Context-Dependent Causal Explanations

Maria Lee* and Paul Compton†

* CSIRO Division of Information Technology, Locked Bag 17, North Ryde, NSW 2113, Australia, E-mail: lee@syd.dit.csiro.au

† School of Computer Science and Engineering, University of New South Wales, PO Box 1, Kensington, NSW 2033, Australia, E-mail: compton@cs.unsw.oz.au

Abstract

We outline a way of generating causal explanations from mathematical models. This approach is derived from the causal ordering theory of Iwasaki and Simon (Iwasaki and Simon, 1986a, 1986b; Iwasaki, 1988). Rather than produce a single causality by propagating causality from variables whose values are determined from outside the model, we allow causality to be propagated from variables within the model which are little influenced from within the model. This allows us to deal with wider range of models including systems with feedback, however, multiple causal explanations may then result. However, with propagation from the "most exogenous variables" a comparatively small number of explanation are produced which include those of interest to domain experts. We have applied this approach to large models including an environment greenhouse effect model. We suggest that the cost of a range of "plausible" models is small compared to the advantages of dealing with a wider range of model types.

INTRODUCTION

Computer simulation of complex systems based on mathematical models has long been an area of interest. With the emergence of high performance computers, simulation has come to play an increasingly important role and current models are very large. The area now even has a name, Computational Science. Simulation is applied to engineering design, scientific development and forecasting. For example environmental modelling is an increasingly critical activity. As knowledge based systems continue to expand in scope and application and their knowledge sources continue to diversify, proper linking of KBS and large mathematical models will become increasing important (Kowalik, 1986).

The link between mathematical models and KBS is problematic because these models are mathematical. Causal reasoning is a core method of reasoning about how physical systems work (Iwasaki, 1988). However, modern physics pays little attention to causality, and mathematics does not attempt to formalise it. One of the limits of the mathematical models is that they provide no explicit knowledge of how to perform analysis or to interpret results (Kunz *et al*, 1989). When we examine a single simulation output we cannot necessarily understand the factors involved. We have to perturb parameters or examine a range of behaviours or have an intimate knowledge of the behaviour of such mathematical equations. However, even a simple mathematical model can have very complex dynamic behaviour (May, 1976).

The interpretation of an equation or a diagram is highly context-dependent. Low-level graphical elements or abstract symbols do not have the precise meanings that words have in natural language. The symbols of x and y in x = y take on different meanings depending on the problem under consideration.

Researchers have worked on constructing causal explanations from mathematical equations. Forbus (Forbus, 1984) suggests that the causal reasoning of an equation should be fixed *a priori*. Iwasaki and Simon (Iwasaki and Simon, 1986, 1986a) assign a causal ordering to variables given only the equations and a list of which variables are *exogenous*. That is, the initial value is influenced from outside the system. Fixing the causal order a priori limits the behaviours generated, since different causal explanations are often possible.

According to Pearl and Verma (Pearl and Verma, 1991), the task of causal modelling can also be viewed as an identification game played by scientists against Nature. The notion of causality is context-dependent, which allows humans to decide on the direction of causality based on the circumstantial knowledge.

We propose a set of heuristics to transform equations to a suitable form to produce reasonable causal explanations. Our method produces a number of explanations. We use some simple heuristics to produce likely explanations from the way people seem to normally construct models. We further use information form the user and other sources to decide which particular models are of interest. In essence we find out the same sort of causality as the method of Iwasaki and Simon (Iwasaki and Simon, 1986a, 1986b). However in order to find out causality for a wide range of models, we have to use different assumptions about the structure of the models and consequently process them in a different way.

CAUSAL ORDERING THEORY

Causal Ordering (Iwasaki and Simon, 1986a, 1986b) is a technique for assigning an ordering to variables given only a set of equations and a list of which variables are exogenous. An exogenous variable is a variable that is influenced from outside the system directly and produces a change to other variables. That is it is a variable whose initial value is fixed by the user. Their approach is based on the theory of causal ordering first presented by Simon in 1952 (Simon, 1952).

The theory of causal ordering defines causal ordering as an asymmetric relation among variables in a set of simultaneous equations.

