

## THE FOUNDATIONAL ROLE FOR THEORIES OF AGENCY IN UNDERSTANDING INFORMATION SYSTEMS DESIGN

Robert B Johnston & Simon K Milton  
Department of Information Systems  
The University of Melbourne  
[r.johnston@dis.unimelb.edu.au](mailto:r.johnston@dis.unimelb.edu.au)  
[s.milton@dis.unimelb.edu.au](mailto:s.milton@dis.unimelb.edu.au)

### ABSTRACT

In this paper we argue that theories of agency form a foundation upon which we can build a deeper understanding of information systems design. We do so by firstly recognising that information systems are part of purposeful socio-technical systems and that consequently theories of agency may help in understanding them. We then present two alternative theories of agency (deliberative and situational), mainly drawn from the robotics and artificial intelligence disciplines, and in doing so, we note that existing information system design methods and ontological studies of those methods implicitly adhere to the deliberative theory of agency. We also note that while there are advantages in specific circumstances from utilising the situated theory of agency in designing complex systems, because of their differing ontological commitments, such systems would be difficult to analyse and evaluate using ontologies currently used in information systems. We then provide evidence that such situational information systems can indeed exist, by giving a specific example (the Kanban system), which has emerged from manufacturing practice. We conclude that information systems are likely to benefit from creating design approaches supporting the production of situational systems.

**KEYWORDS:** Information Systems Design, Theories of Agency, Ontology, Situated Action

### INTRODUCTION

There has been much recent effort towards finding a theoretical basis for the Information Systems discipline. Some of this work draws on ontology and focuses on the representational role of information systems with particular emphasis on design tools and methodologies. This work is based on the assumption that computerised information systems in part record and track states of a 'real world' system and of real world objects of interest to an information system. In this work, design methodologies and tools have been examined using two specific ontologies. This work is important in that it is leading to a deeper understanding of tools and methodologies. However, it is only part of the story.

In this paper, we argue that a more fundamental understanding of information systems design methodologies will emerge from recognising that information systems are not merely representational, but are embedded in socio-technical systems that are designed to automate the achievement of certain goals. The socio-technical wholes, of which information systems are the representational part, can thus be usefully thought of as intentional or goal-directed agents. We therefore propose that information systems design can be studied at a deep level by analysing their design assumption in terms of theories about how intentional action is possible (theories of agency). This proposal cuts deeper than prior ontological analysis of systems since it examines roles other than representation for information systems. However, it also complements that approach since theories of agency contain ontological commitments at a deep level. In particular we will argue that there are two basic theories of agency - the "deliberative theory" and the "situational theory". Only the deliberative theory has had any explicit influence on approaches to information system design. Its inherent ontological commitments are similar to those of the objectivist/realist ontologies that have been mainly used in the systems evaluation work. By contrast, despite its influence within robotics and computer science, the situational approach has had almost no explicit influence on the design of management or information systems, although elements of the approach can be detected in certain non-computerised operational systems. The ontological commitments of this approach are radically different and present a challenge for the ontological analysis of these information systems.

We begin by arguing our substantive proposal, that theories of agency should be viewed as a deep theory of information systems design. We then describe the deliberative theory of agency and show that it is this conception of agency that informs much systems design practice. We then put forward a radically different conception of agency that has been recently articulated in robotics, artificial intelligence and the cognitive sciences. We use an example from the artificial intelligence literature to illustrate this approach, in particular to emphasise the challenge it poses for traditional theories of representation. We give an example of a system from current supply chain management practice (the Kanban replenishment system), which we believe shows traits of the situational agency theory. Our conclusion is that existing ontological studies assume a deliberative theory of agency for information systems design. Further, we conclude that the ontology-based approach is insufficient for a *comprehensive* theoretical understanding of information systems design and that theories of agency form a theoretical basis to which the ontological studies must relate. Moreover, the analysis points to an opportunity to develop new methods for analysing task domains in such a way as to encourage the production of information systems based on the situational approach to agency, which may display the robustness and efficiency that has been demonstrated for this approach in other disciplines. The main contribution of the paper is the novel proposal that information systems should be studied in terms of their explicit or implicit commitments about the

nature of intentional agency and the working through of the consequences of this proposal for ontological evaluation of systems designs.

