Approximation Operators in Distributed Modeling

Sui-ky Ringo Ling & Louis Steinberg
Department of Computer Science
Rutgers University
New Brunswick, NJ 08904
ling@cs.rutgers.edu, lou@cs.rutgers.edu

Abstract

Computer programs which do any task which requires reasoning about physical systems need to use models of those systems with varying accuracy/complexity tradeoffs. This paper describes an approach to model generation in the domain of heat transfer which is capable of producing models that vary greatly along this dimension. This approach is based on the law of conservation of energy, which provides a set of choices for the models in terms of "control volumes" and heat flows. These choices are made by using rules of thumb, which can be seen as instances of two reduction operators, delta-iso and dominance. Various rough models are used to estimate the physical parameters on which these rules depend. That is, the rough models are evaluated in the process of building more accurate ones. The application of these operators is only valid for a specific set of physically meaningful quantities; thus we are really reasoning about physics, not equations.

This method has been implemented in a running system, MSG.

Introduction

Computer programs which do design or diagnosis of physical systems, or indeed do any task which requires reasoning about such systems, often need a way to predict the behavior of the system given its structure, i.e., they need to model the physical parameters and processes of the system. Unfortunately, models which embody all the details of the system are often too complex to use tractably. Thus, we need approximate models which trade off accuracy for simplicity. Furthermore, in many cases there is no single model which is sufficient for the whole task; e.g., in a parameter design task it may be necessary to use a fast but approximate model to get a design that is close to right, and a more accurate but slower model to get to the final design. Since building models is such a key part of reasoning about physical artifacts, it is important that computer systems which do tasks like design be able to build (or select) models as needed. We have implemented a system for automatically building models for a particular class of problems in the domain of heat transfer, where quantitative algebraic, ordinary and *partial* differential equations are involved.

An important issue in modeling is that:

> A system must make a number of choices that determine the accuracy and complexity of the model it is building. How can it make these choices so as to minimize complexity but still (with reasonable likelihood) meet a specified accuracy?

The issue is complicated by the fact that many models in our domain do not have closed-form solutions. Consequently, it is difficult to a priori estimate the error of a model, and to use it to guide the choices of approximations. Our domain experts would rather use a number of heuristics to guide their decisions, which often result in a model close to a specified accuracy. Our system also make use of such heuristics.

Our analysis of these heuristics, which is the main focus of this paper, indicates that they are not just experiential knowledge. They are also based on the domain theory, and there is a structure on these heuristics. In particular:

- The heuristics are structured by the fun-

damental law of the domain, the conservation of energy. The law and the domain theory impose a sequence on choosing approximations.

- Semantically, the heuristics involves the spatial and temporal *variations* of behavior and material properties of artifacts. Syntactically, they can be seen as instances of delta-iso and dominance reduction operators[Falkenhainer, 1992].
- The heuristics base the choices of approximations on variations of behaviors of artifacts, which are then estimated by a set of rough models including order of magnitude reasoning about the heat fluxes.

The following sections will describe what the choices are, and how the heuristics make those choices. The section after that will describe MSG (Model Selection and Generation), a system that implements our approach. The last two sections discuss related work and give a summary.

What Are the Choices

Before we can discuss how the choices are made, we must describe what the choices are. That is what we will do in this section.

In order to understand what the choices are, it is first necessary to understand how the models are derived from the physics of heat flow. An equational model consists of a set of variables and a set of equations relating these variables. In the domain of heat transfer, all mathematical models are based on a fundamental physical law: the law of conservation of energy. This law says that the rate of change of energy stored within any bounded region of space is equal to the net heat flow into the region plus any internal heat generation within the region. But in order to apply this conservation law, it is necessary to first specify a set of *control volumes*, regions of space bounded by *control surfaces*. Then the variables represent the physical parameters of the control volumes such as temperature and energy. The equations include at least one governing equation, representing the interaction of energy and heat flows according to the conservation law, with terms modeling how each energy and flow depends on the physical parameters. There will also be other equations, if necessary, to represent other constraints such as initial and boundary conditions.

