Towards Generality In Modelling Complex Physical Systems.

Conor P. McGann and Jane B. Grimson
Dept. of Computer Science,
Trinity College,
University of Dublin,
Dublin 2,
Ireland.

Tel: +353-1-772941
Fax: +353-1-772204
Email: CPMCGANN@VAX1.TCD.IE

Donal P. Finn
Hitachi Dublin Laboratory,
Hitachi Europe Ltd.,
O 'Reilly Institute,
Trinity College,
University of Dublin,
Dublin 2,
Ireland.

Tel: +353-1-6798911
Fax: +353-1-67988926
Email: DPFINN@VAX1.TCD.IE

Abstract

Qualitative physics has achieved considerable success in small and moderate scale problems but methods for managing the scope and depth of knowledge used by human problem solvers remain elusive. We argue that generality in a modelling approach is essential for extending qualitative physics to complex systems. A novel new model architecture is proposed which combines qualitative and quantitative knowledge in an ontology which draws from existing component and process based approaches. The significance of causal order in knowledge representation and inference is discussed and we introduce a new technique called *dynamic binding* which assigns causal order at run-time. We show how the *context* of a component is used to construct local perspectives of a complicated system and how this locality provides a unique demarcation in both knowledge representation and inference in the system architecture. Our mixed-model approach provides for physical modelling of engineering problems as a precursor to mathematical modelling or for interactive high-level querying by a non-expert user.

1. INTRODUCTION

Research in qualitative physics began with an ambitious agenda - to construct systems capable of reasoning about real-world problems in a manner which engineers and scientists seem to do so successfully [de Kleer and Weld '90a]. Qualitative physics has achieved considerable success in small and moderate scale problems [Kuipers '90] but methods for managing the scope and depth of knowledge used by human problem solvers remain elusive.

Extending qualitative physics to model real-world complex systems is closely bound to the idea of maximising generality in a modelling approach. This seems an intuitive observation: a model which is based as much as possible on commonalities will embrace a wider domain. Equally intuitive is the observation that generality and utility are often opposing characteristics. A model which is widely applicable will typically lack sufficient depth to support the level of detailed investigation required in most engineering applications.

1.1 Domain

Modelling of engineering problems can be divided into a number of stages; geometric modelling, physical modelling, mathematical modelling, numerical modelling and graphical modelling. The general application area of this research is in behavioural analysis of complex physical systems (such as in figure 1) or physical modelling. The position of this work in the overall modelling process is discussed in detail in [Mc Gann et al '91,'91a]. Physical modelling involves identifying and characterising any physical phenomena which are occurring in the system such as heat transfer, fluid flow or stress. Subsequent simplification and abstraction will begin with the physical model. When we speak of complexity in this domain we refer to the range of sub-domains that may interact and the number of physical components that may be included in the system.

1.2 Model Requirements

Any attempt to construct a model must begin with some 'performance requirements'. The first requirement of our model is that it must support a level of detail suitable for engineering applications. For physical modelling this constraint means we must be able to approximate the relative significance of phenomena distributed over a physical system. Secondly, we must be able to support qualitative system specification. This constraint is derived from the observation that engineers perform much initial analysis with symbolic specifications. A temperature may be represented as *atmospheric* or a pressure may be specified *moderate.* They require an approximate prediction of behaviour based on input such as this. The resulting physical model should be as accurate as possible. A completed behavioural analysis of the flat plate solar collector shown in figure 1 will identify heat exchange processes through radiation, conduction and convection throughout the system and stresses arising in the water pipes due to the expansion of the absorber plate. The answers to queries should be presented at the same level of abstraction as they are received. For example, if a user enquires if stress is significant in the water pipes, the system should respond with 'yes' or 'no' with some explanation of why rather than giving the results of a numerical simulation of stress occurring. A more detailed query about the thermal efficiency of the absorber plate will require a response indicating heat absorption from solar radiation and heat transfer to the water pipes.

1.3 Overview

This paper seeks to identify the criteria for maximising generality in a modelling approach which must satisfy the requirements outlined above. In section 2 we take a closer look at what generalisation means in terms of benefits accrued and design constraints imposed. Section 3 examines the ontological issues arising from these design constraints. We consider consider the work of others in the qualitative reasoning community in section 4 and draw on many ideas arising from this work in section 5 where we outline a causal network architecture based on a combination of implicit inference and quantitative knowledge at one level and explicit modelling at another. The central argument of this paper is that generality can be enhanced through close alignment of abstracted knowledge and traditional mathematical equations and through a dynamic, modular ontology. These and other conclusions are discussed in section 6.