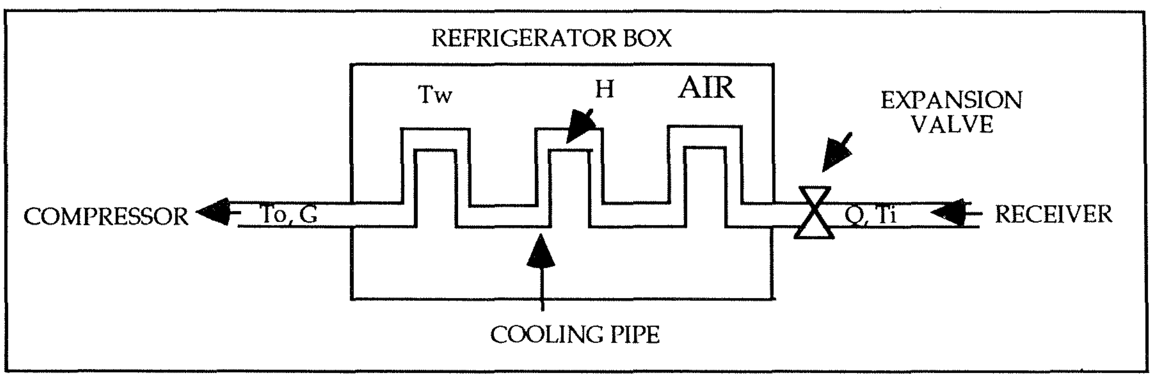

Figure 1. An Evaporator example

Establishing a causal ordering involves finding subsets of variables whose values can be computed independently of the remaining variables and then using those values to reduce the structure to a smaller set of equations containing only the remaining variables. We illustrate the causal ordering procedure by applying it to the evaporator example shown in Figure 1 (adapted from (Iwasaki, 1986a)).

The system is modelled by the equations of Figure 2a. The equations have the following interpretation (the constants are c_i's): (1) The rate of heat gained by the refrigerant, H, is proportional to the temperature difference. Tc is the condensing temperature and Tw is the temperature in the chamber. (2) The sum of the heat absorbed, H, and the energy of the incoming fluid is the energy of the outgoing fluid, where G is the ratio of vapour to total mass in the outgoing refrigerant, Ti and To are the temperatures of incoming and outgoing refrigerant. (3) The condensing temperature of the refrigerant is a monotonically increasing function (f) of the pressure, P. (4) The output temperature of the refrigerant is equal to the condensing temperature in the refrigerator chamber.

The causal ordering procedure assigns causal dependencies between variables by propagation through *self-contained* equations (Figure 2a-2d). Self-contained equations are a system of n equations with exactly n unknowns. Each matrix element is either blank or marked as a "1". A mark in row i column j means that variable x_j appears in equation i. Each row can have one or more marks. In order to make the system self-contained, the causal ordering needs four additional assumptions, in the form of additional equations. The four additional equations are:

$$Ti = c3, \quad (5)$$
$$Q = c4, \quad (6)$$
$$P = c5, \quad (7)$$
$$Tw = c6. \quad (8)$$

Each additional equation defines an exogenous variable, which provides a causal input to the phenomenon and is external to the evaporator. The causal ordering is derived from these exogenous variables and shows how causation propagates to other variables. By substituting the value for all the occurrences of variables, a new self-contained structure is obtained, until there are no more self-contained subsets. Figure 2b - 2c show the derived structures of higher orders. The variable in the minimal complete subset of the matrix is circled in the self-contained matrix. The final causal structure is shown in figure 2e.

In order for causal ordering to produce a "correct" causal structure, each equation must be a *structural equation*, i.e. it represents a conceptually distinct *mechanism* in the system. The term "mechanism" identifies physical processes as a kind of law and each equation is assigned to one mechanism. Iwasaki states that unfortunately there is no simple way to identify that an equation is structural (Iwasaki, 1986a, 1986b). Causal ordering assumes that equations used in the model are structural equations, and does not provide a method for transforming equations to structural equations. The causal ordering theory requires a self-contained structure to describe a system. To make a system self-contained and to assign exogenous variables relies upon an expert's experience and general knowledge of the model (Top and Akkermans, 1991). de Kleer and Brown (de Kleer and Brown, 1986) point out that the method of causal ordering specifies the same ordering for all behaviours. This is problematic as many systems have multiple modes of functioning, each characterised by its own distinct causal interaction. Iwasaki and Simon (Iwasaki and Simon, 1993) declare that the causal ordering theory is not

1. H = c1(Tw - Tc)
2. H = Q(c2G - (Ti - To))
3. Tc = f(P)
4. To = Tc
5. Ti = c3
6. Q = c4
7. P = c5
8. Tw = c6

(a)

	G	To	Tc	H
1			1	1
2	1	1		1
3			1	
4		1	1	

(b)

	G	To	H
1			1
2	1	1	1
4		1	

(c)

	G
2	1

(d)

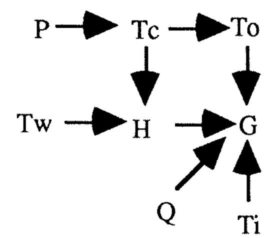

(e)

