Behavioral Aggregation Within Complex Situations: A Case Study Involving Dynamic Equilibria

Shankar A. Rajamoney and **Sang Hoe Koo**
Computer Science Department
University of Southern California
Los Angeles, CA 90089

Abstract

The analysis of large complex situations poses difficult problems for qualitative reasoning due to the complexity of reasoning from first principles and the proliferation of ambiguities. Abstraction is a promising solution to these problems. In this paper, we study a type of abstraction, *behavioral aggregation*—the process of grouping a set of individual entities that collectively behave as a unit. In particular, we show how to build aggregate models of situations involving dynamic equilibria and how to reason about their behavior. Finally, we demonstrate, through several examples, the benefits of reasoning at the aggregate level: a reduction in the complexity of reasoning and a compact, easily interpretable, description of the behavior.

Introduction

The analysis of large complex situations poses difficult problems for qualitative reasoning. Two major reasons are that, as the complexity and magnitude of the situation increases, a) the inherent complexity of reasoning from first principles rapidly increases and, b) the ambiguities due to the qualitative nature of the reasoning [de Kleer, 1979; Kuipers and Chiu, 1987] rapidly multiply, spawning a proliferation of possible behaviors.

Abstraction is a promising solution to both problems [Weld, 1986; Doyle, 1986; Kuipers and Chiu, 1987; Iwasaki and Bhandari, 1988; S. Addanki and Penberthy, 1989]. The two main types of abstraction are [Doyle, 1986]: *approximation abstractions* in which details are ignored or simplifying assumptions are made, and *aggregation abstractions* in which several entities are grouped together and treated as a single entity. Examples of these two types of abstraction are neglecting friction and an OR gate consisting of a collection of transistors. Abstractions may be applied to the model (eg. using a low-frequency transistor model to analyze an electronic circuit) or to the situation (eg. ignoring the presence of the moon when computing the motion of the earth around the sun). In addition, aggregation may be further classified into structural, behavioral, and functional aggregation based on the rationale for collecting the individual entities into a single unit. Examples of aggregates in each of these categories are a river and its tributaries, an oscillatory circuit, and an ADDER, respectively.

This paper studies the application of behavioral aggregation (of situations) to the analysis of large complex situations. We present a framework for developing such aggregates and describe a case study in which we show how to model and reason with aggregates formed from a particular type of behavior—dynamic equilibria. Finally, we demonstrate the simplification in reasoning and interpreting behavior due to the use of these aggregates in the qualitative simulation of several examples from the fluids domain.

Behavioral Aggregation

Behavioral aggregation is defined as the process by which a unit, called an *aggregate*, is formed from a group of individual entities which collectively behave as one. The glue that binds the individual entities together is not their individual interactions (which may differ drastically across the individuals, or may vary unpredictably with time or in response to interactions with external entities), but their sum or collective interactions which result in behavior that can be readily characterized over time or whose response to external interactions can be easily predicted.

We specify three criteria that behavioral aggregation must satisfy to produce useful aggregates that expedite reasoning:

1. *Composability*. Aggregates must be represented in a modular fashion that facilitates the formation of larger aggregates from several smaller aggregates and the original objects in the situation. Very large complicated situations can then be explored by building several layers of aggregates that gradually reduce the complexity of the situation until reasoning is practical.
2. *Uniformity*. Aggregates must be represented uniformly, preferably by adopting the representation of the individual objects in the situation. This further enhances composability, allows easy integration

of reasoning about aggregates and their interactions with other objects in the situation, and obviates the need to devise new reasoning methods or alter existing methods.

3. *Opacity.* The aggregate model must capture the knowledge required for reasoning at the aggregate level; for example, computing the aggregate behavior, determining the behavior when the aggregate is a part of a more complex situation, or identifying interactions with objects external to the aggregate. Reasoning at the aggregate level must not resort to examining the individual basic objects and their interactions since this negates a principal benefit of forming the aggregate—a reduction in reasoning complexity. Consequently, it must proceed independently of (while remaining consistent with) the reasoning at the level of the individual objects.

Aggregates that satisfy these requirements facilitate the analysis of large complex situations by: a) *Reducing the reasoner's burden.* Substituting a single aggregate for several objects and their interactions results in a much simpler situation. b) *Generating simpler output.* Ignoring the individual interactions and focusing only on the simpler and fewer aggregate interactions results in condensed and easily interpretable output. c) *Channeling detailed reasoning.* Dividing the situation into behaviorally distinct portions provides focused access when additional detailed information is requested about the internal behavior of the aggregate.