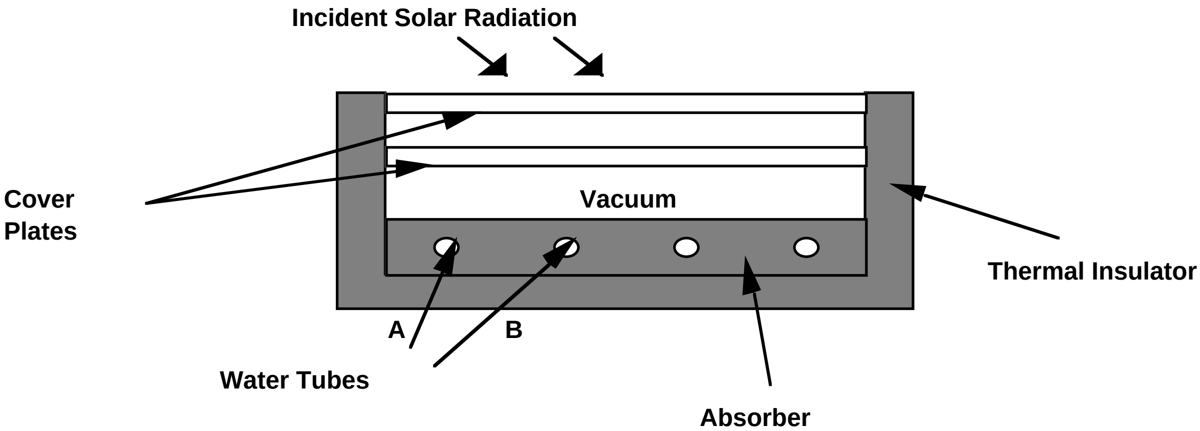

A Typical solar collector configuration is shown above. This is an example of a complex physical system including many sub-domains and (i.e. fluids, thermodynamics, stress) and many components. It consists of an absorbing surface that is thermally insulated on the edges and on the back side. Solar radiation is transmitted through the two glass cover plates before reaching the absorber. The collected energy is removed by circulating water through tubes that are in good thermal contact with the absorber.

Figure 1: Flat Plate Solar Collector

2 OBJECTIVES AND IMPLICATIONS OF GENERALISATION

Generalising a modelling approach is seen as an important step towards extending qualitative physics to model complex physical systems. But what are the benefits of generalisation which will facilitate this extension? To answer this question and position this research we must first look at current obstructions to capturing the scale and complexity of real engineering problems in a qualitative model.

2.1 Domain extensibility

Typically a complex physical system involves knowledge from a wide range of domains. Engineers and scientist can analyse such models because they can draw on *sufficiently* detailed information from a broad base of knowledge. The larger domain of qualitative physics is composed of many sub-domains e.g. thermodynamics, stress and fluid dynamics. In order to migrate to large-scale systems it becomes necessary to integrate such sub-domains which have, to date, been considered in isolation [Skorstad and Forbus '89], [Williams '84], [de Kleer '90], [Murthy and Addanki '87]. Composing an integrated knowledge base which aspires to represent a complete qualitative physics is an enormous task. However, it is an impossible task unless we establish criteria for common representation of knowledge. In this way, individual efforts in a particular sub-domain may contribute to the larger knowledge base as has been the pattern in traditional physics. This issue of domain extensibility rests on identifying a general format for knowledge representation to enable sub-domains to interact freely. Traditional physics achieves this through sharing common parameters. Closely coupled with this common representation is the need for an inference mechanism which supports the chosen knowledge representation and can cope with extensions to its knowledge base. So what are the implications of supporting domain extensibility?

We require a generalised knowledge representation scheme and an inference mechanism which spans many sub-domains. These criteria suggest the relevance of an abstracted quantitative rather than qualitative representation for knowledge. There are a number of reasons for this.

Qualitative values are problem-specific: Qualitative reasoning defines symbolic intervals to represent a continuous real-valued parameter value space. This delineation is typically related to the specific problem at hand. An example of this might be a symbolic value space for temperature. In a central heating system, corrective action may be based on temperature being too low or too high. In another application the temperature of a liquid may be used in reasoning about phase changes so we might base our value space on two thresholds: *freezing* and *boiling*. The objective of a symbolic representation is to link value and meaning in a single representation. The cost of this representation is to introduce explicit problem-specific information to a knowledge base.