The approach of using control volume for formulating conservation laws is the common approach to model formulation in the domains of heat transfer and fluid mechanics. To create such a model, then, we must choose a set of control volumes, a set of energy processes to include, and a specific model for each process included.

The issue in choosing control volumes that has the greatest effect on the complexity of the resulting model is the choice between a *lumped* control volume and a *differential* one. A lumped control volume occupies a finite volume, and we represent a parameter (e.g. the temperature) of the entire volume by a single number. A differential control volume is infinitesimal - actually, we are modeling all the points of the region by a "representative" control volume that has parameters representing the actual xyz position of this volume in the region. There are also combined types, e.g. a control volume that is lumped in the x dimension and differential in y and z.

A lumped control volume leads to algebraic or ordinary differential equations (ODE's), while a differential control volume leads to partial differential equations (PDE's).

After a control volume is chosen, it is then possible to identify all the energy and heat flow processes acting on that volume, and choose which ones to include in the model. One such choice is whether to include heat storage processes in the model. Since these are the only processes whose representation mentions time explicitly, leaving them out results in a *steady state* model in which time does not appear as a separate dimension, while leaving them in results in a *transient* model.

Two other choices involves heat flow processes. In some cases it is possible to ignore heat flow in one or more dimensions (e.g. in x), which reduces the model from 3D to 2D or 1D, thus greatly reducing the cost of numerical solution. If the model does include a conductive heat flow, it is possible to ignore the fact that thermal conductivity of many materials actually depends slightly on the temperature of the material.

How Do We Make These Choices?

Given the kinds of choices outlined above, the question remains as to how we make them. First we will discuss the general approach, and then see how it is instantiated

in a number of specific decisions.

General Approach

Our informal observation of real engineers shows that they use a number of "rules of thumb" to make such decisions, and so our system does too. However the main point of this paper is the following: our analysis shows that these rules can be explained as deriving from two basic reduction operators:[1]

- **Delta-iso:** If $\frac{\Delta_v f}{f_{ref}} \leq \epsilon$ then we can replace $f(v)$ with $f(v_0)$ everywhere in our model.

  Here f is some function of parameter v. $\Delta_v f$ is the amount $f(v)$ varies as v varies over its range (i.e. $|\min f(v) - \max f(v)|$). If this is much smaller than some reference value for f, we can ignore the variation, and treat $f(v)$ as if it were independent of v, and equal to $f(v_0)$ where v_0 is some fixed value in the range of v.

- **Dominance:** If $\frac{A}{B} \leq \epsilon$ then we can replace $A + B$ with B.

Both of these operators say that if some value (A or the variation in f) is "small", then we can ignore the smaller value. One question this raises is "small compared to what?" In an artifact where temperatures vary by hundreds of degrees, a variation of one degree may be small, while in an artifact where temperatures vary by only a few degrees at most, a variation of one degree may be significant. For dominance the answer is straightforward - compare A with the rest of the sum it is being added to. For delta-iso, it turns out that the appropriate f_{ref} can be determined from knowledge of the domain (both analytic and empirical). It is often, but not always, f itself.

But if "small" means "smaller than this other value", we still have to say how *much* smaller, i.e., what ϵ is. The answer here is that ϵ depends on how accurate a model we need. In our domain, for instance, $\epsilon = .1$ generally leads to models with an accuracy of $\pm 3 - 5\%$. Values of ϵ for other accuracies can sometimes also be determined a priori from domain knowledge, but must often be set by trial and error - a model is built with a given ϵ, and its accuracy is evaluated by comparison with a better model. If necessary ϵ is adjusted, and we repeat. Note that such iteration is not a focus of our current work. Our system simply takes the ϵ's as input.

