South African Computer Journal
Number 8
November 1992

Suid-Afrikaanse Rekenaartydskrif
Nommer 8
November 1992

Computer Science and Information Systems

Rekenaarwetenskap en Inligtingstelsels

The South African Computer Journal
An official publication of the Computer Society of South Africa and the South African Institute of Computer Scientists

Die Suid-Afrikaanse Rekenaartydskrif
'n Amptelike publikasie van die Rekenaarverenging van Suid-Afrika en die Suid-Afrikaanse Instituut vir Rekenaarwetenskaplikes

Subscriptions

	Annual	Single copy
Southern Africa:	R45,00	R15,00
Elsewhere:	$45,00	$15,00

to be sent to:
Computer Society of South Africa
Box 1714 Halfway House 1685

Computer Science Department, University of Pretoria, Pretoria 0002
Phone: +27 (12) 420-2504 Fax: +27 (12) 43-6454

Sigma Press

Guest Contribution

The paper below was given as an invited address by Prof Roode at the July 1992 Conference of the South African Computer Lectures' Association. (Editor)

The Ideology, Struggle and Liberation of Information Systems

Dewald Roode

Department of Informatics, University of Pretoria

In 1989, Denning *et al* presented the final report of the Task Force on the Core of Computer Science in an article entitled "Computing as a Discipline" [3]. This was said to present a new intellectual framework for the discipline of computing and proposed a new basis for computing curricula.

In the words of the authors, "an image of a technology-based discipline is projected whose fundamentals are in mathematics and engineering." Algorithms are represented as the most basic objects of concern and programming and hardware design as the primary activities. Although there is wide consensus that computer science encompasses far more than programming, the persistent emphasis on programming "arises from the long-standing belief that programming languages are excellent vehicles for gaining access to the rest of the field" [3].

The new framework sets out to present the intellectual substance of the field in a new way, and uses three paradigms to provide a context for the discipline of computing. These paradigms are *theory*, rooted in mathematics; *abstraction*, rooted in the experimental scientific method and *design*, with its roots in engineering.

Programming, the report recommends, should still be a part of the core curriculum and programming languages should be seen and used as vehicles for gaining access to important aspects of computing.

The following short definition is offered of the discipline of computing [3]:

> The discipline of computing is the systematic study of algorithmic processes that describe and transform information: their theory, analysis, design, efficiency, implementation, and application. The fundamental question underlying all of computing is, *"What can be (efficiently) automated?"*

In the same issue of Communications, tucked away towards the end of the journal, an article by Banville and Landry asked the innocent question "Can the Field of MIS be disciplined?" [1]. It is not clear whether the use of the word "discipline" in both articles was purely coincidental – however, the implications were quite clear: computer science was able to talk about "computing as a discipline," and indeed, could present a report which, in a sense, was a culmination of more than twenty years' efforts. Yet, its sister discipline was still asking questions of a very introvertive nature about itself.

It has become quite clear that the fields (leaving aside for the moment the questions of "disciplines") of computer science and information systems (or MIS, informatics, or whatever other name we want to attach to it) have different aims and objectives, different problems that confront it, and, yes, if we want to be truly scientific, different paradigms. To support the latter statement, it is sufficient to contrast the three paradigms of computing with the four paradigms of information systems development described by Hirschheim and Klein [5]. It can be said that a central activity in information systems is the development of information systems, and that therefore, these paradigms have implications for the field of information systems. The four paradigms can be characterized briefly, as follows:

- The analyst as systems expert
- The analyst as facilitator
- The analyst as labour partisan
- The analyst as emancipator or social therapist.

In the same spirit, Lyytinen sees the "systems development process as an instrument in organizational change" [6] and remarks that analysts' principal problems are "in understanding the goals and contents of such change instead of solving technical problems." Already in 1987 Boland [2] observed that: "designing an information system is a moral problem because it puts one party, the designer, in the position of imposing an order on the world of another."

This is clearly a far cry from Denning *et al*'s statement that the fundamental question is "what can be automated?" At the same time, within the context of the field of computing, there is nothing wrong with this question, and it is probably the right question for practitioners of computing to continually ask themselves. But it is a disastrous question for a practitioner of informatics to ask. And it has taken us quite a long time to realise this – that the two disciplines have fundamentally different roles to play. These roles are complementary and supportive, and not destructively opposed.

The liberation of information systems lies in realising this elemental truth: that information systems are man-made objects designed to effect organisational change and that, as such, they can ill be studied using the paradigms of abstraction and engineering mentioned above.

What then is needed? Banville and Landry offer the consolation that we need not concern ourselves too much about the lack of discipline, and that we can indeed even pride ourselves in being a fragmented adhocracy. It is, in fact, even healthy to continue in all sorts of directions. During this process of finding itself, a discipline should be allowed a considerable degree of latitude, and many avenues should be explored. This obviously makes the field of information systems extremely exciting: it is in the process of discovering remarkable truths, discovering that there are in reality people out there using the systems which analysts design and build, and that the most intriguing problems centre around the role of people in all of this: the analyst, the user, their interaction, the impact of sytems on the work lives of workers on all levels, the impact on organizations. These are questions which have mostly been ignored or lightly treated over the years, but which have emerged as *the* problems to be solved. We do not have the tools to solve them – not yet; but a good starting point would certainly be to first understand more about our field and its research tools, for the empirical, positivist approach so often employed will not suffice to solve the above problems.

In the spirit of contributing to the liberation movement of information systems, we have embarked on a study of research on research in Information Systems, and will report on the results more fully in the near future. We define Information Systems as follows [4]:

> Information Systems is an inter-disciplinary field of scholarly inquiry, where information, information systems and the integration thereof with the organisation is studied in order to increase the effectiveness and efficiency of the total system (of technology, people, organisation and society).