Meaningful qualitative values require explicit interpretation: In a qualitative reasoning system where each parameter has its own quantity space the inference mechanism must reconcile different symbolic values. A fluid whose flow can be described as *laminar, turbulent or transitional* may be related to cross-sectional area which may be *small, medium or large.* Resolving the different representations to perform inference requires explicit interpretation of symbolic values. This explicit transformation conflicts with our objective of a general inference mechanism.

Quantitative Abstraction leads to consistent, objective knowledge acquisition: Traditional physics provides us with an objective, consistent representation for physical behaviour through the vehicle of mathematical equations and, in particular, differential equations. Current qualitative reasoning systems derive the abstract behaviour of a system by simulating hand crafted qualitative versions of the differential equations that characterise it [Sacks ‘90]. This abstraction process is, to a large extent, driven by the inability of symbolic calculus to adequately represent the structure of such equations while supporting qualitative values for parameters. Knowledge acquisition should attempt to parallel the structure of the equation as far as possible sacrificing only the required detail with which we represent values. Using quantitative abstractions should enable the knowledge acquisition process to rely very closely on existing, widely available knowledge found in the equations of traditional physics.

2.2 Modularity

Real engineering problems are often synonymous with large-scale complexity in terms of the number of components in a model. Consider the number of components in an aircraft frame or a PCB circuit board. Frequently there are interactions among many elements of a model. Behaviour is distributed. Qualitative reasoning systems have achieved considerable success in modelling small and even medium scale systems but the application of this technique to more complex problems has been impeded by the exponential relationship between the number of components in a system and its complexity. Any modelling approach which hopes to represent complex physical systems must break this exponential link. The complexity of a system should not be reflected in the architecture which models it. The value of a modular ontology is widely acknowledged [Forbus ‘90],[Davis ‘90] and is fundamental to achieving the composability which is required to cope with complex models. Generality in modelling physical systems relies on modularity in the modelling approach.

3. ONTOLOGICAL ISSUES

The ontology we choose to represent the entities in a physical system and the relationships between them reflect the fundamental abstractions of a real-world system As such, any aspirations of generality in modelling complex systems must begin with a framework suitable for multi-component, wide-domain engineering problems. Furthermore, the abstract model of a physical system provides a framework for subsequent analysis. It is critical that this abstraction provides sufficient utility to meet the requirements of the user. In devising an ontology appropriate to complex systems we must bear in mind the constraints imposed by model requirements (as outlined in section 1) and reconcile them with the need for domain extensibility and modularity. This section examines how generality can be enhanced through ontology and looks at how our model requirements constrain the level of generality we can achieve.

3.1 Implicit vs. Explicit Mechanisms

Traditional Physics does not embody the idea of mechanism. Instead it relies on our commonsense and understanding to invoke appropriate equations or models of behaviour [Forbus and Gentner '90]. Qualitative reasoning systems seek to merge the essential behaviour of traditional physics with a more formal idea of commonsense knowledge. Qualitative physics formalises the idea of mechanism implicitly or explicitly.

Explicit mechanism theories add mechanism as an explicitly defined agent of causation. All changes in the system are stipulated to occur directly or indirectly through this special agent. A process-based ontology is an example of an explicit mechanism theory common in qualitative reasoning systems - [Forbus '84], [Woods '91], [Hendrix '90] - where continuous or discrete processes act as the vehicles of change in the system. Examples of processes include heat conduction and fluid flow. Explicit mechanisms such as processes make explicit the causal direction of behaviour in the model.

Implicit mechanisms do not rely on a separate agent to cause change. Mechanism for change is implied in the relationships among parameters without any reference to external agents. An implicit mechanism for causality is typically synonymous with a non-directed organisational structure. Device-based ontologies [de Kleer and Brown '84] employ implicit mechanisms where the components in the system are described in terms of an external interface and a real-world system is modelled by connecting together collections of device models into a network. Changes arise as a consequence of components interacting with other parts of the network.

The contrasting method for representing mechanism in process-based and device based ontologies provides an important organisational distinction but how can we relate the expression of mechanism to the issues of generality?

Implicit mechanisms promote composability: Device-based ontologies focus on specifying local behaviour patterns and derive global behaviour through observing the interactions of the parts. A multi-component system is constructed through the introduction of self-contained elements to a network. An implicit representation of mechanisms enables this locality of behaviour to be integrated to a more complex system.