The qualitative analysis of a complex situation can now be split into two steps: a) the pre-analysis of the situation by an *aggregator*, working in accord with a qualitative simulator, to identify behavioral aggregates, and b) the qualitative analysis of the aggregated situation, consisting of the identified aggregates and the remaining objects, by a qualitative reasoner, operating normally, to obtain the output described in terms of the aggregate behavior. In this paper, we develop the above framework by focusing on the modeling and the reasoning about the behavior of a class of aggregates.[1]

Aggregating Dynamic Equilibria

An equilibrium state is one in which no changes are occurring. In a static equilibrium, no processes are active and hence there is no change. However, in a dynamic equilibrium, several processes are active; but, the combined effect of all the processes on every quantity is zero and hence there is no change. Dynamic equilibria form an important class of behaviors and are ubiquitous in natural and artificial systems. Examples of systems that exhibit such behavior or rely on such behavior for their successful operation are chemical reactions, the governor in automatic transmissions, cruise control, the price of a commodity determined by market forces, the nitrogen cycle in the ecosystem, and the population of animals in a natural habitat.

Consider the situation shown in Figure 1. Its behavior may be explained as follows:

> In each container, the solution and its vapor reaches a dynamic equilibrium in which the rate of evaporation of the solvent of the solution is exactly equal to the rate of condensation of the vapor. Left alone, each equilibrium situation is maintained and no changes are observed. However, the rate of evaporation is inversely proportional to the concentration of the solution;[2] therefore, the equilibrium rate of evaporation in the container containing the more concentrated solution will be less than that of the other. Since the amount of vapor and, consequently, the vapor pressure depend on the evaporation rate, the vapor pressure will also be less. This vapor pressure difference results in a flow of vapor. The flow constitutes an external disturbance to both the dynamic equilibria; each responds internally to compensate for the disturbance. The loss of vapor at the source of the flow is compensated by internal evaporation to generate vapor, while the excess vapor at the destination of the flow is compensated by internal condensation to consume vapor. These processes result in a gradual decrease in the concentration of the higher concentration solution and a gradual increase in the other solution. The vapor flow and the internal compensatory processes continue to be active until the vapor pressures become equal (which is when the concentrations of the two solutions become equal). In this new equilibrium state, the equilibrium rates of evaporation and condensation in both the containers are equal.

This explanation describes two aggregate systems, their interaction, the internal behavior of each aggregate, and the overall behavior of the system. Notice that as reasoning shifts to the aggregate level, the response of the aggregate to the external gas flow is explained internally in terms of compensation, and the behavior is described in terms of the changes to the equilibrium conditions of the aggregate.

The basic pieces of knowledge required to reason about such situations are an aggregate consisting of

[1]The pre-analysis of a complex situation to determine behavioral aggregates is a difficult and interesting problem that we have also addressed; however, space restrictions prevent us from providing a meaningful description of this problem. For the purposes of this paper, we assume that an intelligent user or another system identifies the aggregates.

[2]In a solution, some of the solute particles occupy a portion of the liquid surface, thereby depriving solvent particles of potential escape sites. Consequently, as the concentration increases the rate of evaporation decreases. In fact, this deep explanation illustrates behavioral aggregation applied to a model: the kinetic theory of matter describes the behavior of the individual particles and provides an explanation for this phenomenon. At the aggregate level, the phenomenon manifests itself as a drop in the evaporation rate with an increase in the concentration of the solution.

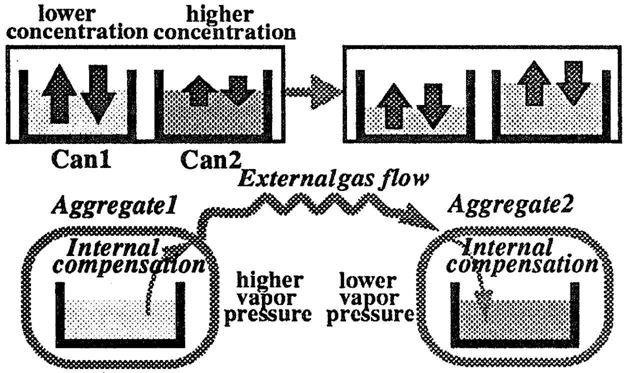

Figure 1: In this situation two containers with solutions of different concentrations are placed in an airtight container. The higher-concentration solution's level increases and that of the lower-concentration solution decreases.

