

Reasoning with multiple abstraction models

Yumi Iwasaki
Knowledge Systems Laboratory
Stanford University
701 Welch Road,
Palo Alto, CA 94304

submitted to: Fourth International Workshop on Qualitative Physics

Abstract

The problem of complexity has kept qualitative physics techniques from being applied to large real-world systems. Use of a hierarchy of abstract models is crucial for managing complexity. Several researchers have proposed ways to use an abstraction hierarchy of models to control the complexity of qualitative simulation [Falkenhainer & Forbus 88, Kuipers 87]. All the approaches proposed require models at pre-defined abstraction levels. Furthermore, the precise relations between different models are not explicitly defined, which makes it difficult to relate the conclusions drawn from different models to generate one coherent description of the behavior of the system as a whole. In this paper, we describe a scheme for generating models at abstraction levels appropriate for a given problem without requiring pre-defined set of abstract models. We also propose means for integrating behaviors produced from different abstraction models into one coherent description.

1. Introduction

When reasoning about the behavior of physical systems, having an appropriate model for a given reasoning goal is crucial. One important factor in deciding the appropriateness of a model is the grain size of the model. Unnecessary details in the model can make the analysis much more complicated than necessary or even impossible. For example, a model of a traveling train including its acceleration and deceleration ability is often unnecessary if the goal is estimating the amount of time the train takes to travel between two cities. Simply using the expected average speed of the train gives a good enough answer for most practical purposes. Recent work on qualitative physics has also shown that the amount of computation required to predict possible behaviors of a qualitative model grows exponentially with the number of variables in the model [Davis 87, Kuipers 86]. The only way even a moderately complex system can be simulated using qualitative reasoning techniques is by suppressing unnecessary details with the use of abstract models.

Use of abstraction hierarchy has been AI's standard answer to the problem of complexity for a long time [Simon 81, Sacerdoti 74, Hobbs 85, Friedland & Iwasaki 85, Falkenhainer & Forbus 88, Kuipers 87]. In all such work, there are a fixed number of pre-defined levels of abstraction, and an abstract model at each level must be carefully prepared by a system builder and given to the system.¹ Unfortunately, many systems with multiple abstraction models do not explicitly define what each level represents.² In other words, it is not clear what one is abstracting over when going from a fine model to a coarse model because they do not make explicit the precise relation between different levels. As a consequence, it is not always easy to decide when it is appropriate to reason at a given level, or how best to combine conclusions at different levels.

In this paper, we propose a scheme for formulating an abstract model that does not require pre-defined hierarchy of abstraction models. The goal is to clarify the notion of model abstraction and provide a way to define abstraction levels explicitly which would allow graceful integration of reasoning at multiple abstraction levels. Since the choice of the right abstraction level depends on the purpose of reasoning, i.e. what question about the behavior is one trying to answer, the approach described requires information about the user's goal in terms of the behavioral aspects of interests such as temporal scope, grain size, types of phenomena. The research described here is part of a larger effort towards constructing a flexible device modeling environment, which, given a representation of the physical structure of a device, can generate a model, analyze its behavior, and give an explanation of the behavior.

Section 2 outlines the model-based reasoning component of the device modeling environment. Section 3 illustrates some difficulties in reasoning with multiple models. We argue that it is important to specify explicitly the dimension along which a model is abstracted also to take the reasoning goals into consideration when generating abstract models. Section 4 describes our scheme for generating process models of appropriate granularity. Section 5 proposes two techniques for reasoning with a detailed model and an abstract model at the same time. Finally, Section 6 gives a summary and discusses problems yet to be solved to implement the approaches discussed in the paper.

¹The system by Falkenhainer and Forbus is an exception in that it generates a process model given a structure and a question.

²Except in Kuipers work, where abstraction is explicitly defined to be along temporal scale.