Figure 2. Equations for evaporator and the derived structure of causal ordering

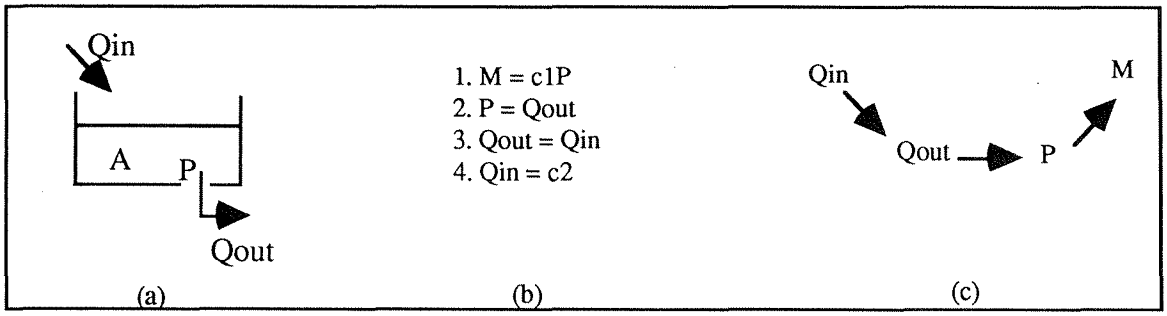

Figure 3. A steady State Bathtub Example

sufficiently developed to interpret all possible causal directions. In particular causal ordering cannot deal with feedback. The following section explains the problem of the causal ordering.

CAUSAL STABILITY AND CAUSAL CONSISTENCY

Electrical engineers use Ohm's law, and Kirchhoff's voltage and current laws to describe the fundamental relations among voltage, current and resistance in a circuit. The laws are presented as algebraic equations, which can be manipulated (e.g. V = IR, I = V/R, R = V/I). The equation V = IR represents electrical conduction in a resistor, where a voltage *V* volts produces a current of *I* amps through a resistor *R* ohms. Depending on the context, the equation can be causally explained as either the voltage *V* is causally dependent on current *I* and resistor *R* or current *I* on voltage *V* and resistance *R*. The third alternative, resistor R is causally dependent on voltage V and current I, does not make sense (Nayak, 1992). There is no set way of looking at it, engineers must be able to think about it in all possible ways, but only some of which make sense. Sometimes people use the following triangle to help remember the three forms of the formula:

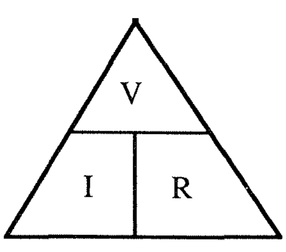

No one form of this equation is used more than the others.

White and Frederiksen (White and Frederiksen, 1990) state that the problem-solving process that students are taught does not necessarily facilitate an understanding of the physical system under study. Hence their view that qualitative theories are not consistent concerning basic causal relations between voltage, current, and resistance. They argue that our mental models should be consistent in the assumed direction of causality among resistance, voltage, and current in a circuit example. However many mental models have no mapping to the physical world - hence the mental models won't have the same sort of causality (in contrast to many biological models where the causality comes first (Feldman and Compton, 1989)). We state intuitively that adding water (Qin) to a bathtub (Figure 3a), increases the mass (M) of water and increases the pressure (P), which in turn increases the output flow rate (Qout). However, there are a lot of things in the physical world which are NOT intuitive - and are often counter-intuitive. The causal structure of the steady state bathtub may seem counter-intuitive (figure 3c). It shows that the output flow rate directly depends on the input flow rate, the pressure depends on the output flow rate, and the mass of water depends on the pressure. The idea of current leading voltage is another counter-intuitive example.

Skorstad (Skorstad, 1992) states that one of the limitations of the causal ordering theory of Iwasaki and Simon (Iwasaki and Simon, 1986, 1986a) is *context sensitivity*. He argued that the causal dependencies produced by the causal ordering theory may change depending on the context or scenario in which the underlying physical system

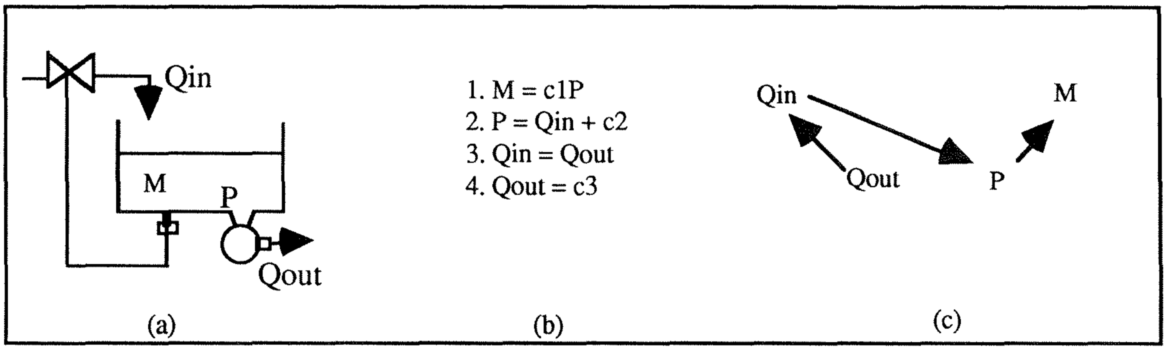

Figure 4. An example of Bathtub drain attached to a pump.