In Information Systems then, we see the fundamental question underlying the entire discipline, to be the problem of balancing the need to contribute, through information systems, to the achievement of the mission of the organisation with the moral responsibility to develop and implement socially accepted information systems.

Each of the fields, computer science and information systems, benefits enormously from the activities of the other. Nonetheless, we must recognize the different approaches used by the two disciplines and allow them to complement each other. It should not be our business to convince one another that the universal truth is that which we use in our discipline – whether that be computer science or information systems. Instead, we should seek out the opportunities for synergy, and for complementing each other. If we succeed in doing this at SACLA, then we could indeed do ourselves proud.

References

1. C Banville and M Landry. 'Can the field of MIS be disciplined?'. *Communications of the ACM*, **32**(1):48–60, (1989).
2. R J Boland and R A Hirschheim, eds. *Critical Issues in Information Systems Research.* John Wiley & Sons Ltd., 1987.
3. P J Denning, D E Comer, D Gries, M C Mulder, A Tucker, A J Turner, and P R Young. 'Computing as a discipline'. *Communications of the ACM*, **32**(1):9–23, (1989).
4. N F Du Plooy, L D Introna, and J D Roode. 'Notes on research in information systems'. Unpublished research report, Department of Informatics, University of Pretoria, (1992).
5. R Hirschheim and H K Klein. 'Four paradigms of information systems development'. *Communications of the ACM*, **32**(10):1199–1216, (1989).
6. K Lyytinen. 'New challenges of systems development: A vision of the 90's'. *Data Base*, pp. 1–12, (Fall 1989).

Editor's Notes: To Compete or Collaborate

Human interaction invariably brings with it a blend of competition and collaboration. Competition means that one enjoys the exhilaration of winning while the other endures the shame of loosing. Because of this reward/punishment mechanism, it is a widely assumed that competition enhances performance and efficiency. This dogma pervades not only commerce, sport and politics, but is found in practically all areas of human endeavour, including research.

The competitive spirit in research is found in the well-known saga of Watson and Crick racing to unravel the double helix structure of DNA. Not so well-known, though equally illustrative, is the intensity of Newton's stratagems to oust Leibnitz from receiving any credit for differentiation. Recently there have been reports of scientists who have either tolerated or manufactured fraudulent results in order to win some or other scientific race. The space race, the arms race, the race for an AIDS cure, the scurry for faster smaller hardware, the race for awards, the drive for publications, Nobel prizes: all of this attests to a profoundly competitive international research culture.

But while competition might be the handmaiden of commerce and sport, it is the harlot of research – an unfortunate concomitant of the silly side of human nature. The archetypal researcher not only rises above the incidentals of human accolades; he disdains them. By tradition, the definitive research qualification is a PhD – a Doctor of Philosophy – a lover of thought. Discovery and thought are not only by their very nature rewarding, they are also humbling. When the archetypal researcher moves outside his interior thought-world, it is to share his discoveries. If he is childish, it is not the little boy flexing his biceps and saying: "I'm stronger than you" but the child rushing to tell everyone: "Wow – look at this!" He is forgetful of self: Pythagoras, oblivious of the invading enemy and his impending death while he researches in the sand; Archimedes shouting "Eureka" without care for his nudity. The competitive spirit is a crass intrusion into this ancient legacy of innocence and selflessness.

By its nature, collaboration thrives in a climate of easy social intercourse. It may initially feel uncomfortable for researchers, who are inclined to be socially inept and are wont to bury themselves in work away from society. However, once the plunge to collaborate is taken there is ample evidence that it leads to successful research. In maximizing the use of available talent, it brings about a synergy in which two heads are better than one. All participants enjoy its rewards and no individual has to endure the full weight of its failures. In fact, the notion of collaboration is now so commonplace that significant research seems impossible without it. The tendency, however, is to encourage research collaboration within an organisation, but to emphasize competition in relation to outside organisations.

During a forum discussion at the July South African Computer Lecturers' Association (SACLA) conference, an appeal was made for greater collaboration between universities. Not surprisingly, the information technology disciplines at local universities have always had both a competitive and a collaborative relationship. The competitiveness usually takes the form of friendly rivalry, while the very existence of SACLA bears testimony to a rather unique collaborative relationship. In latter years the competitiveness seems to have intensified, while electronic mail and other developments have improved the prospects for collaboration. At issue, then, is whether there is an imbalance between these dual forces. The appeal at the SACLA forum implied that there is, and I would strongly agree. It is my view (my prejudice, if you will) that competition between universities is a self-indulgent and wasteful dissipation of energy.

Those who are inclined to compete should seriously examine what is to be gained. It is unconvincing to argue that winning makes a significant impact on the way in which students select universities: in the main, this is a matter of geography and language preference. To some extent, the same might be said about staff, although research reputation perhaps plays a more important role here. Neither are research funding agencies (e.g. the FRD) influenced by whether X is "better" in some or other sense than Y. On the contrary, it has wisely been decided to fund on the basis of criteria that are believed to be objective, without any reference whatsoever to the performance of competitors. True enough, funds are limited, but it is precisely for this reason that it is wasteful to divide the little there is between divergent research efforts.

It seems to me that there is a wealth of research talent out there, but that each researcher selects an area of interest almost as a matter of whim. There is an urgent need for well-coordinated collaboration on focussed research areas that have been carefully selected as directly relevant to the country. It is especially incumbent on those who finance, manage and lead research to identify such areas and to encourage collaboration in every possible way.

I look forward to the manifestation of such collaboration in SACJ publications authored by researchers from different university departments. To date there have been none of consequence. If we fail to collaborate, we are in danger of becoming little Don Quixotes who spend our lives attacking windmills and defending castles of xenophobia and irrelevance.