Explicit mechanisms make modelling assumptions explicit: The need to select different models based on certain operating and simplifying assumptions being valid is important for reasoning about complex models [Nayak et al. '90], [Forbus and Falkenheimer '91], [Davis '90]. Explicit mechanisms allow multiple models to be organised and invoked dynamically depending on the operating context of the model.

Explicit representation of causality is domain-specific: The 'directedness' of explicit representations clarifies causal order by making a clear distinction between dependent and independent variables. There are situations in modelling where such clarification is important e.g. justifying conclusions. However, the imposed selection of dependent and independent variables might not be appropriate in all domains. Consider ohm's law relating current, voltage and resistance: V = R * I. [Forbus and Gentner '90] make the point that this equation, which implicitly describes the relationship among three parameters in a non-directed interpretation, may be formalised to explicitly capture the intuitive use of the equation which typically holds R constant. If this directedness were explicitly encoded in the knowledge base it would render the equation inappropriate in a domain where temperature dependencies of materials lead to changes in resistivity. To make this intuition explicit in a knowledge base detracts from the essential physics of the equation.

From this discussion it is clear that both implicit and explicit mechanism theories have some contribution to make. Implicit mechanisms support the modularity required to migrate to multi-component systems. The non-directed nature of implicit mechanisms is an appealing representation for a generalised domain which seeks to parallel the representations of traditional physics. Explicit models are necessary to apply the sort of commonsense reasoning which engineers and scientist employ to abstract and simplify complex models. However, they tend to be more domain-specific than implicit representations and acquiring knowledge for explicit models is a task that is harder to generalise. Expertise in modelling is more subjective than the first principles equations common in traditional physics. The performance requirements of our model outlined in section 1 require some compromise to be made. We recognise the need for explicit mechanisms in a model capable of tackling complex systems but emphasise the need to minimise explicit representations and maximise the use of implicit mechanisms. The ontology we propose will incorporate this idea.

3.2 Dynamic Binding - an implicit mechanism for causal direction

The example of Ohm's law illustrates the loss of generality as a result of making causal direction explicit but we must accept the need for explicit mechanisms to determine causal order. A critical step in creating a complete qualitative physics relies on domain extensibility (section 2) which is compromised by the use of explicit knowledge. How can we reconcile these conflicting issues? The answer seems to be through *Dynamic binding* of explicit information. The relationship between dynamic binding and generality extends to many areas of computer science. In the context of modelling physical systems, it means specifying a causal order at run-time within implicit non-directed equations. The mechanism for inference is based on successive substitution to solve for unknown variables. The assertion of known and unknown variables becomes the principle mechanism for assigning causal order. Such assertions are interpreted based on scope and precedence of variables. For example, if we model heat transfer through a block, independent variables are asserted in the boundary conditions of the block. In one instance we may specify a temperature difference, the next time we may specify an input heat flux. The domain equations remain the same and the behavioural analysis can cope with both situations because the causal order is bound dynamically through the vehicle of scope (i.e. global scope has greater precedence than local scope). [Iwasaki and Simon '90] use exogeneous variables to resolve ambiguities in causal order. [de Kleer and Brown '90] point out the limitations of this approach when extended to multi-component systems where the system is an assembly of sub-systems and thus confusion arises about what is an exogeneous variable. We suggest that scoping of variables provides a mechanism to overcome this deficiency, dynamically binding causal order to parameters whose relationships are fundamentally non-directed.

3.3 Representation

Model Requirements outlined in section 1 stipulate the inclusion of qualitative knowledge both in terms of physical system specification and analysis. The further requirement that the model must support a level of detail suitable for engineering applications, coupled with the objective of domain extensibility (section 2), suggest a role for quantitative knowledge. The result is a hybrid of qualitative and quantitative knowledge. Users may specify initial structural behaviour qualitatively or quantitatively. A preprocessing of the structural description maps qualitative values into appropriate quantitative intervals. Subsequent simulation is based on approximate quantitative, rather than symbolic values. The results of simulation may be presented qualitatively or quantitatively. The clear demarcation between qualitative and quantitative knowledge confines domain-specific symbolic representations to higher levels of the model architecture.