operates. Causal ordering uses a set of exogenous variables to place a system in different situations, which may change the interactions between the system and its environment. Thus, this restricts the qualitative modeller by providing a fixed causal interpretation.

Skorstad (Skorstad, 1992) defines the meaning of *causal stability* "A set of algebraic equations at a particular modelling viewpoint is causally stable if and only if its causal ordering is invariant with respect to its scenario space. Such a set of equations is unidirectional with respect to the modelling viewpoint". A modelling viewpoint means that the modeller makes decisions about ontology, perspective, and assumptions when conceptualising the phenomenon. A scenario space is a set of possible situations which are consistent with the equations.

Figure 4 shows a bathtub in a steady state where a bathtub drain is attached to a pump and the input stream is attached to a control valve (adapted from (Skorstad, 1992)). Skorstad notes that the context sensitive equations in the bathtub example are:

Qin = Qout, (equation 3 of figure 4b)
Qout = Qin. (equation 3 of figure 3b)

The causal dependency of the above equations changes depending on the circumstances. In the simple scenario of figure 3a, the output flow rate Qout is causally dependent on the input flow rate Qin. However, in figure 4a the output flow rate Qout has become exogenous. The input stream is attached to a control valve, thus the input flow rate is no longer independent of the system and cannot be treated as exogenous. Skorstad argued that if an equation is unidirectional with respect to the modelling viewpoint, then the equation is causal stable equation, such as:

P = c1M. (equation 1 of figure 1b and 2b)

de Kleer and Brown (de Kleer and Brown, 1986) state that ambiguity is the single key advantage in qualitative causal analysis. It is not necessary to have a unique solution in the n independent equations with n unknowns. In fact, unique solutions occur only rarely. Thus each solution potentially reflects a different global functioning with a distinct causality. Most systems are indeterminate. Therefore qualitative causal reasoning should be able to interpret all its possible behaviours. This is crucial for using the model to explain a physical system's operation.

CONTEXT-DEPENDENT CAUSAL EXPLANATIONS

In this section, we propose an approach to overcome some of the limitations of the causal ordering theory. We use equations that are a finite set of simultaneous equations and are from a mathematical model that describes the dynamic behaviour of the system. We identify variable dependencies first, then restructure equations to be asymmetric causal equation. We use the term "asymmetric causal equation" instead of "structural equation", because a structural equation should express the "real" causality; our equation are structural-like, but express "reasonable causality" (Lee *et al* , 1992a, 1992b). Causality can then be explicitly represented in asymmetric causal equations. If the dependency of a variable can not be fully specified, then multiple plausible causal behaviours are generated.

An asymmetric causal equation is an equation which can be understood as containing independent and dependent variables. This asymmetry manifests itself in that variables on the LHS are dependent on the variables on the RHS. Hence:

(a) The output variable appears on the LHS of the equation, it is the variable whose behaviour is of interest.
(b) The input or independent variables appear on the RHS of the equation in the model.
(c) Dependent variable occurring once and only once on the LHS of the equation in the model.
(d) For a differential equation, the derivative is on the LHS of the equation, with only one derivative in each equation.

important variables will be used frequently in the model, while other variables are added to fill in the gaps. In contrast to the causal ordering theory by propagating from exogenous, we deal with the causal influences from these "lesser" variables or "least caused" variables first. Once the greatest number of least occurring variable are chosen. The equation is manipulated so that the least occurring variables are on the RHS in the model. The remaining of the equations are then manipulated to give an asymmetric causal form. The first equation is now considered fixed and the process repeated for the rest of the equations etc.

The identification of appropriate causal explanations from equations is highly dependent on the problem under consideration. As the need for second-generation expert systems are to express

Equations	Asymmetric Causal Equations
1. H = c1(Tw - Tc)	1. H =[Tw, Tc]
2. H = Q(c2G - (Ti - To))	2. G = [H, Ti, To, Q]
3. Tc = f(P)	3. Tc = [P]
4. To = Tc	4. To = [Tc]
(a)	(b)

Causal Graphs

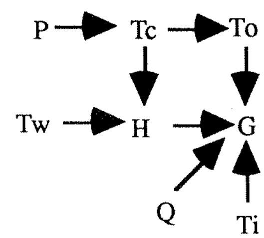

(c)

Figure 5. Equations for evaporator and the derived structure of our method.