Qualitative Reasoning: An Introduction

John Bradshaw
Department of Computer Science, University of the Witwatersrand, Johannesburg 2050.

Abstract

This paper serves as an introduction to the areas of qualitative reasoning. The three main streams of qualitative reasoning are described: the component centered approach, the process centered approach and the constraint based approach. There is also some discussion on how ideas from these different approaches were adopted in the design of the Doris system. The Doris system evaluates the behaviour of mechanical systems and has been designed for use in a variety of tasks, for example design evaluation and fault diagnosis.
Keywords: *Qualitative reasoning, Process centered, Component centered, Constraint based.*
Computing Review Categories: *I.2.3*

Received July 1992

1 Introduction

The intention of this paper is to provide the reader with an initial feel for the relevant areas in the qualitative reasoning literature. In this paper some the ideas which have influenced the eventual shape of the Doris system [3, 4, 5] are introduced. The Doris system was designed to evaluate the behaviour of mechanical systems using knowledge of purpose and knowledge of structure. It is through an understanding of the relationships which exist between purpose, structure and behaviour that we are able to perform tasks such as fault diagnosis and design evaluation. The Doris system gives some indication that using the same reasoning strategies it is possible for computers to perform the same variety of tasks using the same representation and inference strategies as ourselves.

The general aim of device understanding systems is to enable computers to engage in the same kinds of reasoning tasks as those currently performed by experts within the domain of mechanical causal systems. The approach utilized in device understanding system is different from that employed by first generation expert systems. First generation expert systems do not attempt to represent the underlying nature of the system (the structure and behaviour of the device) or the domain (the basic laws which determine behaviour). Rather, they present the high level problem solving strategy of experts who provide the knowledge contained within the knowledge base. Thus the expert's deeper understanding of the system and domain is not explicitly recorded and is therefore unavailable to the reasoning mechanisms of these first generation systems. It can be thought that these systems operate at a shallow level whereby they naively apply the rules which they hold to the situation which they 'see'. These systems are limited both in terms of what they can do and how they explain the recommendations they make. It is very difficult to adjust this type of expert system to perform any application apart from that for which it was specifically built, since its knowledge base has been constructed to fulfill a particular purpose and all the knowledge has been organized to achieve a clearly defined single aim. Attempts to convert the MYCIN expert system [8] to an intelligent tutoring system illustrates the problems involved [10].

Consequently, the generally accepted approach has been to investigate new theories in order to formulate and organize the type of knowledge which is needed if knowledge bases are to provide a better (more general) basis for understanding the behaviour of the represented system. The qualitative reasoning community has taken the lead in this work. Rather than representing high level information and problem solving strategies in the knowledge base their aim is to capture and represent the fundamental knowledge which describes the behaviour of the system. Qualitative reasoning systems manipulate qualitative variable values, such as hot and cold, in much the same way as humans do, rather than precise quantitative values, eg 65 degrees, which are more traditionally used in computer applications. With this 'deeper', more fundamental knowledge it is anticipated that systems could be constructed which would reason more like human experts than first generation expert systems [28]. These systems have been called second generation expert systems.

Second generation expert systems are knowledge based systems which use deep knowledge to solve problems. Kuipers [26] defines deep knowledge as knowledge about the underlying mechanism of a system which accounts for the visible behaviour of the system. Chandrasekaran and Milne [9] define deep or causal models as models of a domain which may be used to reason about the expected behaviour of a system operating in that domain. Qualitative reasoning and deep models naturally combine together since both aim at providing mechanisms for representing and describing the nature and behaviour of a system rather than simply the precise variable values existing at any one time. The behaviour of causal systems can be explained by examining the behaviour of the parts of the system and the way in which they influence both the behaviour status of the parts and the variable values within the system. The explanation of behaviour can be given at a number of different levels. Price and Lee [28] describe the process of going through these levels as being equivalent to trying to answer a child who asks a succession of 'why's'. Each 'why' forces the person giving the answers to regress back through the causal chain using deeper and deeper knowledge, and so using increasingly fundamental knowledge about the system and the domain.

2 Background

The first Naive Physics manifesto [20] was an early expression of this 'new' thinking about deep knowledge base systems. Although only published in 1979, Hayes had written this paper several years earlier and it had been well circulated and discussed prior to its publication. As such its influence was significant well before 1979. In this paper Hayes articulated the desire for systems which encapsulate the nature of the behaviour of the system rather than, for example, the precise mathematical detail which is contained in descriptions of their behaviour used in traditional physics. Hayes advocated the capture and use of the commonsense knowledge which enables 'the man in the street' to predict the behaviour of physical devices even though he in all probability could not even begin to articulate in detail the real physics involved. In this paper Hayes articulated his commonsense reasoning ideas for the first time. The publication of the second Naive Physics manifesto [21] is a reworking of the 1979 paper in the light of comment and other work which had been done and which showed that there were problems with some of the ideas in the first paper. The second paper is also more oriented to the realization of workable theories rather than the articulation of his grand vision which was described in the first paper.