Model-based reasoning about device behavior

The central components of the device modeling environment (DME) are the model generation, simulation, and explanation modules. Given a description of the structure of a device and a question about its behavior, the system will generate a model of its behavior, simulate it, and generate an explanation that answers the question appropriately. The model generation module must, first, determine the appropriate level of abstraction to model the device. Then, it will formulate a model at the level, which is given to the simulation module to predict its behavior. Sometimes, the level selected initially must be changed or multiple models must be employed if it becomes apparent that more detailed or abstract behavior must be studied. In such cases, the system should be able to formulate different models. Finally, when the system finishes analyzing the behavior, conclusions drawn from different models must be integrated into a coherent explanation of the behavior of the system as a whole.

The model generation module of the system currently being implemented is based on Qualitative Process Theory [Forbus 84]. It has a library of physical processes, and given a structural description of a system, it detects active processes and generates constraint equations from them. In QPT, Forbus defines the concept of physical processes as "something that acts through time to cause changes". In our scheme, the concept of processes is extended to include the following;

steady-state process: A phenomenon that does not result in an observable change of state but that can be said to occur in the same sense as dynamic processes of QPT. For example, a steady current flow in a close circuit with a constant voltage source such as battery represents a steady-state process if the flow does not result in appreciable discharge of the battery.

instantaneous change: A process that happens over a very short period of time but that results in an appreciable change in the state. For example, opening or closing of a relay can be perceived as happening instantaneously for most purposes.

component function: Component functions can also be described as processes which activate when certain conditions are satisfied and cause changes in the state of the world according to some constraints.

The types of phenomena listed above are represented as processes because they can be represented and reasoned about largely in the same manner as QPT's dynamic processes. Furthermore, they actually represent the same physical phenomena as QPT processes at different granularity. The distinction between steady-state process, dynamic processes, and instantaneous changes is a matter of grain size. The same heat flow process can be regarded as a steady state process if the period of interests is relatively short and the heat source and sink capacities are large; as an instantaneous change that equates the source and sink temperature if the temporal grain size observation is very large; and as a dynamic process otherwise. This means that the same phenomena should be represented as instantaneous, dynamic, as well as steady-state process in the knowledge base, and we plan to do just this in our knowledge base of processes.

Alternatively, one could try to make the system generate these alternative representations of processes automatically from one representation. However, since the information content of the different representations is not equivalent, such conversion is not possible in general without additional information or assumptions. This is obvious from

the fact that one can obtain equilibrium equations from differential equations describing dynamic behavior of a system while one cannot derive correct differential equations from equilibrium equations without making some assumptions about how the system behaves when disturbed out of equilibrium³.

Once a process structure is determined, the model generation module produces a qualitative equation model by instantiating the constraints and influences associated with active processes as well as objects. This set of equations and a description of the initial state is given to QSIM [Kuipers 86] to predict the behavior.

In order for the system to generate models and to reason at appropriate levels of abstraction, there are several issues that must be addressed:

- How can one define abstraction levels in such a way that information about the goal can be used to select an appropriate level?
- How can one characterize different goals of modeling, i.e. the types of answers sought by the question, in a way that will help select an appropriate abstraction model?
- How can conclusions at one level be related to conclusions at another level?

The next section discusses the difficulties of reasoning with multiple models in more detail with an example.

3. Difficulty with reasoning with multiple models

We illustrate the difficulty of combining reasoning at different abstraction levels, using an example of a rechargeable, nickel-cadmium battery. When we are interested in the behavior of the battery over hours, there are two types of processes, charging and discharging, whose preconditions and effects are given below. C represents the amount of electrical charge currently stored in the battery, and C_{MAX} is the maximum amount that can be stored.

Charging-process :
 precondition : $C < C_{MAX}$
 effects : $dC/dt > 0$

Discharging-process :
 precondition : $0 < C$
 effects : $dC/dt < 0$

Let M_{NICD-0} denote the process model at this level of detail.
 $M_{NICD-0} = \{\text{Charging-process, Discharging-process}\}$

If we observe the behavior of the battery over a large number -- thousands -- of charge-discharge cycles, the maximum capacity of the battery slowly decreases. This phenomenon is called aging.