the participating objects and their interactions, the external stresses that can affect the equilibrium, and the internal compensatory processes that tend to maintain the equilibrium. Briefly, reasoning about the behavior of a dynamic equilibrium situation involves determining the external stresses in the situation, activating the appropriate internal compensatory processes, and computing the net changes to the equilibrium. We first describe the representation of each of these elements and then more fully describe reasoning about complex situations involving dynamic equilibria. The representation is developed bearing in mind the three basic principles described earlier: composability, uniformity, and opacity. For concreteness, we adopt the Qualitative Process (QP) theory [Forbus, 1984] as our representation language and the Qualitative Process Engine (QPE) [Forbus, 1989] as our qualitative reasoner.

Modeling Dynamic Equilibria

An aggregate for a dynamic equilibrium includes features essential to reasoning about the changes to the equilibrium due to external influences. The aggregate consists of the objects and the processes participating in the equilibrium. If there are no external influences then the equilibrium is maintained and, consequently, no changes occur. If there are external influences then the behavior is computed by finding the combined changes due to the external processes and the internal compensatory processes. At the aggregate level, the equilibrium is assumed to be constantly maintained.[3] Consequently, the equilibrium processes are not required in reasoning and are accordingly suppressed. However, the aggregate must include information about the rates of the equilibrium processes in order to reason about the constraints that must be satisfied if the equilibrium is to be maintained and to reason about shifts in the equilibrium position. For the former, the relations that must be satisfied by the rates of the equilibrium processes to preserve equilibrium are required. For the latter, the relations describing the dependence of the rates of the equilibrium processes on other quantities are required. Both these pieces of information are collected by the aggregator when the equilibrium is identified: the former is obtained from the equilibrium condition and the latter from the equilibrium processes. Figure 2b shows the aggregate representation for the dynamic equilibrium between the condensation of the vapor and the evaporation of the water in the situation shown in Figure 2a. The equilibrium is maintained when the amounts and the derivatives of the two rates are equal. In addition, the equilibrium rate of condensation depends on the vapor pressure.[4]

Modeling External Stresses

An external stress on a dynamic equilibrium is produced when an external process introduces changes affecting the rates of the internal processes participating in the dynamic equilibrium in such a manner that the equilibrium is destroyed. All potential stresses on a dynamic equilibrium can be identified by the aggregator by determining which process instances influence (directly or indirectly) the rates of the equilibrium processes. The disturbance to the equilibrium produced by the external stress is qualitatively proportional to the rate of the external process. When the rate is zero the disturbance is also zero, and when the rate is positive the disturbance is also positive. Figure 2c shows the representation of an external stress. An external process, such as a flow of vapor from the container, would produce an external stress since it indirectly influences the equilibrium rate of condensation. In such a case, the stress would be activated with a disturbance proportional to the rate of the vapor flow.

Modeling Internal Compensatory Processes

When an external stress is imposed on a dynamic equilibrium, the system responds to minimize the effects of the stress by introducing internal compensatory processes. The influences of these processes are modeled on the equilibrium processes since, intuitively, the compensatory changes required at the aggregate level correspond to the net changes produced by the equilib-

[3] Hence we also assume stable dynamic equilibria.

[4] In this simple example, the equilibrium rate of evaporation, since it is not affected by any quantity, must remain constant. To maintain equilibrium, the rate of condensation must also remain constant. Accordingly, this relation constrains the vapor pressure to remain constant. In more difficult examples (for instance, Figure 1), both rates may change while preserving equilibrium, thereby, signaling shifts in the equilibrium position.

(a)

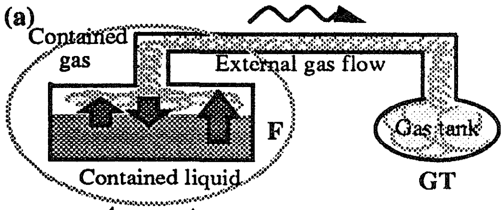

Aggregate

(b)

Individual-View Aggregates
Individuals
  cg a contained-gas
  cl a contained-liquid
  evap a process evaporation
  cond a process condensation
Quantityconditions
  (Active cg)
  (Active cl)
Relations
  Rate(evap) = Rate(cond)
  Rate(cond) $\propto_Q$ Vapor-pressure(cg)

(c)