An integral part of the search for commonsense reasoning was the need to develop techniques which could represent and manipulate qualitative variables rather than the quantitative variables used in traditional physics. People use qualitative values as often as they can and are able to dispense with knowing exact mathematical relationships and precise values by knowing the significant qualitative values for each variable and understanding how the qualitative behaviour of these variables is affected by other values and devices within the system. Qualitative values represent broader ranges of values than traditionally used in physics. Temperature can be represented as hot and cold rather than as a set of precise figures in degrees centigrade. It is usually sufficient to use qualitative variables in place of their precise quantitative counterparts to perform whatever reasoning is necessary. For example, it is usually adequate to reason about heat flow using the qualitative logic — If the radiators are hotter than the room then the room will warm up — without having to know the exact temperatures and the relevant laws of thermodynamics. Qualitative values can be thought of as large, perhaps ill defined, ranges of quantitative values. One of the advantages which may be derived from using qualitative values, aside from the more natural behavioural representation issues, is that the number of values which need to be represented is greatly reduced. This in turn reduces the size of the search space during problem solving which impacts positively on the computational efficacy of the systems using these values. The relationship between qualitative variables is often specified in terms of qualitative differential equations. These are to normal differential equations what qualitative values are to quantitative values.

Since the early Naive Physics manifesto [20] the commonsense reasoning cause has been taken up by a large number of people under the qualitative reasoning banner. Essentially, however, three basic approaches have emerged. These are:

- the constraint centered approach of which Kuipers [27, 25] is the foremost proponent;
- the component centered approach of deKleer and Brown [15]; and
- the process centered approach is identified with the work of Forbus [18].

The three 1984 papers [15, 18, 25] all appear as part of a special issue of Artificial Intelligence [2] and a book [1] edited by D. Bobrow. This early collection of qualitative reasoning work served to establish the early state of the art, as well as to provide the cornerstones on which most subsequent research has been built. The purpose of this paper is to introduce and highlight the main points of the different systems. There are a number of more substantial overview papers e.g. [12] and collections of papers e.g. [29], which introduce and discuss qualitative reasoning in more detail. If more expansive introductory and comparative information is required the reader should refer to these.

3 Constraint Based Reasoning

Perhaps conceptually the simplest of the three approaches is constraint based reasoning. This is based on an articulation of the constraints which determine the behaviour of identified variables [27, 25]. The model, in common with other approaches, is articulated as a set of qualitative differential equations.

The status of every variable is represented as a pair of values: one to establish the value (eg low or high) of the variable and the other its first order derivative (increasing, decreasing or steady). Every variable has a set of values determined by landmarks. A landmark is a significant, named point on the linear spectrum of possible values. As the value of the variable passes through a landmark value the qualitative value of the variable changes. It is possible to establish new landmark variables dynamically as the behaviour of the model is generated. For example, if there is a system which involves decreasing oscillation — such as a mass and spring on a table — then new landmark values would be generated to represent the changing turning points in the oscillation. The full set of values for any variable is therefore formed from the landmark points, the regions between the landmark values and the two regions at the extremes. Of course not all systems will constantly create new landmark values, a central heating system would maintain a constant set of landmark values upon which the behaviour of the different components of the system would be based.

In the Doris system a simpler view of variable values is maintained. Only regions are represented: if points are important these can be represented as small regions. Indeed, even in quantitative systems this approach is used: integers actually represent a set of real values but are usually considered as points for simplicity. Understanding points in this way has perhaps meant that the use of points and regions in qualitative systems was a natural development but the view held here is that the differentiation between these two types of qualitative values adds a disproportionate level of complexity when compared to the benefits which are produced.

The qualitative values and differential equations are evaluated and manipulated by a qualitative simulation system – QSIM. Kuipers has defined qualitative simulation rather differently from the conventional understanding of simulation systems. In qualitative simulation, time is advanced by an amount necessary for a qualitative variable to change value. Thus each step is not of a predetermined time span but rather constitutes an amount of time necessary to allow an interesting qualitative change to occur. Since the detail in the model has already been much reduced by introducing qualitative variables, little is lost by doing away with a quantitative sense of time. What is interesting in qualitative models is the relative behaviour of individual values with respect to each other rather than a precise modeling of their behaviour with respect to any regular notion of time. For example, typically we can say that when we turn on a heater temperature will increase over time until it reaches some new higher, constant value but we can not describe the precise relationship between time and temperature.

The QSIM algorithm selects the next active state from a list of potential qualitatively distinct next behaviour states and from this next state generates a set of candidate child states, new potential next states. Various consistency checks, transition rules and other filters are applied to these candidate child states to eliminate behaviour states which have already been found and those which the system can decide are definitely unattainable. The remaining child states are added to the active list and in turn used to generate their own set of child states. The simulation is complete when there are no active states remaining. However, owing to the dynamic creation of landmark values it may not be possible for QSIM ever to terminate in this fashion. For example, in the decreasing oscillation scenario QSIM would not be able to determine when the system would definitely come to rest and thus would continue generating new landmark values closer and closer to the midpoint of the oscillation as well as generating the 'at rest' behaviour state. This is a trade-off which has had to be set against the advantages which the dynamic creation of landmark values gives. Therefore, with the creation of new landmark values it is possible to model the behaviour of increasing, decreasing and constant oscillation. It was decided, however, that the Doris system should guarantee termination and so the generation of landmark values has not been implemented, resulting in a reduction in the power of the behaviour simulation system.

Kuipers has shown that all possible behaviour states are described by the QSIM algorithm but in addition other "impossible" states and scenarios may also be generated (as we have seen above). This result is basically unsatisfactory since given only the product of a QSIM session all that can be stated with certainty is that some, but perhaps not all, generated states are actually reachable. Kuipers provides no method for distinguishing between reachable and unreachable states and as such the powers which QSIM provides in terms of predicting the behaviour of systems and prejudging their performance is significantly prejudiced. Kuipers [27] has provided some discussion and guidance for recognizing the possible occurrence of these spurious behaviours. DeKleer and Brown [14, p. 1162] have described this approach as qualitative mathematics, not qualitative physics, arguing that "[Kuipers'] model requires a list of constraints bearing no connection to the description of the physical system. He only provides a representation of qualitative differential equations and solution methods for them — it is a qualitative mathematics not a qualitative physics." Indeed Kuipers himself uses this term in [27] to describe the mechanics of the QSIM algorithm.