³ de Kleer and Brown derive dynamic behavior using confluence equations obtained by differentiating algebraic equations. However, they do so under the assumption that the system is quasi-static -- the relation represented by each equilibrium equation is always maintained [de Kleer & Brown 84].

Aging-process :

precondition :

A large number (> 1000) of discharge-charge cycles take place.effects : $dC_{MAX}/dt < 0$

Let MNICD-1 represent the model at this more abstract level, containing the aging process but not the individual instances of charging and discharging processes.

MNICD-1 = {Aging-process}

In MNICD-0, C_{MAX} is a constant and will remain so forever, while MNICD-1 predicts that C_{MAX} will decrease steadily. If we wish to use both models to predict the behavior from hour to hour over a period of several weeks, we must find a way to combine the conclusions from different models into a coherent explanation of the behavior. There are several causes for the difficulty:

- (1) There is no precise definition of what is meant by a more "abstract" model. This leads to the following two problems:
- (2) The scope of applicability of each model is not clear.
- (3) There is no way to compare and relate the conclusions from the two models based on the represented information.

In Following Section 3.1, we argue that the dimensions of abstraction must be specified before we can define levels clearly. Section 3.2 discusses the importance of reasoning goals for choosing abstraction dimensions and levels.

3.1 Dimensions of abstraction

One reason that it is not clear what each level in an abstraction hierarchy represents is that there are many dimensions along which a model can be abstracted, and often each "step up" in the hierarchy of models can involve abstraction along several dimensions though it is seldom explicitly stated as such. Some important dimensions are;

- | | |
|---------------|--|
| structural: | Abstraction by lumping together a group of components that are physically close. |
| functional: | Abstraction by lumping together a group of components that collectively achieve a distinct function. |
| temporal: | Abstraction by ignoring behavior over a short period of time. |
| quantitative: | Abstraction by ignoring small differences in variable values. |

These dimensions are not necessarily independent. For example, structural abstraction often resembles temporal abstraction because physical proximity tends to correspond to the speed of interaction between parts of a system. In defining abstraction, the first thing which must be determined is the primary dimension along which to abstract. In our effort to define abstraction clearly, we will initially concentrate on abstraction along one dimension, namely temporal. This means that abstraction levels of models will be defined in terms of their temporal grain size. The temporal grain size will be defined more precisely in Section ???. For now, we will just state that it is the unit of time such that any

change taking place over a time period smaller than the unit will be considered instantaneous. We will denote the temporal grain size of model M by T_M , where T_M is measured in log scale of seconds. In other words, if $T_M = n$, M is a model formulated by ignoring any time delay smaller than 10^n seconds. M describes behavior in terms of changes that can be observed at this temporal grain size. Any changes that take place in much less time will be considered instantaneous while any changes that take place over a much (orders of magnitude) longer period of time will be ignored in M .

3.2 Explicit representation of the reasoning goals

Any reasoning activity has some explicit or implicit purpose, in light of which the quality of the outcome must be judged relative to the goal. In the case of model formulation, the quality of a model must be judged relative to the particular types or aspects of behavior one wishes to study. Therefore, information about the user's goal should be used to decide what kind of model to formulate. In our scheme for model generation, the natural place for incorporating this information is while selecting the relevant set of processes. We will develop a simple language which a user can use to characterize his/her goal in terms of the types and aspects of behavior of interest. Some relevant characteristics are;

- types of question: For example, determining stability, comparative statics, dynamic transient behavior.
- types of phenomena of interests: For example, electrical, thermodynamic, structural, kinetic.
- precision required: The degree of quantitative precision required for the answer. What orders of magnitude change in variable values are negligible or significant?
- temporal grain size: The degree of temporal precision desired for the answer. Is one interested in the changes from second to second or from day to day?
- temporal scope: The length of the time period over which the behavior must be analyzed. Is one interested in the behavior over a few seconds or over years?

Eventually, the system should be able to determine heuristically the appropriate temporal grain size and scope for the model using these characteristics of the user's goal, if such information is not provided explicitly. For the purpose of this paper, we will assume that the desired temporal scope and the grain size as well as the quantitative grain size have already been determined.