To reduce the number of causal alternatives in feedback system loops, an important heuristic is to reorganise equations by propagating causality from the parameters which are "the least caused". Those variables are independent variables, which are set by the user. We look rather for the "most" independent variable. We start from the equation where there are the greatest number of such variables. We hypothesise that this has something to do with parsimony in scientific explanation. Model builders often want to discover (create) the smallest number of entities and causal connections to explain the behaviour of a system. More

how the things work; how different mechanisms interact; and to *explain* evidence in terms of structure and behaviour (Kuipers and Williams, 1988). In order to explain the evaporator internal behaviour of the total mass of the outgoing refrigerant, G, and the outgoing temperature, To, our system propagate variables within the equations, which come from an evaporator mathematical model (figure 5a), and generates the asymmetric causal equations (figure 5b) of the evaporator model. In figure 5b, circles show the variables of interests. Once the system constructs this asymmetric causal structure for the model, the causal graph is from the RHS of variables direct to the LHS of variables (figure 5c).

In order to express the different phenomena interact at pressure within chamber, P, and the total mass of outgoing refrigerant, G, multiple causal asymmetric causal equations are generated (figure 6b). Based on the multiple asymmetric causal equations, the equations could produce all plausible causal directions. In this evaporator example, the context sensitive equations are:

To = Tc, (equation 4 of figure 6.2a)
Tc = To. (equation 4 of figure 6.2b)

Its causal dependency varies depending on its situational context. In figure 6.2a, the condensing temperature is shown to be causally dependent on the temperature of the outgoing refrigerant. However, in figure 6.2b the temperature of the outgoing refrigerant is causally dependent on the condensing temperature. According to the causal stability definition by Skorstad (Skorstad, 1992) , the equation of the evaporator:

H = [Tw, Tc] (equation 1 of figure 5b, 6.2a and 6.2b)

is causally stable or unidirectional, which holds in any scenario that might be encountered.

We impose the restriction that a caused or dependent variable must appear only once on the LHS of the equation. That is, each equation expresses all the causal influences on a particular parameter. Thus once the dependent variable is determined, the causal relationships within the equation can be fully specified. In a feedback system, sometimes it is difficult to identify the dependency of variables, then we apply our causal heuristic. Since the causal ordering limits the equations on a self-contained structure, i.e n equations with n unknowns, this restricts the causal ordering theory in finding a unique solution for the model. Thus if the self-contained structure could not maintain during the causality constructing process, the implementation of causal ordering eventually ceased. If there is no minimal complete subset within equations, there is no substitution between variables in the causal ordering theory. The self-contained structure is not only limited multiple possible causalities but also restricted causal behaviour generation in feedback loop systems.

Further, our method is able to be used for model revision. It has the capabilities of allowing the users to make their hypotheses, we then backtrack through the proposed causal graph to construct a new set of