This inability to distinguish between reachable and unreachable states is a result of the way in which information is represented in the knowledge base. If the view is taken that information held in the knowledge base is simply data to be used by some external process to determine the behaviour of the represented system then the role of the knowledge base is essentially passive. If, however, the knowledge base is viewed as an active, integral part of the evaluation system, then rather than using the information as data for the evaluation system (as in QSIM) the knowledge base can be thought of as the generator of the next behaviour states. The knowledge base itself holds the information and procedural logic required to generate new behaviour states. Thus the model (the knowledge base) has procedural as well as declarative properties. The knowledge base now has two different roles: it serves as a description of the system and also as the tool for generating behaviour states. The approach which is taken in the Doris system is that the evaluation routines outside the knowledge base are used to provide the right environment (the support and infrastructure) in which the evaluation and execution of the procedural elements of the knowledge base take place. For example, consider a heat exchanger. The actual device contains, in its structure, the functional information which determines the way it behaves. It does not require an external evaluation system to determine its behaviour. The knowledge representation of the heat exchanger must therefore contain both the structural information as well as that information which is used to determine its behaviour so that its behaviour can be simulated.

The two main tasks which must be performed by device understanding systems are, behaviour simulation and behaviour evaluation. In the Doris system behaviour simulation is performed using the information contained in the device representation. The role of the Doris system thus is to provide the framework in which simulation can take place. The Doris system provides the system support required during evaluation, and to fulfill the behaviour evaluation tasks. Since simulation and evaluation are now separated, given that there is a common language for describing behaviour, then a number of different behaviour evaluation tasks can be performed using the same representation. Also the same evaluation systems can be applied to a number of different devices which can be modeled using the same representation strategy. This understanding between the different roles of simulation and evaluation has been behind much of the development of the Doris system. [3, 4, 5] discuss this approach more fully.

The other two qualitative reasoning approaches focus much more on the physical aspects of the devices which they represent. DeKleer and Brown's component centered approach is the more straightforward of these to understand.

4 Component Centered Approach

The essence of the component centered approach is the notion that since a device consists of a number of components the behaviour of the device is a function of the behaviour of its components. DeKleer and Brown [15] have taken the work of systems dynamics [11] as their starting point and have described a qualitative version of these theories. "System dynamics is used to model the dynamic behaviour of systems of all types starting with a description of a system in terms of ideal elements connected via ideal connections." [15].

DeKleer and Brown have developed a qualitative physics which they intend should provide the commonsense basis for second generation expert systems. They describe the behaviour of a device by establishing qualitative differential equations (called confluences) which describe (i) how variables change over time, (ii) the influences of other variables and (iii) the behaviour state of different devices. This representation is used to build an envisionment graph. The term envisionment was originally introduced by deKleer [13] to describe the process of determining behaviour from structure. The envisionment graph is the product of the envisionment process. This graph represents the behaviour of the device and contains all potential behaviour states and the paths between them. Thus by examining the graph it is possible to deduce such behaviour types as cycles and describe the behaviour path which a device might follow. Envisioning is not an end in itself but is a means to an end insofar as it provides a description of device behaviour. From this description of behaviour deKleer and Brown [15] state that other application tasks can be performed, such as explanation, fault diagnosis and design evaluation. In [3] this claim is discussed further since the view held here is that the restricted descriptive information held in the deKleer and Brown model combined with the envisionment graph does not provide a sufficiently rich knowledge base on which these other applications can usefully be performed.

A central feature of the component centered approach is the need to be able to describe a library of components. In order to ensure that the parts contained in the library could be used in any number of different systems (possibly unknown at the time the parts are described) the description of the parts should not make any reference to, or any assumption about the environment around them. This is the basis of deKleer and Brown's No Function in Structure (NFIS) principle. They permit class wide assumptions to be made but have tried through this principle to maintain the same flexibility which is held by the physical parts which they describe; the parts can be connected in anyway the designer of the system wishes, irrespective of the original intentions of the designer of the parts. DeKleer and Brown define class wide assumptions as those assumptions which can be applied to a whole class of devices as opposed to those which are "idiosyncratic to a particular device" within the wider class.

The view held by this author is that so long as the knowledge base is to be used for analyzing behaviour, the information allowed under this principle is much too restrictive. In order to understand and comment on behaviour it is important that one includes within the description of the parts knowledge of their context. It is the context in which a device exists which affects its behaviour. It is, therefore, not possible to describe and comment on behaviour without using knowledge of its context. If the aim was to represent the structural knowledge of devices only then the No Function in Structure Principle would hold since to describe the structure of a device no reference need be made to the world outside the device. Amongst others, Keuneke and Allemang [24] have identified problems with this principle and focus on the early goals of deKleer and Brown which were to perform qualitative simulation rather than any 'behaviour evaluation' task. Indeed, in response to Iwasaki and Simon's [22] comment on their early work, deKleer and Brown [16] admit that the NFIS principle was "not as well worked out as it should [have been]" but claim that some equivalent principle is needed to ensure that assemblages of components do not simply behave according to the assumptions built into their description in the model. DeKleer and Brown, in both their 1984 [15] and 1986 [16] papers, state that their model represents just one of many which experts use and that many of these other models do not hold to the NFIS principle. Therefore the fact that the Doris system does not hold to the NFIS principle is not seen as a significant departure from deKleer and Brown's views.