4. RELATED WORK

Issues pertinent to modelling complex systems have received much attention in recent years in the qualitative reasoning community. Many ideas resulting from this effort are reflected in this research. The modularity we spoke about in section 2 is captured in the device-based ontology of [de Kleer and Brown '84] but the primitive components of such a network (i.e. a device) require a pre-defined functional specification inappropriate for modelling at the level of physical phenomena. Instead, network elements correspond to containers in which physical behaviour may be simulated. Modelling physical systems where behaviour is phenomenological calls for a network of contained spaces to represent geometrically distinct components which act as sites for phenomena to occur. Geometric and material properties of a component, together with the phenomena it contains, characterise the behaviour of this contained space. Hayes' contained-stuff ontology [Hayes '90] offers a suitable analogy to our 'spatially contained' phenomena. By adapting the device-based ontology to a network of 'contained-spaces' we get the benefits of a modular architecture but avoid commitment to functional specification of each node in model construction.

A process is an appealing metaphor for transport of physical phenomena from one contained space to another. In section 3 we point out the need for explicit mechanisms to enable the management and selection of an appropriate model based on the operating context of the model. A process-based ontology [Forbus '86] enables different process models to be invoked dynamically based on certain operating conditions being valid. [Forbus and Falkenheimer '91] build on the explicit mechanisms of QP theory to organise multiple models of varying levels of approximation and abstraction based on explicit simplifying and operating assumptions. Context dependent behaviours [Nayak et al '90] offer an interesting perspective on management of multiple models based on the structural context, behavioural context and expected behaviour. Both these approaches are geared towards the device being the principal unit of composition. In the 'contained-space' network described earlier, these approaches for organising multiple models can be adapted to select appropriate models for physical phenomena 'in-transit' through the network.

Knowledge representation provides a foundation for modelling. In section 2 we identified the relevance of abstracted quantitative knowledge to underpin an extensible knowledge base and support analysis of real-world engineering problems. [Sacks '90] describes the use of piecewise-linear approximations to abstract the differential equations at the core of traditional physics and couples this with qualitative knowledge to interpret simulation results. This approach is a step towards the generalised knowledge representation we describe. Piecewise -linear approximation is felt to be an appropriate knowledge acquisition technique for a generalised domain model.

[Davis ‘90] derives global behaviour through localised propagations. Layered paths of interaction are the media for propagating local changes. This ontology provides modularity and corresponds closely to context dependent transfer models which are the explicit agents of causality in the model. The reliance of this approach on the functional organisation of the model makes it an inappropriate ontology for modelling physical phenomena but the focus on locality and the idea of paths of interaction contribute to the evolution of our hybrid ontology which is described in the next section.

5. AN ARCHITECTURE FOR MODELLING COMPLEX PHYSICAL SYSTEMS

In this section we present an architecture for modelling complex physical systems. The complexity of the system is reflected in both the number of components and the range of phenomena that may interact. It is useful to consider the main elements of the discussion so far as they are the formative ideas of this architecture. First, locality is the relevant organising principle deriving global behaviour through the interactions of the parts.Second, we must maximise the use of implicit knowledge on the basis that it promotes composability. Third, non-directed, abstracted quantitative equations are the basis for a generalised knowledge base. Finally, explicit process models provide the medium for local propagations.

5.1 A Causal Network Ontology

A physical system comprising many physically distinct components (figure 1) may be conceptualised as a network of contained spaces . Figure 2 shows a partial view of this network. We use the term ‘component’ to represent a contained space. Network level inference must orchestrate the interactions of components. Network management is performed by a *Boundary Engine.* The behaviour of each component is context dependent. The context of a component is a local perspective of the system behaviour and is constructed by referring to the interactions between the component and its neighbours in the network. Local inference establishes the behaviour of the particular component in the current context. The context is presented to the component in terms of boundary conditions and asserted parameter values. Behaviour is derived through a straight forward method of iteratively solving equations with one unknown until no new information is derived. The simplification of the local reasoning process is based on the assumption that the boundary conditions represent the components perspective on the integrated physical system. This emphasises the importance of the network management mechanisms which must dynamically construct locally appropriate perspectives for each component in the network.