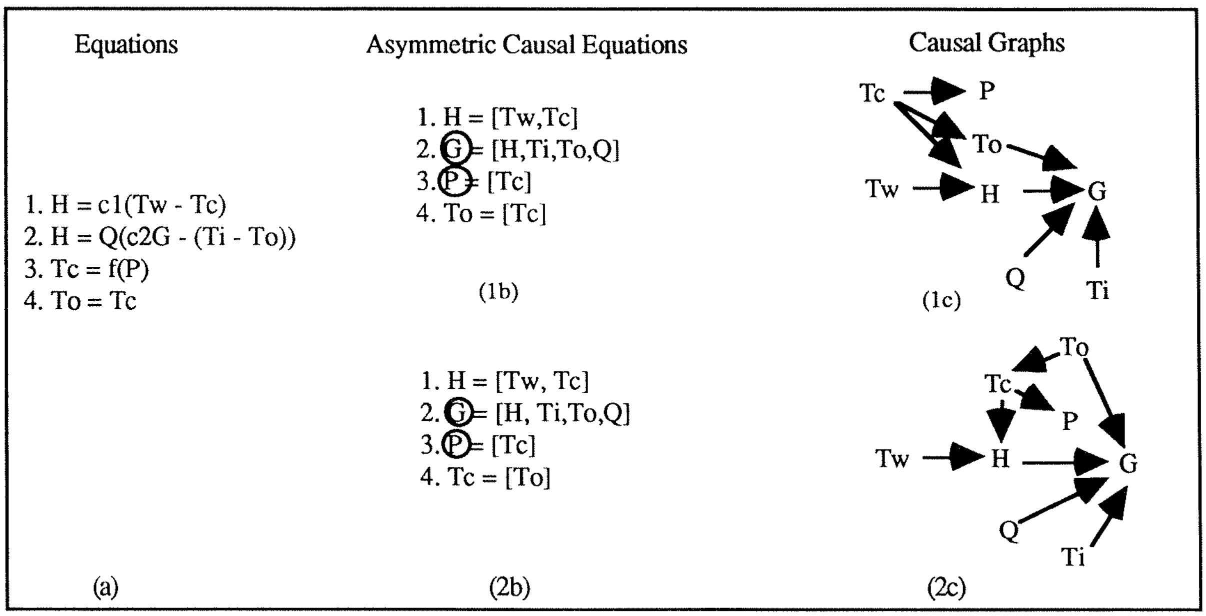

Figure 6. An example of generating multiple causal asymmetric causal equations.

1. G = L - T
2. I = P - K
3. T = I * G
4. K = I
5. L = I
6. P = G

(a)

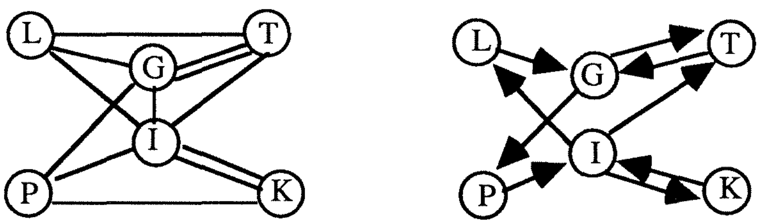

(b) (c)

Figure 7. Equations with it's undirected and a minimal directed spanning

asymmetric causal equations. Thus the user can check the consistency of the hypothesis.

The Algorithm

We present the basic algorithm that formulates a set of equations as asymmetric causal equations in an efficient way. The algorithm delivers all possible causal models if the set of equations together with the causality information are consistent, otherwise it generates appropriate error messages and stops. The interpretation of the causality information is up to the user. Unlike the causal ordering approach, this algorithm can be used in a system with feedback (Lee *et al* , 1992a).

```
procedure genAsyCausalEqu
   % MultiSol is a stack of storing all plausible
   equation models
   % FoundEqu is the equations with great number of
   least occurring variables
   TotalVar := empty;InputDevices :=
   empty;MultiSol := empty;FoundEqu := empty;
   PUSH all the variables occurring in the set of
   equations INTO TotalVar
    while TotalVar is not empty do
        search (least-occurrence-variables)and put
                into InputDevices
        for all InputDevices do
            search (greatest-number-of-least-
            occurring-variables-in-a-equ) and
            put into FoundEqu
            if no more FoundEqu then stop
            for all equations in the FoundEqu do
                set the inputDevices on
                RHS and pop theinputDevices
                if there are more then one
                emaining variable on the
                equation then
                  save all the possible situation
                  on MultiSol
                else put it to LHS and push the
                variable to Inputdevices
                if (this conflicts with causality
                information) then
                  backtrack to next solution
                else "model inconsistent" stop
            end
        end
end
```

Complexity

In order to determine the order of magnitude of the time-complexity of the algorithm, we represent a set of abstract equations (figure 7(a)) together with the set of variables in an undirected spanning tree (figure 7(b)). The vertex represents the variable and the edges are the connection between vertices. In order to detect the edges, each vertex of the undirected graph needs to be visited and the edges incident upon each visited vertex needs to be directed. Visiting each vertex of a graph is equivalent to obtaining its spanning tree. A spanning tree is any tree consisting solely of edges in a graph G and including all vertices of G. Thus the complexity of our algorithm is comparable with the Kruskal algorithm (Horowitz and Sahni, 1976) for obtaining the minimal spanning tree, a spanning tree with minimum cost. Figure 7(c) show the minimal directed spanning tree of the model. Although our algorithm does not look for a lowest-cost edge, we search for the greatest number of least occurring variables for consistency checking. Hence the complexity of our algorithm seems to be $O(n \log n)$ where n is the number of edges. However, if the users provide more causal information, its complexity will be significantly reduced.