The other important aspect of deKleer and Brown's work, so far as the development of the Doris system is concerned, is their understanding of time. DeKleer and Brown described two forms of time — normal time and mythical time. During normal time variables change value and during mythical time the behaviour of the devices change state. The need for this distinction is so that it is possible to derive and describe the behaviour within a state as well as between states. During normal time variable values are allowed to change state and during mythical time the behaviour of the devices in the system moves from one equilibrium state through a number of non-equilibrium states, to a new equilibrium state. Thus during mythical time the effects of the changes generated in normal time are allowed to propagate through the devices. Changes in variable values will cause device behaviours to change, which will have an effect on connected device behaviour. Thus variable value change will have a rippling effect on device behaviour change. The system imposes no absolute measure on the length of time spent in each of these 'time zones'. In general, however, normal time is usually long compared to mythical time since variable values tend to change more slowly than device behaviour states.

In a high profile discussion on mythical causality and the need for the distinction between normal and mythical time Iwasaki and Simon [22, 23] contend that additional information must be provided to establish the causal ordering. They state that "establishing a causal ordering involves finding sets of variables whose values can be computed independently of the remaining variables and using those values to reduce the structure to a smaller set of equations containing only the remaining variables" [22]. Iwasaki and Simon argue that this can be achieved by analyzing the qualitative differential equations. For systems which contain feedback loops Iwasaki and Simon suggest that further information must be added to the knowledge base. DeKleer and Brown [16] provide several counters to the Iwasaki and Simon arguments, amongst which are that their domain (physical systems as opposed to economics) is different and thus poses different problems and yields different forms of equations. They also argue that their aim was to develop ENVISION (a computer program which produces an envisionment graph) whereas Iwasaki and Simon do not address this issue at all.

Although causality is a much debated issue in the literature, insofar as the work represented by the Doris system is concerned, mythical time is a practical means towards achieving a executable end. The distinction between these two 'time zones' handles the parallel interaction between variables and devices. Mythical time can be thought of as a form of 'time out' from normal time during which the status of devices are allowed to change in accordance with the variable values existing at the end of the previous normal time period. Problems do arise if devices never reach equilibrium during mythical time, and the Doris system is ineffective in these cases. This is one of the limitations of the Doris system and it is left to the responsibility of the knowledge engineer to handle this special case explicitly as and when it occurs, and so to ensure that the represented device will stabilize during mythical time.

In the component centered approach we see that the focus is very much on the physical delineation of the system. Forbus [18, 19, 17], however, has articulated an alternative approach which focuses on the processes which are involved and not on the devices which perform these processes. He originally called this Qualitative Process Theory (QPT).

5 Process Centered Approach

Unlike the two approaches described above Forbus has been more concerned with following Hayes' [20, 21] approach for representing commonsense reasoning about the physical world, rather than developing a qualitative physics. He has broadened Hayes's ideas by addressing both computational and representational issues. The core idea upon which QPT is based on an understanding that solving problems of the physical world "in general requires knowing what kinds of things can happen and how they can affect each other — in other words a theory of processes" [18]. The fulcrum about which QPT turns is the sole mechanism assumption which specifies that "all changes in a system are caused directly or indirectly by process". Given that the closed world assumption can be made and that the model contains the definitions of all the relevant processes, the claim is that the representation holds sufficient knowledge for us to reason about how the system will behave.

Forbus has developed the notion of views and processes which are used to represent the system. Views are frame-like devices used for describing objects, or groups of objects. Views describe the object in a particular position/state and as the object changes state the view which represents it also changes. For example, a spring could be represented by three different views: extended, compressed and normal. The types of knowledge contained within the view structure are:

individuals – the components which combine to create the view;

quantity conditions – internal conditions, which are influenced by individuals within the view, which must be satisfied for the view to exist;

preconditions – external conditions, which are influenced by external factors or individuals, which must hold true;

relations – other internal relationships which must also hold true.

Processes are similar to views in the information which they hold, except that they also contain information on influences. Influences describe the changes which occur in the model. Influences are either positive or negative and specify which quantities are affected by which processes. A process is active when both quantity conditions and preconditions are satisfied. For example, consider a heat exchanger system where heat flows from the hotter to the cooler material. This process could be described as I+(heat(destination), heat(flow rate)) and I-(heat(source), heat(flow rate)). The heat of the source is negatively affected by the flow rate and the heat at the destination is positively affected so the destination warms up when there is a heat flow. If there is only one influence then when the process is active this will determine the behaviour of the system. However, there may be other conflicting (or competing) influences which could prevent this process from occurring. The qualitative reasoning system must resolve any competing influences to determine the next state. If no clear resolution can be made then the alternatives will cause the envisionment to fork. The envisionment process used by Forbus is similar to that of deKleer and Brown, as is the graph which it generates. An envisionment graph consists of nodes which each represent a unique process structure. A process structure represents all the system's processes and views at a given time when certain instances are active as determined by the prevailing conditions at that time.

The problem with this approach, with respect to the desire in the Doris system which is to relate behaviour to structure, is that by focusing on processes and not on components it consciously moves away from a natural device centered view of devices and their structural interrelationship. There is no denying that representing processes does provide a useful (and natural) means for representing behaviour, but the problem lies with relating this behaviour to the structure of the originating system. Forbus has tried to extend the ideas initially articulated by Hayes in his Naive Physics manifesto, namely trying to provide a commonsense representation of the world. As such, the representation of processes does describe the behaviour of the world in a commonsense manner. The problem with this approach is not so much a criticism of the approach per se. Indeed, the contention is that this view does not contradict Cohn's view that "Qualitative Process Theory certainly has the most sophisticated set of modeling features and the most complex ontology" [12]. Instead, the view employed in the Doris system which focuses on components rather than processes, provides a more accessible means of satisfying the aims and objectives of this project which is to develop a means of relating behaviour to structure and so provide a suitable basis for a number of different applications. Forbus' approach is perhaps well suited to explanation but, as regards other tasks such as design evaluation and fault diagnosis, its powers are less obvious.