The final question is how we can determine the value of the ratio, i.e. $\frac{\Delta_v f}{f_{ref}}$ or $\frac{A}{B}$. As we will see, often f, A, and B are values such as temperatures which are not given as part of the problem statement. Indeed, often they are the unknowns which the model, once built, will be used to find. However, we are able to escape from this apparent circularity for two reasons. First of all, it is sometimes possible to give bounds on the *variations* of some expression, which is what we need for the delta-iso rule, without knowing the detailed *values* of the expression. For example, the temperature in an artifact with no internal heat generation will lie between the maximum and minimum temperature in the initial condition and current condition of its boundary environment. Thus, we can bound the variation of temperature given only information in the problem. Also, we sometimes have constraints on the *derivative* of a function, which can give us bounds on the variability. The second reason we escape from the apparent circularity is that when we are building a model at one level of accuracy, we can often base our modeling decisions on the results of a much *less* accurate model. Often a crude order-of-magnitude result can tell us enough. These less accurate models appear in the forms of domain specific compiled rules, sometimes using "dimensionless numbers", such as Biot number, which are numbers often used by engineeers to characterize a system.

Our current system only goes one level deep in this process of using a less accurate model to gather information needed in building a more accurate model, but we speculate that in general the process may be recursive, e.g. with decisions about a highly accurate model being based on results from an intermediate model, which in turn is built based on results from a still less accurate model, and so on. Furthermore, we have observed that the same process is also used by engineers in the domain of ship design.

[1] The names of these operator are derived from the terminology of Falkenhainer[Falkenhainer, 1992] - the dominance operator is exactly the same as Falkenhainer, and delta-iso is a variation on his iso operator.

Example Rules

Now let see how a number of example rules are instances of this general scheme.

Choice of control volume Choice of control volume is guided by the degree of temperature variation in spatial dimensions. The justification for choosing a differential control volume is to capture temperature variation in a region. If temperature is uniform, then a lumped model is as effective as a distributed model, with less computation cost.

For a solid whose boundary touches a fluid, the Biot number[Incropera and DeWitt, 1990] is used to estimate that temperature variation. It is a "dimensionless number", i.e. a ratio in which all physical dimensions such as length and time cancel out. Such numbers are common in these rules of thumb. The Biot number is derived from the balance of heat flux at the surface of an object. It is defined as

$$\frac{\Delta T_{solid}}{\Delta T_{boundary}} = \frac{h * L}{k} \equiv \text{Biot Number}$$

where h, k, L are convection coefficient, conductivity, and thickness of an object along one spatial dimension. The formula for Biot number is derived from an estimate of the ratio of internal thermal resistance of a solid to the boundary layer thermal resistance, but it can be shown that this ratio is also roughly the ratio of temperature drop across thickness of the solid to temperature drop across the boundary. Thus, a small Biot number value implies a small temperature variation within the solid. If it is less than the specified threshold value, ϵ, a lumped model is chosen.

This decision is an instance of the delta-iso operation, with $f(v) = T_{solid}(x)$ and $f_{ref} = \Delta T_{boundary}$.

Reduction of Transient to Steady State Reduction of transient to steady state is guided by the relative amount of temperature variation over time, defined as $\Delta_t T / T_{min}$, where $\Delta_t T$ is the difference in average temperature at the start and end of the time interval we are modeling, calculated by using an algebraic lumped model of heat flow, and T_{min} is the minimum temperature of the two values. If $\Delta_t T / T_{min} \leq \epsilon$, then T is assumed to be constant with respect to time t. Therefore, $\partial T / \partial t = 0$, and the term representing energy storage can be dropped out, resulting in a steady state equation. For example,

$$k\frac{\partial^2 T}{\partial x^2} = \rho C \frac{\partial T}{\partial t} \Longrightarrow k\frac{\partial^2 T}{\partial x^2} = 0$$

This decision is an instance of the delta-iso operation.

