behavioral description to aggregate an exponentially exploding tree of behaviors into a few distinct possibilities. The second method draws on additional mathematical knowledge, and assumptions about the smoothness of partially known functional relationships, to derive a correspondingly stronger result. Higher-order derivative constraints are automatically derived by manipulating the structural constraint model algebraically, and applied to eliminate impossible branches. These methods have been implemented as extensions to QSIM and tested on a substantial number of examples. They move us significantly closer to the goal of reasoning qualitatively about complex physical systems.

26. B. J. Kuipers, C. Chiu, D. T. Dalle Molle & D. R. Throop. 1991. Higher-order derivative constraints in qualitative simulation. *Artificial Intelligence* **51**: 343-379.

Abstract

Qualitative simulation is a useful method for predicting the possible qualitatively distinct behaviors of an incompletely known mechanism described by a system of qualitative differential equations (QDEs). Under some circumstances, sparse information about the derivatives of variables can lead to intractable branching (or "chatter") representing uninteresting or even spurious distinctions among qualitative behaviors. The problem of chatter stands in the way of real applications such as qualitative simulation of models in the design or diagnosis of engineered systems.

One solution to this problem is to exploit information about higher-order derivatives of the variables. We demonstrate automatic methods for identification of chattering variables, algebraic derivation of expressions for second-order derivatives, and evaluation and application of the sign of second- and third-order derivatives of variables, resulting in tractable simulation of important qualitative models.

Caution is required, however, when deriving higher-order derivative (HOD) expressions from models including incompletely known monotonic function (M^+) constraints, whose derivatives beyond the sign of the slope are completely unspecified. We discuss the strengths and weaknesses of several methods for evaluating HOD expressions in this situation.

We also discuss a second approach to intractable branching, in which we change the level of description to collapse an infinite set of distinct behaviors into a few by ignoring certain distinctions.

These two approaches represent a trade-off between generality and power. Each application of these methods can take a position on this trade-off depending on its own critical needs.

27. B. J. Kuipers and J. P. Kassirer. 1984. Causal reasoning in medicine: analysis of a protocol. *Cognitive Science* **8**: 363 - 385.

Abstract

The ability to identify and represent the knowledge that a human expert has about a particular domain is a key method in the creation of an expert computer system. The first part of this paper demonstrates a methodology for collecting and analyzing observations of experts at work, in order to find the conceptual framework used for the particular domain. The second part develops a representation for qualitative knowledge of the structure and behavior of a mechanism. The qualitative simulation, or envisionment, process is given a aualitative structural description of a mechanism and some initialization information, and produces a detailed description of the mechanism's behavior. The simulation process has been fully implemented, and its results are shown for a particular disease mechanisms in nephrology. This vertical slice of the construction of a cognitive model demonstrates an effective knowledge acquisition method for the purpose of determining the structure of the representation itself, not simply the content of the knowledge to be encoded in that representation. Most importantly, it demonstrates the interaction among constraints derived from the textbook knowledge of the domain, from observations of the human expert, and from the computational requirements of successful performance.

28. W. W. Lee and B. Kuipers. 1988. Non-intersection of trajectories in qualitative phase space: a global constraint for qualitative simulation. In *Proceedings of the National Conference on Artificial Intelligence (AAAI-88).* Los Altos, CA: Morgan Kaufman.

Abstract

The QSIM algorithm is useful for predicting the possible qualitative behaviors of a system, given a qualitative differential equation (QDE) describing its structure and an initial state. Although QSIM is guaranteed to predict all real possibilities, it may also predict spurious behaviors which, if uncontrolled, can lead to an intractably branching tree of behaviors. Prediction of spurious behaviors is due to an interaction between the qualitative level of description and the local state-to-state perspective on the behavior taken by the algorithm.

In this paper, we describe the *non-intersection* constraint, which embodies the requirement that a trajectory in phase space cannot intersect itself. We develop a criterion for applying it to all second order systems. It eliminates a major source of spurious predictions. Using it with the curvature constraint tightens simulation to the point where *system-specific* constraints can be applied more effectively. We demonstrate this on damped oscillatory systems with potentially nonlinear monotonic restoring force and damping terms. Its introduction represents significant progress towards tightening QSIM simulation.