Individual-View Negative-stress
Individuals
  ep an external-process
  agg an aggregate
Preconditions
  (Stresses-negatively ep agg)
Quantitiyconditions
  (Active ep)
  (Active agg)
Relations
  disturbance $\propto_Q$ Rate(ep)
  Correspondence((disturbance,ZERO), (Rate(ep),ZERO))

(d)

Process Evaporation-compensation
Individuals
  str a negative-stress
  ep an external-process
Quantityconditions
  (Active ep)
  (Active str)
Relations
  compensation-rate $\propto_Q$ disturbance(str)
Influences
  I- [Amount-of(cl),A(compensation-rate)]
  I+ [Amount-of(cg),A(compensation-rate)]

Figure 2: A dynamic equilibrium between the evaporation of a contained liquid and the condensation of its vapor, and the representation of the aggregate, an external stress and an internal compensatory process.

rium processes when disturbed from equilibrium at the detailed level. In the example shown in Figure 2a, a flow of vapor from the container disturbs the equilibrium by decreasing the amount of the vapor and hence, indirectly, the rate of condensation. Consequently, the changes observed are due to the excess evaporation. Hence, the internal compensatory process at the aggregate level must capture the same changes.

The type of compensation required depends on the nature of the stress: for example, a flow of vapor out of the container would require an internal compensation modeled on evaporation whereas a flow of vapor into the container would require an internal compensation modeled on condensation. The aggregator determines the nature of the stress and the type of compensation required by comparing the change produced by the stress to those produced by the equilibrium processes to determine which process tends to oppose the stress. The rate of the internal compensatory process is directly proportional to the disturbance produced by the external stress and, therefore, indirectly proportional to the rate of the external stress-producing process. As the disturbance decreases and reaches zero, the rate of the internal compensatory process also decreases and reaches zero (at which time the process instance becomes inactive).

Figure 2d shows the internal compensatory process modeled on evaporation that is activated by a corresponding stress produced by an external process. In the example situation, if a flow of vapor from the container is active, then a stress that activates the internal compensatory process modeled on evaporation is produced. The rate of the internal process depends on the disturbance which, in turn, depends on the rate of the flow of vapor. When the flow of vapor ceases to be active, the stress and the internal process also become inactive.

Reasoning with Aggregates

We adopt the Qualitative Process Engine (QPE) [Forbus, 1989] without any changes (the uniformity principle) to generate the envisionment of the aggregated situation. QPE generates the active aggregates, identifies external stresses, and introduces internal compensatory processes when required. The change to any quantity is the net change due to the influences of the external processes and the internal compensatory processes. The change in the equilibrium position is modeled by changes to the rates of the equilibrium processes. The aggregate-level envisionment describes the behavior of the equilibria. Below, we describe several examples of varying complexity, and show how reasoning with aggregates effectively curtails reasoning and compacts the overall envisionment.

A) Single Equilibrium, Single Stress. Our earlier example in Figure 2a shows a simple situation involving an equilibrium between evaporation and condensation. If the vapor pressure in the container is greater than that in the tank a flow of vapor out of the container occurs resulting in an external stress to the dynamic equilibrium. Figure 3a shows the envisionment generated for this situation and this aggregate-level envisionment may be compared with the detailed-level envisionment shown in Figure 3b.

B) Single Equilibrium, Multiple Stresses. When multiple external stresses are imposed on an equilibrium its behavior is the result of the combined effect of all the corresponding internal compensatory processes and the external processes. If the external stresses are of conflicting types then the resulting behavior is ambiguous and depends on which type dominates. QPE generates all the possible cases. Figure 4a shows a situation involving two conflicting stresses and a brief description of its aggregate-level envisionment.

C) Multiple Equilibria Connected by an External Process. Figure 4b shows the envisionment for the example described in the previous section. In this example, two similar dynamic equilibria are affected

(a)

state	s0	s1	s2
aggregate at F	Active	Active	Active
gas flow from F to GT	Inactive	Active	Inactive
evap-compensation at F	Inactive	Active	Inactive
gas flow from GT to F	Inactive	Inactive	Active
cond-compensation at F	Inactive	Inactive	Active
Ds[amount-of(liquid at F)]	0	-1	1
Ds[amount-of(gas at F)]	0	0	0

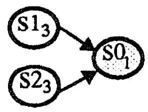

(b)

state	s0	s1	s2	s3	s4
gas flow from F to GT	I	A	I	I	I
gas flow from GT to F	I	I	A	I	I
evaporation at F	A	A	A	A	A
condensation at F	A	A	A	A	A
rate(evap) ?? rate(cond)	=			<	>

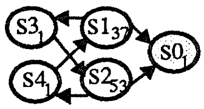

Figure 3: A brief description of the envisionment, at the aggregate and the detailed levels, of the situation shown in Figure 2a. (The subscripts indicate the number of sclasses collapsed into the sclass shown.)

by an external flow of vapor due to the difference in the equilibrium vapor pressures in the two containers. This flow results in a shift of equilibria until the equilibrium vapor pressures become equal.