There is, however, a similar but different view of processes which can be taken. Rather than placing the emphasis on the process itself, shifting attention to the variable attributes which are the manifestation of these processes allows us to describe how the behaviour of separate variables is affected by the behaviour of components within the structure and by other variables. For example, the heat exchanger system can be modeled by describing the temperature of the two materials and the relationship between them. If one is hotter than the other and they are 'in contact' then two statements can be made: first that the lower temperature will increase, and second that the higher will decrease. This is the form in which the laws of thermodynamics are most often couched. Describing the process involved provides a 'deeper' level of knowledge and understanding of the mechanics of the represented system. Although the aim of the project to develop the Doris system was to move away from the use of shallow rules to describe behaviour, the approach taken has been to concentrate on a component (device) centered approach where the benefits are derived from a more fundamental understanding of the behaviour and interaction of devices, rather than on a process centered approach.

6 Qualitative Reasoning and the Doris system

The aims of the project can perhaps be criticized for being too restrictive. However, the position which was been adopted was that there is a need within the area of device understanding for the currently useful first generation expert systems to be enhanced by the more practical and workable ideas within qualitative reasoning. It is known what current expert system technology can do and what properties are desired in future systems. The problem is to advance current systems in the direction of what is desired. This project is a contribution towards this effort. As such there are notions which have been developed by the deep reasoning community, of which the process centered view embodied in QPT is just one example, which for the purposes of this project have not been found to be directly useful. This is not surprising given the different goals and the different starting points of those who have developed these ideas. The qualitative reasoning community, Forbus in particular, is striving to represent the deep underlying physical knowledge of the world and they have taken Naive Physics as their starting point. This project on the other hand looks to provide better representations of real mechanical systems which will provide the flexible type of knowledge base which more sophisticated applications require.

7 Other work

The Bobrow collection of papers [2] provided the qualitative reasoning community with a statement of the state of the art at that time and has served as an important reference point for subsequent work. It is interesting that the three main approaches described then and outlined above have not been substantially extended. Are these the only viable approaches? No convincing alternative has been developed and recent work has tended to either broaden the applicability of one or more of these approaches or to investigate and further refine their properties. There has been some effort to combine the different approaches in order to secure the "best of all worlds", for example the work of Bredeweg [6, 7]. The recent publication by Weld and deKleer [29] provides a more up to date picture of the qualitative reasoning literature. This single volume brings together the most important papers in the field with a short personal historical overview of the development of qualitative reasoning techniques by Johan deKleer.

References

1. D Bobrow. *Qualitative Reasoning about Physical Systems.* MIT Press, Cambridge, Mass, 1985. published book of [2].
2. D E Bobrow. 'special edition'. *Artificial Intelligence,* **24**, (1984).
3. J Bradshaw. *Using Knowledge of Purpose and Knowledge of Structure for Evaluating the Behaviour of Me-*

chanical Systems. PhD thesis, Computer Laboratory, University of Cambridge, New Museum Site, Cambridge CB2 3QG, 1991. Also Tech. report No. 236.

4. J Bradshaw and R M Young. 'Evaluating design using knowledge of purpose and knowledge of structure'. *IEEE Expert*, **6**(2):33–40, (1991).
5. J Bradshaw and R M Young. 'Evaluating the behaviour of the BAe 146 hydraulic system using the Doris system'. In *10th. European Conf. on Artificial Intelligence (ECAI'92)*, pp. 739–743, Vienna, Austria, (1992). John Wiley and Sons Ltd.
6. B Bredeweg, M Reinders, and B Wielenga. 'GARP: A unifying approach to qualitative reasoning'. Technical Report VF-memo-117, University of Amsterdam, Department of Social Science Informatics, (1990). revised and extended version of [7].
7. B Bredeweg and B Wielenga. 'Integrating qualitative reasoning approaches'. In *8th. European Conf. on Artificial Intelligence (ECAI'88)*, pp. 195–201, (1988).
8. B Buchanan and E Shortliffe. *Rule Based Expert Systems*. Addison Wesley, Reading, MA, 1984.
9. B Chandrasekaran and M R (eds). 'Special section on reasoning about structure and behaviour and function'. *ACM SIGART*, **93**:4–55, (1985).
10. W Clancey. *Use of MYCIN's Rules for Tutoring*, pp. 464–489. Addison Wesley, Reading, MA, 1984.
11. I Cochin. *Analysis and Design of Dynamic Systems*. Harper and Row, New York, 1980.
12. A Cohn. 'Approaches to qualitative reasoning'. *Artificial Intelligence Review*, **3**:177–232, (1989).
13. J de Kleer. 'Multiple representations of knowledge in a mechanics problem solver'. In *3rd. International Joint Conf. on Artificial Intelligence (IJCAI'77)*, pp. 299–304, Cambridge, Mass., (1977).
14. J de Kleer and J Brown. 'The origin, form and logic of qualitative physical laws'. In *5th. International Joint Conf. on Artificial Intelligence (IJCAI'83)*, pp. 1158–1169, Karlsruhe, (1983).
15. J de Kleer and J Brown. 'A qualitative physics based on confluences'. *Artificial Intelligence*, **24**:7–84, (1984).
16. J de Kleer and J Brown. 'Theories of causal ordering'. *Artificial Intelligence*, **29**:33–62, (1986).
17. B Falkenhainer and K Forbus. 'Setting up large scale qualitative models'. In *7th. National Conf. on Artificial Intelligence (AAAI'88)*, pp. 301–306, St. Paul, (1988).
18. K Forbus. 'Qualitative process theory'. *Artificial Intelligence*, **24**:85 – 168, (1984).
19. K Forbus, P Nielsen, and F B. 'Qualitative kinematics: A framework'. In *10th. International Joint Conf. on Artificial Intelligence (IJCAI'87)*, pp. 430–436, Milan, (1987).
20. P Hayes. 'The naive physics manifesto'. In D Michie, ed., *Expert Systems in the Microelectronic Age*, pp. 242–270. Edinburgh University Press, (1979).
21. P Hayes. 'The second naive physics manifesto'. In *Formal Theories of the Commonsense World*, pp. 1–36. Ablex Norwood, (1985).
22. Y Iwasaki and H Simon. 'Causality in device behaviour'. *Artificial Intelligence*, **29**:3–32, (1986).
23. Y Iwasaki and H Simon. 'Theories of causal ordering: Reply to deKleer and Brown'. *Artificial Intelligence*, **29**:63–67, (1986).
24. A Keuneke and D Allemang. 'Exploring the 'no-function-in-structure' principle'. *Journal of Experimental and Theoretical Artificial Intelligence*, **1**:79–89, (1989).
25. B Kuipers. 'Common sense reasoning about causality: Deriving behaviour from structure'. *Artificial Intelligence*, **24**:169–203, (1984).
26. B Kuipers. 'The limits of qualitative simulation'. In *9th. International Joint Conf. on Artificial Intelligence (IJCAI'85)*, pp. 128–136, Los Angeles, (1985).
27. B Kuipers. 'Qualitative simulation'. *Artificial Intelligence*, **29**:289–338, (1986).
28. C Price and M Lee. 'Deep knowledge tutorial and bibliography'. Technical Report IKBS3/26/048, Alvey Report, (1988).
29. D Weld and J de Kleer. *Readings in Qualitative Reasoning about Physical Systems*. Morgan Kaufman, 1990.