DYNAMIC STRUCTURE

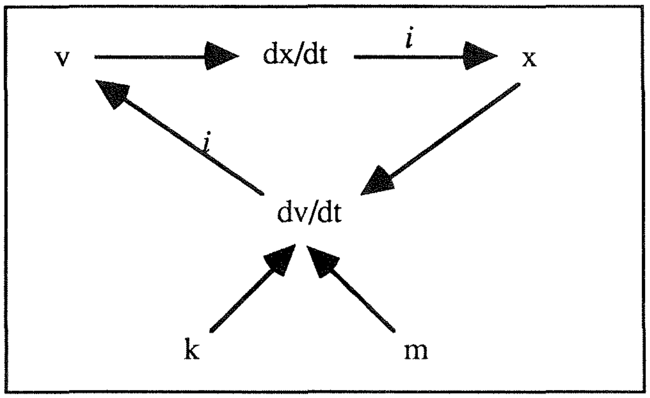

Figure 8. The causal graph of the spring-mass system.

Our approach can be used to represent systems of equations that are of a high order. We reduce higher order system to first order as first order changes with time are reasoned about causality. It seems that normal reasoning about causality is first order. People reason primarily about effects over time.

Consider the following differential equations of a spring-mass system:

$$dx^2/dt^2 + (k/m) * x = 0 \qquad (1)$$

where

x = displacement,
k = Hooke's spring constant,
and m = mass.

Solving for dx^2/dt^2 yields:

$$dx^2/dt^2 = -(k/m) * x$$

Velocity (v) is the first derivative of displacement (i.e. v = dx/dt). We can replace the second order differential, dx^2/dt^2, with a first order differential in terms of v:

$$dv/dt = -(k/m) * x$$

Now there are two first order differential equations:

$$dx/dt = v \qquad (1)$$
$$dv/dt = -(k/m) * x \qquad (2)$$

The causal relations are explicitly represented on the above equations such that the variables on the LHS are dependent on the variables on the RHS. The illustration below represents the causal diagram corresponding to the spring-mass system.

In figure 8, an integration link, which is an edge connecting a derivative of a variable to the variable itself, is marked by *i*. The rate of change of v is determined by x, k and m.

DISCUSSION

Iwasaki and Simon in the retrospective on "Causality in device behaviour" (Iwasaki and Simon, 1993) state that the causal ordering theory was not sufficiently developed in that equations are not able to interpret in all possible causal direction; that the theory does not show how to formulate equation models; that the theory defined causality sometimes are not consistent with the underlying perceived causal direction in dynamic systems.

This paper is aimed at addressing the first point that our proposed theory could interpret all plausible situations that are consistent with the equations, subject to the underlying modelling viewpoint. Our method is not only to identify equations, which are context sensitive depending on the circumstances but also uncover equations, which are causally stable or unidirectional. The causally stable algebraic equations could provide the qualitative model builder with a unique qualitative component which holds in any scenario that might be encountered.

Model formulation is a difficult problem in qualitative physics (Forbus, 1984; Falkenhainer and Forbus, 1991; Weld, 1990). In order to successfully produce causal relations that reflect our intuitive perception in the causal ordering theory, the equations must be structural. Structural equations represent conceptually distinct mechanisms in the system being modelled. However, deciding which equations are structural in a given situation is an essential problem of model formulation. Iwasaki

(Iwasaki 1988; Iwasaki and Simon, 1993) points out that there is no simple way to identify that an equation is structural. However, our proposed method seems to have identified a structural equation and to assign a direct relationship to physical components of the equation by using independent and dependent variables. The method parses the defined independent and dependent variables within equations and reconstructs the equations to be asymmetric causal equations. We use the term "asymmetric causal equation" instead of "structural equation", because a structural equation should express the "real" causality; our equation are structural-like, but express "reasonable causality" (Lee *el al*, 1992a, 1992b). Causality then can be explicitly represented in the asymmetric causal equations. We will explore this issue more and address it later as we haven't discovered from where a mathematical model comes from and from a cognitive point of view, how a model builder constructs a mathematical equations.

In terms of the causal relations in dynamic systems, our proposed method and the causal ordering theory require the differential equations in the model to be in a *canonical form*, where the derivative is on the LHS of the equation and with only one derivative in each equation (Iwasaki, 1988; Lee *et al,* 1992a). The direction of causality is from the variables on the RHS to the derivate on the LHS. However, sometimes in the dynamic physical systems where a change in a quantity is perceived as the cause of some other quantity, such as the Faraday's law of induction, $E = -d\O_B/dt$, where a change in magnetic flux ($d\O_B/dt$) produces electromagnetic force (E), but not vice versa (Iwasaki and Simon, 1993). Also, in order to represent a higher order differential equation, our approach reduces higher order to first order. We do so because most higher order equations are derived from first order equations and causality seems to be explicitly represented in the first order equations. However, in some dynamic systems, such as an equation to describe the bend of a beam:

$$\partial^4 V/\partial X^4 - \partial^4 U/\partial Y^4 = 0,$$

it is difficult to reduce higher order equations to first order equations.