D) Complex Equilibria. Figure 4c shows a situation in which a complex equilibrium consisting of two dynamic equilibria (evaporation and condensation, precipitation and dissolution) is affected by an external stress. The figure also shows the aggregate-level envisionment. This example illustrates the composability principle: the complex equilibrium is composed of two simpler equilibria and hence the situation can be analyzed at three levels.

Figure 4d shows the significant compaction of the detailed envisionment due to aggregation. These benefits are due to several factors including fewer processes (several equilibria processes are replaced by a few internal compensatory processes), fewer consistent combinations of processes (the internal processes are active only when the external stress-producing processes are active), and stronger constraints on the values and changes of quantities (from the equilibrium condition).

Related Work

Abstraction has been extensively investigated in several areas of AI: planning, design, diagnosis, and learning to name a few. Here, we discuss recent research on this topic in qualitative reasoning, concentrating on work that addresses behavioral aggregation. Closely related to our work is Kuipers' work on time-scale abstraction [Kuipers, 1987]. Kuipers shows how large complex situations involving widely separated time-scales can be decomposed into layers of equilibrium mechanisms in which a particular mechanism views a much faster mechanism as instantaneous and a much slower one as being constant. Our work is very much in this spirit and we extend Kuipers' work by build-

(a)

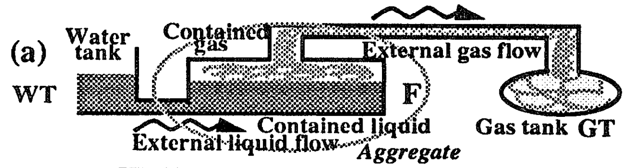

state	s0 ←	s1
aggregate at F	A	A
gas-flow from F to GT	I	A
liq-flow from WT to F	I	A
evap-compensation F	I	A
Ds[volume(cg at F)]	0	-1
Ds[amount-of(cg at F)]	0	-1
Ds[vapor-pressure(cg at F)]	0	0

(b)

state	s0 ←	s1
aggregate at Can1	A	A
aggregate at Can2	A	A
gas flow from Can1 to Can2	I	A
evap-compensation at Can1	I	A
cond-compensation at Can2	I	A
Ds[amount-of(gas at Can1)]	0	-1
Ds[amount-of(liquid at Can1)]	0	-1
Ds[concentration(Can1)]	0	1
Ds[amount-of(gas at Can2)]	0	1
Ds[amount-of(liquid at Can2)]	0	1
Ds[concentration(Can2)]	0	-1

(c)

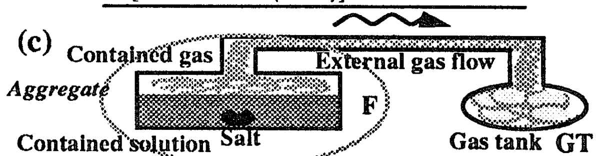

state	s0 ←	s1
evap/cond/dissolve/precipitate aggregate at F	A	A
gas flow from F to GT	I	A
evap-compensation at F	I	A
precipitate-compensation at F to increase solid salt	I	A
Ds[amount-of(liquid at F)]	0	-1
Ds[amount-of(dissolved-salt at F)}	0	-1
Ds[amount-of(solid-salt at F)}	0	1
Ds[concentration(solution F)]	0	0

(d)

	Detailed envisionment			Aggregates envisionment		
Examples	#of sclasses	#of transitions	CPUtime in sec.	#of sclasses	#of transitions	CPUtime in sec.
Fig. 2a	70	252	111.91	13	4	32.51
Fig. 4a	451	?	>1000	13	4	15.45
Fig. 1	>1000	?	>1000	19	162	69.43
Fig. 4c	328	?	>1000	16	60	39.88

Note: 1. CPU time is measured on SUN-4/490.
2. Experiments stopped after 1000 cpu seconds.
? indicates that limit analysis did not complete.