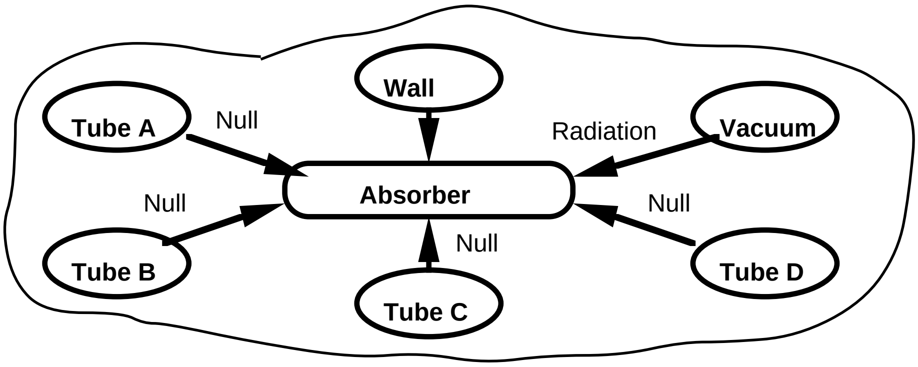

A partial view of the network representation of the collector configuration of figure 1 is shown above. The links indicate incoming transfer parameters to the context of the absorber. A radiation effect is transmitted through the vacuum to the absorbing plate. No other phenomena are currently under consideration.

Figure 2: Partial Network - Absorber Context

Links - a transport medium for local interactions. The boundary between two physically connected components is represented in a causal network through a link. Causality proceeds locally through these links. If a components local behaviour produces side-effects then these may be registered in each of its links as output parameters. At the receiving end, the information coming into a component through its links is a raw account of the interactions between that component and its neighbours. This information must be collated and interpreted before presenting it to the local component.

Transfer Models - constructing side effect scenarios. The modularity of this network representation relies on self-containment. Local contexts must be formed without reference to other components. Interactions must be fully captured in link transfer parameters. Such containment requires careful selection of transfer parameters when propagating side-effects. If you derive a local temperature increase should it be propagated as a surface temperature at the neighbour or as an input flux? Thermal expansion of a constrained component leads to stress. How can we represent this interaction? This problem is handled by the use of transfer model which introduce and initialise specially selected transfer parameters based on explicit model conditions. Heat transfer from one component to many can be represented by an equivalent electric circuit transfer model. Based on certain material and geometric properties the net flux to be transferred may be distributed over to all neighbouring components. The receiving component need refer only to its context which presents a heat influx at a boundary. This information is easily incorporated into the equation based inference performed locally.

The Boundary Engine - a context manager. Information contained in a number of links must be aggregated and interpreted to present the side-effect phenomena to the local component in a relevant format. A simple example of this might be summing a number of heat fluxes. The context of the component is determined by geometry, material properties and external interactions. The boundary engine must formulate a context for a component on this basis and present its conclusions to the component in terms of parameters which can be implicitly incorporated into local equations. The management of transfer models also falls to the boundary engine using context to select an appropriate model.

5.2 Dynamic Binding

A causal network ontology as outlined in section 5.1 is a dynamic modelling approach. Transfer models represent the interactions of the components of the system. Such models are created based on run-time contexts and may be modified or removed as model conditions change. Figure 3 hows the transfer model for heat transfer between the absorber plate and pipes from figure 1. This model assumes a perfect vacuum, that the boundary wall is a perfect insulator and steady state conditions. This approach enables the global model to evolve incrementally and iteratively. Assumptions and conditions may alter based on feedback from other components. In contrast to a device-based ontology, interactions between components are bound dynamically rather than predefined in the functional specification of a component.

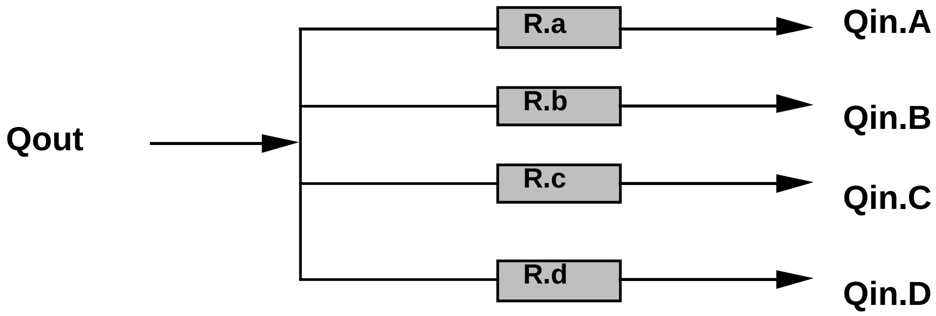

This diagram shows the equivalent electric circuit for the heat transfer model between the absorber and the water tubes. Heat transfer is by conduction. A surface pressure may also be derived due to possible expansion of the absorber as it heats up, thus exerting force on the water pipe. Transfer of heat from the pipe wall to the fluid will be via conduction and convection, depending on the spped of the water.