4. Selecting a process model

Once the desired grain size and the temporal scope for a model and the quantitative grain sizes for variables are determined, the first step in formulating a behavioral model is to select the set of processes to be considered for inclusion in the model. For this, one must determine the grain size of each process. Let v_1 be the variable whose value is directly influenced by a dynamic process P . The temporal grain size of P denotes the time required for the change in v_1 caused by P to become non-negligible -- i.e. larger than the granularity required for v_1 . Let $s(v_1)$ represent the grain size of v_1 , and $r(P)$ be typical rate of change in v_1 caused by P . We can compute the approximate temporal grain size of P as $s(v_1)/r(P)$. Since it may not be possible to specify $r(P)$ precisely, the grain size is likely to be a range in orders of magnitude. For some processes such as traveling light, it is possible to know $r(p)$ precisely. For other processes, it cannot be determined a priori, because it depends on

other parameters. For example, heat conduction rate between two objects depends on their temperature difference. Such processes will be assigned $r(P)$ with a large interval. Thus, every process in the knowledge base will lie in some interval along the entire range of temporal grain sizes as shown in Figure 4-1.




Figure 4-1: Processes with temporal grain sizes.

Given this knowledge of processes associated with approximate temporal grain sizes, and the desired temporal and quantitative granularity for the problem, the system will generate a process model as follows:

1. For a given grain size n_0 , classify the processes into the following three categories;
 - slow: Processes whose grain sizes are much larger than n_0 .
 - medium: Processes whose grain sizes include n_0 .
 - fast: Processes whose grain sizes are much smaller than n_0 .

The processes in the slow category are ignored because they have no detectable effect within time period of interest. The processes in the fast categories are treated as instantaneous processes with no time delay. Thus, they can cause discontinuous changes in variable values. The processes in the fast and medium categories form the process space with respect to the grain size n_0 and denoted by $Sp(n_0)$.
2. From $Sp(n_0)$, find all the processes whose preconditions are satisfied with respect to the current state of the structure. Instantiate these processes.
3. Generate equations from the description of functional relations and effects associated with the instantiated processes.
4. Perform orders of magnitude reasoning [Raiman 86] to determine the approximate rate of changes of processes. Use this information to refine the temporal grain sizes of processes.
5. Repeat steps 1 through 4 with the refined grain size information until the model no longer changes.

The idea of classifying processes into the three categories is based on the theory of aggregation [Simon & Ando 61, Iwasaki & Bhandari 88]. By differentiating among long-term, short-term, and middle-term phenomena, attention can be directed to the dynamics of specific subsystems without dealing with the entire system at once, reducing the degree of

complexity one must deal with when reasoning about the behavior of a dynamic system. This procedure will produce an equation model of the desired temporal grain size in multiple iterations. In each iteration, the estimates of the grain sizes of the processes are improved to refine the model to fit the desired granularity. We believe this is similar to the way humans build models. People rarely build a right model the first time, but a model's failure to produce the expected behavior enables them to detect its inadequacies and to improve it subsequently.

5. Reasoning with multiple models

When the temporal scope of interest is large compared to the temporal grain size, it will be necessary to take into consideration processes whose temporal grain sizes differ greatly. When processes of significantly different granularity must be taken into consideration, it is better to create separate models of different grain sizes than to create one large model. Creating separate models, each consisting of processes of similar granularity, helps to keep each model small -- therefore, to keep the complexity in reasoning about the behavior down. However, having multiple models requires conclusions drawn from different models to be integrated into one coherent description of behavior.




Figure 5-1: Sliding block

Consider the situation, where a block is continuously sliding right and left on a surface along the x axis as shown in Figure 5-1. Suppose that at the level of abstraction 1, there are processes block-moving-right and block-moving-left:

Block-moving-right-process

effect: $dx/dt > 0$
granularity: 0

Block-moving-left-process

effect: $dx/dt < 0$
granularity: 0

Let B_1 be the process model containing these two processes.