Spatial Dimension Reduction Reduction from a 3D to a 2D or 1D model is based on inferring that heat flow in some dimension, say x, is much smaller than in other dimensions. If so, we can drop the term for heat flow in x from the heat equation, and for example,

$$k\frac{\partial^2 T}{\partial x^2} + k\frac{\partial^2 T}{\partial y^2} = 0 \Longrightarrow k\frac{\partial^2 T}{\partial y^2} = 0$$

We estimate the ratio of heat flow in one dimension to heat flow in the other dimensions using a kind of order-of-magnitude reasoning; the resulting rule of thumb is based on the aspect ratio of the object, i.e. the ratios of its lengths in x, y, and z.

This decision is an instance of dominance operation, using the the dominance of heat flux to eliminate terms of second order partial derivative in equations. Since we drop the only term that mentions, say, x, this results in a drop of spatial dimension.

Non-linear to Linear Equations The heat equation in its full form is a *non-linear* PDE. If we can assume that thermal conductivity is independent of location, we can transform this into a *linear* PDE which is much easier to solve. For example,

$$\frac{\partial}{\partial y}(k\frac{\partial T}{\partial y}) = 0 \Longrightarrow k\frac{\partial^2 T}{\partial y^2} = 0$$

This simplification is based on estimating conductivity variation with respect to temperature. The conductivity variation is defined by $\Delta_T k / k_{min}$, where $\Delta_T k$ is difference of conductivity at maximum and minimum temperatures. If $\Delta_T k / k_{min} \leq \epsilon$, then the conductivity is assumed to be independent of temperature, and therefore, also independent of x, y, and z (assuming a homogeneous material), which justifies the transformation above.

This decision is an instance of delta-iso operation.

Implemented System

We will now discuss MSG, the system we have implemented for thermal modeling based on our approach.

The input to MSG is

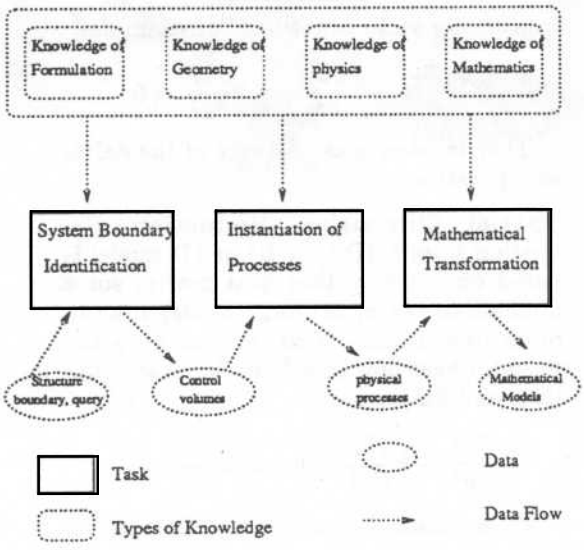

Figure 1: Tasks in Model Generation

- The physical system to model: a boundary representation of the objects' geometries, and their material properties, environment, and initial conditions.
- A query that a model is required to answer, expressed in terms of a variable, and its distribution across space and time. (E.g., $T(x,t)$ would ask how temperature varied over the x dimension and over time.)
- A set of thresholds ϵ for the approximation rules

The output is:

- A mathematical model satisfying the input query and the set of thresholds, and is expressed in terms of a set of variables, their attributes, and a set of equations, and
- A set of assumptions which the system makes in choices of control volumes, and energy terms.

The architecture of MSG follows the structure of the model building process outlined above. See Figure 1.

Status

The system described in this paper has been implemented. The first two subtasks, i.e. choice of control volumes and reasoning of energy and heat flows, are written in Common Lisp and CLOS, and the third subtask, mathematical simplification, is interfaced with Maple, a symbolic mathematical system.

The system is currently being tested for both single-component and multi-component objects. It has generated 12 examples, with 2 algebraic models, 2 ODE models and 8 PDE models (of which one of them involves 2 coupled PDE, and another one is a non-linear one). The examples cover heat treatment of flat plate, cooling fin problem, extrusion process, and oven. The types of heat flow include conduction and internal heat generation within solids, radiation and convection at surfaces of solids. The CPU time for these examples running on a Sparcstation 1 ranges from 1 second for an algebraic model to 4 seconds for a model of two coupled PDEs.