Figure 4: Various examples and brief descriptions of the aggregate-level envisionments. For simplicity, all the liquid flow and gas flow instances are assumed to be uni-directional and only a few relevant Ds values are shown.

ing aggregate models and by showing how reasoning about the same phenomenon can be conducted at different aggregate levels, each with a different perspective. Weld's work on the aggregation of cyclic behavior [Weld, 1986] is also closely related. Weld's system detects cyclic processes in a history and builds an aggregate continuous process description of this behavior. While we share the same broad goal of performing behavioral aggregation to summarize behavior and facilitate advanced reasoning, there are several differences. Some of the important ones are: we build aggregate models that are used for reasoning, we allow external influences on the aggregate, and we reason about changes to the aggregate behavior (as in shifting equilibrium). Iwasaki and Bhandari [Iwasaki and Bhandari, 1988] describe a method that manipulates a system of near-decomposable equations to find aggregate subsystems of equations. Our focus is on building qualitative descriptions of the aggregates, and it is not clear how their method, which involves several non-qualitative steps such as the transformation of numerical matrices, may be applied to our problem.

Recently, several researchers have addressed large-scale model aggregation. Rajamoney and Koo [Rajamoney and Koo, 1990] and Amador and Weld [Amador and Weld, 1990] describe models for representing microscopic particles and their interactions, while Liu and Farley [Liu and Farley, 1990] describe models for representing electrical behavior at the level of electrical charges and fields. Such multiple models of a domain at different levels of aggregation are often useful for reasoning about phenomena that cannot be satisfactorily explained by a single model. While our current work is related, our emphasis is on aggregating a situation and not a model, and we focus on a much smaller scale of aggregation. Conceivably, the aggregate models developed for a situation may be generalized and incorporated into the domain model; however, the aggregates developed for one situation need not be even remotely useful in another situation. Unless care is taken to selectively learn such aggregates, the performance of the system can degrade considerably.

Discussion

In this paper, we described how to build aggregate models of dynamic equilibria and how to reason about their behavior. Several other classes of behaviors exist: oscillation, cycles, and feedback to name a few. Our ongoing research consists of identifying and modeling such behavior, and building an aggregator which, equipped with a library of such behavior types, will decompose a complex situation into behaviorally characterizable aggregates that can be reasoned about much more easily at the aggregate level. Our future work will also address several related issues including the specification of the relationship between the aggregate-level and the detailed-level envisionments more formally, the development of methods for probing the detailed-level envisionment in a focused manner, and the investigation of the trade-offs involved in aggregation (for example, the loss of behavioral resolution with the gain of reasoning power). We believe that the incorporation of abstraction into qualitative reasoners is essential to the analysis of complex situations.

Acknowledgements

We thank Prasanta Bose, Hee-Youn Lee, and Nicolas Rouquette for helpful discussions. We are also grateful to Ken Forbus for making the QPE code available.

References

Amador, F. G. and Weld, D. S. 1990. Multi-level modeling of populations. In *Fourth International Workshop on Qualitative Physics.*

de Kleer, J. 1979. The origin and resolution of ambiguities in causal arguments. In *Proceedings of International Joint Conference on Artificial Intelligence.*

Doyle, Richard 1986. Constructing and refining causal explanations from an inconsistent domain theory. In *Proceedings of AAAI-86.*

Forbus, Ken 1984. Qualitative process theory. *Artificial Intelligence.*

Forbus, Ken 1989. *The Qualitative Process Engine.* Technical Report, University of Illinois, Urbana, IL.

Iwasaki, Y. and Bhandari, I. 1988. Formal basis for commonsense abstraction of dynamic systems. In *Proceedings of the Seventh National Conference on Artificial Intelligence.*

Kuipers, B. and Chiu, C. 1987. Taming intractable branching in qualitative simulation. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence.*

Kuipers, B. 1987. Abstraction by time-scales in qualitative simulation. In *Proceedings of AAAI-87.*

Liu, Z. and Farley, A. M. 1990. Shifting ontological perspectives in reasoning about physical systems. In *Proceedings of AAAI-90.*

Rajamoney, Shankar A. and Koo, Sang Hoe 1990. Qualitative reasoning with microscopic theories. In *Proceedings of AAAI-90.*

S. Addanki, R. Cremonini and Penberthy, J. S. 1989. Reasoning about assumptions in graphs of models. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence.*

Weld, D. S. 1986. The use of aggregation in qualitative simulation. *Artificial Intelligence.*