R.a = L/kA........L:pipe wall thickness; A: pipe surface area; k: thermal conductivity; h: convection coefficient.
Qin.A = Q.out * R.total/R.a
R.total = R.a * R.b * R.c * R.d/ (R.a + R.b + R.c + R.d)

Figure 3: Transfer Model - Equivalent Circuit

Causal order in the network is also determined dynamically rather than predirected. This is achieved through the use of scope to give precedence to parameters. There are four levels of scope which are given in descending order of precedence: *global,network,context and local.* Global scope applies where the parameter is exogeneous to the whole system. This is determined when the initial structural description is specified. Network scope applies to all transfer parameters in the links. Context scope is for parameters asserted as boundary conditions for a component. Material properties have a context scope. This reflects the fact that they may change based on context. Finally, local scope is the lowest in precedence and refers to parameters which are only valid locally. When the component goes out of scope, the local parameters revert to their unknown status. A greater permanence for a local parameter may be achieved by elevating it to a network parameter (i.e. output parameter) or a context parameter (e.g. a derived surface temperature). Such decisions are handled by the boundary engine before moving on to the next component in the network. This use of scope to implicitly establish a precedence hierarchy for parameters binds causal order at run-time and allows it to be changed as simulation proceeds. We confine making causal direction explicit to the manipulation of transfer model which are inherently directed seeking to determine output parameters based on certain explicit input conditions.

5.3 Knowledge Representation

There are two forms of knowledge representation incorporated into this architecture. Local knowledge is based on implicit, non-directed constraint equations. These equations are derived through quantitative abstraction of traditional equations of physics. A parameter may have an exact real numbered value or it may be represented approximately by a closed interval. No precedence is given to parameters in terms of knowledge representation so there are no explicit directional dependencies.

Knowledge at the network level is concerned with context management. Context management requires explicit representation of mechanism. Transfer models use explicit model conditions to produce values for identified transfer parameters. Compositional modelling [Forbus and Falkenheimer ‘91] offers an organisational structure which we can utilise to control the invocation of appropriate transfer models.

6. CONCLUSIONS

In this paper we have looked at the issues involved in modelling complex physical systems. In particular, we argue that generalisation in a model architecture is a critical step towards extending the application of qualitative physics to larger, more integrated problems. In complex systems, where the behaviour of many components integrates to give a global behaviour, composability of sub-domains and physical structure (i.e. number of components) is a necessary model attribute. Generalisation enables such composability. With this in mind we propose a new model architecture. Support for generality is provided through ontology in the form of a modular network representation and through a knowledge representation which closely reflects the physics of traditional mathematical domain equations. We point out that symbolic values and explicit mechanisms are problem specific characteristics of a model architecture. However, the suitability of explicit mechanisms in representing multiple models of interaction and the need for qualitative interpretation of behaviour dictate their inclusion in our model. By confining qualitative representations to higher levels of the model and making a clear distinction between implicit knowledge (within component) and explicit knowledge (context management) we achieve an adequate compromise.

The approach taken in this research is more of a ‘first principles’ method of deriving behaviour. There is a minimum of predefined function or structure. The concentration on phenomena rather than device implies a finer granularity than is appropriate in current qualitative reasoning systems. Applications of device-based ontologies assume that the external interface of the component can be predefined. We derive such interactions dynamically. Aggregating low-level behaviour to generate higher level interpretations will be an important area for future work. We advocate provision of an approximate quantitative model of complex systems at the level of physical phenomena as a basis for subsequent abstraction and approximation.

Initial implementation focusses on the sub-domains of stress, thermodynamics and fluid dynamics.

ACKNOWLEDGEMENTS

This work has been sponsored by the Hitachi Dublin Laboratory. We would like to thank Nobotushi Sagawa, Neil Hurley and Eamonn Cahill of the AI-DEQSOL group in HDL and Canice Lambe of the Knowledge and Data Engineering Group, Trinity College Dublin for various discussions and comments on earlier drafts of this paper.