Notes for Contributors

The prime purpose of the journal is to publish original research papers in the fields of Computer Science and Information Systems, as well as shorter technical research papers. However, non-refereed review and exploratory articles of interest to the journal's readers will be considered for publication under sections marked as Communications or Viewpoints. While English is the preferred language of the journal, papers in Afrikaans will also be accepted. Typed manuscripts for review should be submitted in triplicate to the editor.

Form of Manuscript

Manuscripts for *review* should be prepared according to the following guidelines.

- Use wide margins and $1\frac{1}{2}$ or double spacing.
- The first page should include:
  - title (as brief as possible);
  - author's initials and surname;
  - author's affiliation and address;
  - an abstract of less than 200 words;
  - an appropriate keyword list;
  - a list of relevant Computing Review Categories.
- Tables and figures should be numbered and titled. Figures should be submitted as original line drawings/printouts, and not photocopies.
- References should be listed at the end of the text in alphabetic order of the (first) author's surname, and should be cited in the text in square brackets [1, 2, 3]. References should take the form shown at the end of these notes.

Manuscripts accepted for publication should comply with the above guidelines (except for the spacing requirements), and may be provided in one of the following formats (listed in order of preference):

1. As (a) LaTeX file(s), either on a diskette, or via e-mail/ftp – a LaTeX style file is available from the production editor;
2. As an ASCII file accompanied by a hard-copy showing formatting intentions:
   - Tables and figures should be on separate sheets of paper, clearly numbered on the back and ready for cutting and pasting. Figure titles should appear in the text where the figures are to be placed.
   - Mathematical and other symbols may be either handwritten or typed. Greek letters and unusual symbols should be identified in the margin, if they are not clear in the text.

   Further instructions on how to reduce page charges can be obtained from the production editor.
3. In camera-ready format – a detailed page specification is available from the production editor;
4. In a typed form, suitable for scanning.

Charges

Charges per final page will be levied on papers accepted for publication. They will be scaled to reflect scanning, typesetting, reproduction and other costs. Currently, the minimum rate is R20-00 per final page for LaTeX or camera-ready contributions and the maximum is R100-00 per page for contributions in typed format.

These charges may be waived upon request of the author and at the discretion of the editor.

Proofs

Proofs of accepted papers in categories 2 and 4 above will be sent to the author to ensure that typesetting is correct, and not for addition of new material or major amendments to the text. Corrected proofs should be returned to the production editor within three days.

Note that, in the case of camera-ready submissions, it is the author's responsibility to ensure that such submissions are error-free. However, the editor may recommend minor typesetting changes to be made before publication.

Letters and Communications

Letters to the editor are welcomed. They should be signed, and should be limited to less than about 500 words.

Announcements and communications of interest to the readership will be considered for publication in a separate section of the journal. Communications may also reflect minor research contributions. However, such communications will not be refereed and will not be deemed as fully-fledged publications for state subsidy purposes.

Book reviews

Contributions in this regard will be welcomed. Views and opinions expressed in such reviews should, however, be regarded as those of the reviewer alone.

Advertisement

Placement of advertisements at R1000-00 per full page per issue and R500-00 per half page per issue will be considered. These charges exclude specialized production costs which will be borne by the advertiser. Enquiries should be directed to the editor.

References

1. E Ashcroft and Z Manna. 'The translation of 'goto' programs to 'while' programs'. In *Proceedings of IFIP Congress 71*, pp. 250–255, Amsterdam, (1972). North-Holland.
2. C Bohm and G Jacopini. 'Flow diagrams, turing machines and languages with only two formation rules'. *Communications of the ACM*, **9**:366–371, (1966).
3. S Ginsburg. *Mathematical theory of context free languages*. McGraw Hill, New York, 1966.

South African Computer Journal

Number 8, November 1992
ISSN 1015-7999

Suid-Afrikaanse Rekenaar-tydskrif

Nommer 8, November 1992
ISSN 1015-7999

Contents

Sigma Press