We have successfully tested our proposed method on more than 20 mathematical models, which include all of those in the relevant literature and some greenhouse effect models, i.e. a very large Global Energy Model (Edmonds and Reilly, 1983) with 41 equations and 73 variables. In all cases the method discovers the "correct" causality. However, equations in mathematical models do not model any original causality. We need to further investigate what class of equations (if any) have implicit causality and what classes of equations do not. Our heuristic does manage to recapture the original causality where relevant, or at least a reasonable causality where the equation was not based on an initial causal model. We attempt to understand on what classes of mathematical models this heuristic works or in what way causality is implicit in these models. Also, we need to assign causal effects (+ or - signs) in the causal directions. If, for example, air temperature increases as solar radiation increases, then the causal link between the two is positive (+) or proportional. Conversely, if the level of water in a lake decreases as solar radiation increases, the causal link between the two is negative (-).

CONCLUSION

Decision support systems may require the use of existing complex mathematical models. It is desirable to reduce the equations of such a model to an explanatory causal form to support decision making. We have shown that fixing causal dependencies in a specific context is extremely limiting to behaviour generation. An asymmetric causal explanation approach has been proposed to generate causal knowledge in context. We have shown that it is possible to support the model builder's problem solving by generating all the plausible causal directions in a physical system. Our approach has overcome some limitations of the causal ordering theory in a feedback system. We have presented the algorithm that formulates equations as asymmetric causal equations.

This approach also applies to temporal knowledge in a dynamic system.

REFERENCES

Edmonds, J. and Reilly, J. (1983). A Long-Term Global Energy-Economic Model of Carbon Dioxide Release from Fossil Fuel Use, in Energy Economics, April, 1983.

Falkenhainer, B. and Forbus, K.D. (1991). Compositional modelling: finding the right model for the job, Artificial Intelligence 51:95-143.

Feldman, B., Compton, P., and Smythe, G. (1989). Hypothesis testing: an appropriate task for knowledge-based systems, in Knowledge Acquisition for Knowledge-Based Systems Workshop, Banff, Canada.

Forbus, K. (1984). Qualitative Process Theory, Artificial Intelligence, 24:85-168.

Horowitz, E. and Sahni, S. (1976). Fundamentals of Data Structures, Computer and Software Engineering Series.

Iwasaki, Y. and Simon, H. (1986a). Causality in Device Behaviour, Artificial Intelligence 29: 1-33.

Iwasaki, Y. and Simon, H. (1986b). Theories of Causal Ordering: Reply to de Kleer and Brown, Artificial Intelligence 29: 73-117.

Iwasaki, Y. (1988). Model Based Reasoning of Device Behaviour with Causal Ordering, a PhD Thesis from Carnegie Mellon University.

Iwasaki, Y. and Simon, H. (1993). Retrospective on "Causality in device behaviour", Artificial Intelligence 59: 141-146.

Kowalik, J.S. (1986). Coupling Symbolic and Numerical Computing in Expert Systems, edited by Kowalik, J.S., North-Holland, Elserier Science Publishers, B.V.

Kuipers, B. and Williams, B. (1988). AAAI-88 Qualitative Simulation and Causal Models Tutorial program.

Kunz, J., Stelzner, M. and William, M. (1989). From Classic Expert Systems to Models: Introduction to a Methodology for Building Model-Based Systems, in Topics in Expert System Design, Elsevier Science Publishers.

Lee, M., Compton, P. and Jansen, B. (1992a). Modelling with Context-Dependent Causality, in the Proceedings of JKAW'92, pp357-369.

Lee, M., Compton, P. and Jansen, B. (1992b). Causal Explanation From Mathematical Models, in the Proceedings of AI'92, pp38-43.

May, R. (1976). Simple Mathematical Models with very Complicated Dynamics, in Nature vol 261.

Nayak, P. (1992). Causal Approximations, in the proceedings of AAAI-92, pp 703-709.

Pearl, J. and Verma, T.S. (1991). A Theory of Inferred Causation, in Allen, J.A. et al (Eds) Principles of Knowledge Representation and Reasoning, Morgan Kaufmann.

Simon, H. (1952). On the Definition of the Causal Relation, Journal of Philosophy 49:517-528.

Skorstad, G. (1992). Finding Stable Causal Interpretations of Equations, in Recent advances in Qualitative Physics, Faltings and Struss (Eds), MIT Press, pp399-413.

Top, J. and Akkermans, H. (1991). Computational and Physical Causality, in the proceedings of IJCAI-91, pp1171-1176.

White, B. and Frederiksen, J. (1990). Causal Model Progressions as a Foundation for Intelligent Learning Environments, Artificial Intelligence 42: pp99-157.

Weld, D.S. (1990). Approximation reformulations, in proceedings AAAI-90, pp407-412.