REFERENCES

[Addanki et al. '89]: Sanjay Addanki, Roberto Cremonini, J. Scott Penberthy. Contexts: Dynamic Identification of Common Parameters in Distributed Analysis of Complex Devices. In *Proceedings of IJCAI '89.* International Joint Conference On Artificial Intelligence

[Davis '90]: Randall Davis. Diagnosis via Causal Reasoning: Paths of Interaction and the locality principle. *In [de Kleer and Weld '90]*

[de Kleer '90]: Johan de Kleer. *In [de Kleer and Weld '90]*

[de Kleer and Brown '83]: Johan de Kleer and John S. Brown. 1983. The Origin, Form and Logic of qualitative Physical Laws. In *Proceedings of IJCAI '83.*

[de Kleer and Brown '84]: Johan de Kleer and John S. Brown. 1984. A Qualitative Physics Based on Confluences. International Joint Conference On Artificial Intelligence, 1158 - 1169.

[de Kleer and Brown '90]: Johan de Kleer and John S. Brown. 1990. Theories of Causal Ordering. In *Readings in Qualitative Reasoning about Physical Systems* 631 - 645. Morgan Kaufmann Publishers.Artificial Intelligence

[de Kleer and Weld '90]: Johan de Kleer, Daniels S. Weld. *Readings IN Qualitative Reasoning About Physical Systems.* Morgan Kaufmann Publishers. 1990.

[de Kleer and Weld '90a]: Johan de Kleer, Daniels S. Weld. *Personal View in [de Kleer and Weld '90].*

[de Witt and Incropera '81]: D. P. de Witt and F.P Incropera, Fundamentals of Heat Transfer. Published by John Wiley & Sons.

[Forbus '84]: Kenneth D. Forbus. Qualitative Process Theory.
Artificial Intelligence. 24:85 -168

[Forbus '90]: Kenneth D. Forbus. Qualitative Physics: Past, Present and Future.*In [de Kleer and Weld '90]*

[Forbus and Falkenheimer '91]: B. Falkenheimer and K.D. Forbus. Compositional Modelling: Finding the right model for the job. In *Artificial Intelligence* 51:95 - 144

[Forbus and Gentner '90]: Kenneth D. Forbus and Dedre Gentner, *In [de Kleer and Weld '90]*

[Hayes '90]: Patrick Hayes. Naive Physics I: Ontology for Liquids. *In [de Kleer and Weld '90]*

[Hendrix '90]: Gary Hendrix. Modeling Simultaneous Actions and Continuous Processes.*In [de Kleer and Weld '90]*

[Iwasaki and Simon '90]: Yumi Iwasaki and Herbert A. Simon. 1990. Causality in Device Behaviour. In *Readings in Qualitative Reasoning about Physical Systems* 631 - 645. Morgan Kaufmann Publishers.

[Kuipers '86]: Benjamin Kuipers. Qualitative Simulation. In *Proceedings of IJCAI '86.* International Joint Conference on Artificial Intelligence.

[Kuipers and Chiu ’90]: Benjamin Kuipers and Charles Chiu. Taming Intractible Branching in Qualitative Simulatiion. *In [de Kleer and Weld ‘90]*

[McGann et al. ‘91]: Conor McGann, Jane Grimson and Donal Finn. A Conceptual Architecture for Modelling Physical Systems. In *Proceedings of IMACS ‘91.* World Congress on Computation and Applied Mathematics 3:929 - 934.

[McGann et al. ‘91a]: Conor McGann, Jane Grimson and Donal Finn. Modelling Physical Systems.*In Proceedings of Computer Graphics ‘91,* Alexander Palace, London.

[Murthy and Addanki ‘87]: Seshashayee S. Murthy, Sanjay Addanki. PROMPT: An Innovative Design Tool. In *Proceedings of AAAI ‘87*

[Nayak et al. ‘90]: P. Pandurang Nayak, Sanjaya Addanki and Leo Joskowicz. Modeling with Context Dependent Behaviours. Working Paper

[Sacks ‘90]: Elisha P. Sacks. Piecewise Linear Reasoning, *In [de Kleer and Weld ‘90]*

[Skorstad and Forbus ‘89]: G. Skorstad and K.D. Forbus. Qualitative and Quantitative Reasoning about thermodynamics. *In Proceedings of Eleventh Annula Comference of the Cognitive Science Society,* Ann Arbor, MI (1989) 892-899.

[Williams ‘84]: B.C. Williams, Qualitative Analysis of MOS Circuits, *In Artificial Intelligence.* 24 (1984) 281-346.

[Woods ‘91]: Erling A. Woods. The Hybrid Phenomena Theory. In *Proceedings of IJCAI ‘91.* International Joint Conference On Artificial Intelligence, 1138 - 1143