Discussion

This section will make several additional points.

Inferring Global Error from Local Error

Note that our models can be viewed either as systems of mathematical terms and equations or as systems of physical processes (temperatures, heat flows, etc), but that either way they are systems with parts that can interact in complex ways. We are concerned with the cost and accuracy of solving the whole system of equations, i.e. with determining the *overall* behavior of the physical system. However, our simplification methods are based on reasoning about the effect a given approximation will have on individual terms or equations, i.e. on the *local* behavior of *parts* of the system.

In general, one cannot simply assume that a bounded change to a part of a system of equations will lead to a bounded change in the overall solution. Furthermore, our domain experts tell us that the systems we are dealing with are too complex for them to be able to do much formal reasoning on how a local error affects the overall solution. So how can we justify our approach?

The answer is that we do not try to apply the operators to arbitrary terms in arbitrary expressions. Rather, we apply them only to very specific terms, representing specific physically meaningful quantities in a specific class of equations, modeling a particular kind of physical system. For

these particular quantities in these particular systems, experts in the domain know from experience that bounded local error leads to bounded global error, and in particular that a local error of 10% usually leads to a global error of 10% or less.

Thus, although the operators have been presented in mathematical form, they should not be seen as syntactic operations on equations, but rather as operations on specific physical parameters and processes, for which we have empirical validation that the operators lead to correct results. In fact, as can be seen from the discussion of MSG's architecture above, the rules of thumb derived from these operators are applied by MSG at a stage where the representation is still an explicitly physical one, before we have created the initial set of equations. This is consistent with our approach, discussed in [Ling and Steinberg, 1990], of reasoning as long as possible in a physical representation.

Independence and Symmetry

It is interesting to note that the kind of approximations our system makes are independence assumptions, that $f(v)$ is independent of v or that whatever physical term $A + B$ equals is independent of A.

In our context, this independence leads to simpler models in at least two ways. First of all, dropping a dependence on a dimension like x or t can lead to a problem of smaller dimensions, which is thus much more tractable numerically. Secondly, dropping of dependence can lead to a more decomposable problem. For instance, dropping the dependence of heat conductivity on temperature turns the equations into *linear* partial differential equations, which can be decomposed into separate solutions for each dimension.

Need for Validation and Revision of Models

It should be noted that there are several kinds of heuristics in this method. The validity of using the reduction operators in the places they are used and also the values of ϵ used are based on at least somewhat uncertain domain knowledge, and some of the methods used to evaluate the ratios in the rules are approximate. Because modeling decisions are based on heuristics that are not always correct, the model that results from our method, like the models that human engineers produce, may not in fact have the required accuracy. It is only a good first guess, and must be validated and, if necessary, revised. Our current work does not focus on this problem of validation and revision.

Related Work

Both Addanki *et al* [Addanki *et al.*, 1991] and Weld [Weld, 1990] [Weld, 1992] focus on model selection and switching. The work of Addanki *et al* uses graphs to represent a set of user-provided models, and uses domain specific parameter change rules to select models which resolve the conflicts between predictions and observation. Our approximation operators are used to construct models, rather than to select one from a set of user-provided models. Weld's work uses a domain independent technique of comparative analysis to select appropriate models from a set of user-provided models for query of the form "Is parameter X > parameter Y?". The forms of our approximation operators are domain independent, but the operators require domain specific methods to infer behaviors of objects. Our approximation operators can work on a more general form of queries, but cannot provide guarantee on the accuracy of approximate models. Finally, our operators cover not only terms eliminations in equations by fitting approximations, but also simplification of the overall form of equations by different choices of control volumes.

Nayak [Nayak, 1991] presents a domain-independent technique of validating an approximate model of a set of algebraic equations, by comparing the prediction of the model with those from an accurate model, and see if the discrepancy falls within a specified tolerance.