$B_1 = \{\text{Block-moving-right-process}, \text{Block-moving-left-process}\}$

Suppose further that the surface of the table slowly wears out as the block moves back and forth many times. Thus, the block gradually sinks. This is represented in the following process;

Surface-wearing-out-process

precondition:

a large number ($> 1,000$) of block-moving-right and block-moving-left processes

effect:

$dy/dt < 0$

granularity:

4

Let B_4 be the model containing the surface-wearing-out processes.

$B_4 = \{\text{Surface-wearing-out-processs}\}$

If we are interested in the behavior at grain size 1 over a period of more than 10^4 seconds, we must use both models. The first model, B_1 , will predict that y stays constant for any length of time. While the second model, B_4 will predict that y decreases. The conclusions of the two models contradict each other. Unless we find means to communicate the conclusions at one level to other levels, the behavior predicted at each level will continue to diverge. In this example, the block in B_1 will continue to move back and forth at the same y coordinate forever, while the block in B_4 will continue to sink. In the following sections, we propose two methods for handling such interactions between quantities at different levels in order to allow graceful integration of reasoning at multiple levels.

5.1 Use of relative value measurement

One way to avoid contradictory conclusions about one quantity being drawn from different levels is to redefine the variables in the finer grain size model to represent the relative value with respect to the same variable in the coarse model. This is similar to the use of local coordinate systems in spatial reasoning. When describing the motion of a finger, it is easier to do so relative to the coordinate system attached to the hand. The movement of a hand with respect to the arm is described in the coordinate system attached to the arm. Thus, the movement of a finger with respect to the arm can be computed by combining the two descriptions.

In the example of the sliding block, we will first reason at level 1, which predicts $y = 0$ for all time. When we introduce B_4 , we can change the interpretation of y in B_1 . y in B_1 originally represented the coordinate position of the block with respect to some global y axis. Now, y in B_1 represents the displacement with respect to y in B_4 . If we need to know the y position of the block with respect to the global coordinate system, we can compute it from the values of y in B_4 and of y in B_1 , and its new interpretation. This technique is useful when processes in a finer model cause rapid fluctuation in the value of a variable, while in the long run the variable moves in the general direction determined by long-term processes.

5.2 Changing landmark values

Using relative value does not solve all the problems. Consider, again, the rechargeable battery example in Section 3. The detailed model will predict that C_{MAX} remains constant, while the more abstract model will predict the level decreases over many charge-discharge cycles. If we interpret the change in C_{MAX} in M_0 to be the change in the displacement of C_{MAX} in M_0 with respect to that in M_1 , and the value of C_{MAX} in M_0 to be C_{MAX} in M_1 , the predictions by the two models can be interpreted without contradiction. However, when the aging process of the battery finally causes C_{MAX} to become 0, the more detailed

model must detect that the conditions for the charging and discharging processes can no longer be satisfied.⁴

In the sliding block example, the long-term process, surface-wearing-out, did not interfere with the short-term, horizontal movement of the block. However, in other situations, changes caused by a long-term process can affect activity of short-term processes. This interference of long-term and short-term processes happened when a change caused by long-term processes invalidates an assumption about the ordering of landmark values implicit in the definitions of short-term processes. For example, the definition of charging process in Section 3 makes implicit assumption that $C_{MAX} > 0$, where C_{MAX} and 0 are both landmark values of C .

The quantities that act as landmark values at the finer level are treated as constants at the level. However, in the longer-term behavior, these quantities may change, altering the ordering among such landmark values. If the conditions for short-term processes implicitly assume certain ordinal relations among landmarks, whenever a change in the ordering is detected in a coarse model, the preconditions of short-term processes should be reexamined.

In order for this to take place, the following must be enforced:

1. In process definitions, one must make explicit the assumptions about the ordinal relations among symbolic constants (i.e. the landmark values) that appear in the preconditions.
2. When a coarser model treats as a variable a quantity which is a landmark value in a finer model, the quantity space of the variable must include the constants whose ordinal relations with the quantity are among the assumptions described in 1.

In the battery example, an implicit assumption in the precondition of the charging process in M_0 is that $C_{MAX} > 0$. The quantity space of C_{MAX} in M_1 should include 0. When C_{MAX} reaches 0, this fact can be detected by M_1 and communicated to M_0 , which can reexamine the preconditions of the charging process to discover that it can no longer be satisfied.

6. Summary and Discussion

This paper discussed difficulties involved in generating models of appropriate granularity. In order to limit the size of a model and the complexity of reasoning about its behavior, it is important to formulate a model that is focused on the behavior of interest instead of using a comprehensive, complex model. Therefore, it is essential to make use of the information about the user's goal in terms of the grain size and scope of the behavior of interest. The paper proposed an approach for generating a model at an appropriate temporal grain size given such information.

When the temporal scope of the behavior of interest is large, it becomes necessary to formulate and reason with multiple models each of different granularity. The paper discussed ways to integrate potentially conflicting conclusions about behavior drawn from such models into one consistent description of behavior.

⁴ This is basically the same problem as that of noticing interference from other processes Weld discusses in his work on aggregation. [Weld 86].

A number of problems remain to be solved before these approaches can be fully implemented. We do not have good answers to questions of how to automatically reformulate a process description so that a variable will represent a relative value instead of an absolute value or even of deciding when it is appropriate to do so. We hope to gain insights into these questions by experimenting with these approaches within the Device Modeling Environment. DME is currently being implemented by How Things Work project at Knowledge Systems Laboratory using CYC [Lenat & Guha 90] and QSIM [Kuipers 86].

The discussion in this paper focused on temporal dimension as the abstraction dimension, but it is only one of many possible dimensions along which an abstraction hierarchy can be constructed. In the future, we hope to generalize our approach by extending it to other dimensions.

REFERENCES

- [Davis 87] Davis, E.
Constraint Propagation with Interval Labels.
Artificial Intelligence 32(3), 1987.
- [de Kleer & Brown 84] de Kleer, J., and Brown, J. S.
Qualitative Physics Based on Confluences.
Artificial Intelligence 24, 1984.
- [Falkenhainer 88]
Falkenhainer, B. and Forbus, K.
Setting up Large-Scale Qualitative Models.
In *Proceedings of the Seventh National Conference on Artificial Intelligence*.
1988.
- [Forbus 84] Forbus, D. K.
Qualitative Process Theory.
Artificial Intelligence 24, 1984.
- [Friedland 85]
Friedland, P. and Iwasaki, Y.
The Concept and Implementation of Skeletal Plans.
Journal of Automated Reasoning 1(2), 1985.
- [Hobbs 85] Hobbs, J. R.
Granularity.
In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*. 1985.
- [Iwasaki & Bhandari 88] Iwasaki, Y. and Bhandari, I.
Formal Basis for Commonsense Abstraction of Dynamic Systems.
In *Proceedings of the National Conference on Artificial Intelligence*. 1988.
- [Kuipers 86] Kuipers, B.
Qualitative Simulation.
Artificial Intelligence 29, 1986.

- [Kuipers 87] Kuipers, B.
Abstraction by Time-Scale in Qualitative Simulation.
In *Proceedings of the Sixth National Conference on Artificial Intelligence*. 1987.
- [Raiman 86] Raiman O.
Order of Magnitude Reasoning.
In *Proceedings of the Fifth National Conference on Artificial Intelligence*. 1986.
- [Sacerdoti 74] Sacerdoti, E. D.
Planning in a hierarchy of abstraction spaces.
Artificial Intelligence 5, 1974.
- [Simon 81] Simon, H. A.
The Sciences of the Artificial.
MIT Press, 1981.
- [Simon and Ando 61]
Simon, H. A. and Ando, A.
Aggregation of Variables in Dynamic Systems.
Econometrica 29, 1961.
- [Weld 86] Weld D. S.
The Use of Aggregation in Causal Simulation.
Artificial Intelligence 30(1), 1986.