Falkenhainer [Falkenhainer, 1992] presents a method of building a set of likely useful models, based on input estimates of probability distribution of problem parameters. The models are then evaluated by actually executing them on a number of examples, rather than by the kind of a priori rules our system uses. An interpolation function is obtained for these accuracy. The interpolation function is then derived from these measurements and used to guide future model selection. The approach is promising when the cost of numerical evaluation can be amortized over much use of the resulting models, and when

the cost of evaluating the models is not too high, i.e. models are not too complex, and the number of likely-useful models is not too high.

Other relevant work are by Falkenhainer and Forbus [Falkenhainer and Forbus, 1991], by Finn *et al* [Finn *et al.*, 1992], by Nayak [Nayak, 1992], and by Gelsey [Gelsey, 1989]. This work, however, does not focus primarily on the issue of automatically choosing an approximation with a desired numerical accuracy.

Summary

Modeling is important in automated reasoning about complex physical devices. This paper describes an approach to model generation for algebraic, ordinary and partial differential equations in the domain of heat transfer. The structure of the conservation law, and the control volume formulation, provide a set of choices for the models in terms of control volumes, and heat flows. These choices are made by using rules of thumb, which can be seen as instances of two reduction operators, delta-iso and dominance. Various rough models are used to estimate the physical parameters on which these rules depend. That is, the rough models are evaluated in the process of building more accurate ones. The application of these operators is only valid for a specific set of physically meaningful quantities; thus we are really reasoning about physics, not equations.

This method has been implemented in a running system, MSG.

Acknowledgments

We would like to thank Saul Amarel, Tom Ellman, Andrew Gelsey, Yogesh Jaluria, Elaine Kant, Steve Lee, Chun Liew, Gerry Richter, Don Smith, Bob Vichnevetsky, R.S. Viswanath, Ke-thia Yao and other members of the CAP project for discussions and comments.

This work has been supported by NSF under Grant Number DMC-8610507, by the Rutgers Center for Computer Aids to Industrial Productivity, and by DARPA under Grant N00014-81-K-0394 and by DARPA and NASA under DARPA-funded NASA grant NAG2-645. The opinions expressed in this paper are those of the authors, and do not reflect any policies, either expressed or implied, of any granting agency.

References

S. Addanki, R. Cremonini, and J. S. Penberthy. Graphs of models. *Artificial Intelligence*, 51(1-3):145–177, 1991.

Brian Falkenhainer and Ken Forbus. Compositional modeling: Finding the right model for the job. *Artificial Intelligence*, 51(1-3):95–144, 1991.

Brian Falkenhainer. A Look at Idealization. In *Working notes of the AAAI-92 workshop on Approximation and Abstraction of Computational Theories*, 1992.

D. P. Finn, J. B. Grimson, and N. M. Harty. An intelligent modelling assistant for preliminary analysis in design. In *Proceedings of the Second International Conference on Artificial Intelligence in Design*, 1992.

Andrew Gelsey. Automated physical modeling. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, Detroit, Michigan USA, August 1989.

Frank P. Incropera and David P. DeWitt. *Fundamentals of Heat and Mass Transfer-3rd Ed.* John Wiley and Sons, 3rd edition, 1990.

Ringo Ling and Louis Steinberg. Automated Generation of Approximate models for Thermal Casting Processes-Some Preliminary Results. In *Working Notes of the AAAI-90 workshop on Automatic Generation of Approximations and Abstractions*, 1990.

Pandurang Nayak. Validating approximate equilibrium models. In *Proceedings of the AAAI-1991 Workshop on Model-Based Reasoning*, 1991.

Pandurang Nayak. Causal approximations. In *Proceedings of Tenth National Conference on Artificial Intelligence*, 1992.

Daniel S. Weld. Approximation reformulations. In *Proceedings of Eight National Conference on Artificial Intelligence*, pages 407–412, Cambridge, Ma, 1990. AAAI.

Daniel S. Weld. Reasoning about model accuracy. *Artificial Intelligence*, 56, 1992.