### PRIOR WORK

Previous attempts to judge the soundness of design methodologies have focused on identifying a meta-theory that encompasses information systems design principles as a subset, and which thus provides an external reference theory against which the methodologies can be judged and refined. Noting that these methodologies are concerned with the possible entities and states of affairs that can exist in a task domain, which is also the topic of a branch of philosophy known as Ontology (Honderich, 1995), several researchers have made use of ontologies from the philosophical literature for this purpose. Specifically, Bunge's ontology (Bunge, 1977; Bunge, 1979) has been adapted and used to study information systems design and development tools ranging from data flow diagrams through to case tools (Wand, 1989; Wand and Weber, 1993; Wand et al., 1995; Wand and Weber, 1995; Weber, 1997). This ontology is *realist* (assumes the existence of an external reality independent of human experience) and also, in Lakoff's (1987, p159) sense, *objectivist* (assumes that the world consists of entities, properties and relations that can be put in one-to-one correspondence with symbolic set-theoretical counterparts). Chisholm's ontology (Chisholm, 1996), which is a less formalised but more expressive *commonsense realist* ontology, has been used to evaluate specific data modelling tools (Milton et al., 2001; Milton and Kazmierczak, 2002). One difficulty with these approaches is that there are a number of philosophical ontologies available, with differing levels of formality and a wide range of philosophical commitments. The problem of how to select the appropriate one has not been resolved. The choice is thus largely a matter of taste and philosophical bias, which introduces a degree of arbitrariness in this line of research.

Recently, we (Johnston and Milton, 2001) proposed an alternative meta-theory of information systems design. Our point of departure is the observation that information systems are never designed to represent what is in the world *for its own sake*; they are designed to be used by human actors to *achieve certain goals*. Thus, an information system is just one component of a larger socio-technical system intended to allow one or more human actors to bring about certain desirable states of affairs and which is therefore a *goal-directed* system. This point of view is particularly compelling for operational systems where the information system is used to control or direct human action in repetitive, time-constrained and largely routine activity, in pursuit of persistent goals such as: "maximize return on investment", "minimize late delivery", "minimize stock-outs". An appropriate meta-theory for information systems is thus the theory of intentional (goal-directed) behaviour in general, or *theories of agency*. It is important to note that we are viewing the combination of human actors and the information systems they use, as a single intentional agent denoted "the socio-technical system". Many kinds of complex systems ranging from robots, software agents, humans and animals, to organisations can be characterised as intentional agents (Dennett, 1978), and this approach is particularly useful in designing artificial systems, because different theories of intentional agency suggest different decompositions of these complex systems into parts (Brooks, 1986; Agre, 1997).

A theory of agency is a theory about:


- how an agent can obtain data about (sense) the world in which it acts;
- how an agent can represent the world it senses;
- how an agent can manipulate this representation to select actions that will allow it to close in on its goal; and
- how actions are effected.

Recently, a number of authors (Brooks, 1986; Agre and Chapman, 1987; Suchman, 1987; Hendriks-Jansen, 1996; Johnston and Brennan, 1996; Agre, 1997; Clancey, 1997) from diverse disciplines, have argued that there are essentially two radically different main views on the nature of intentional action, which we will call the "deliberative" and the "interactional" theories of agency. What they have in common, though, is that agent actions are presumed to be derived from a representation of the agent's condition in the world of action. To represent the world of action requires commitment to an ontology – a theory about what is, and can be, in the world. The two main theories have radically different ontological commitments and thus differ on the type of representational scheme that is most efficient for intentional action. Our argument is that the fundamental connection between ontology and system design is through the place of ontology in theories of agency and their role in system design. It is because of this connection that ontology is relevant to system design and the evaluation of system design tools. We will argue that existing approaches to computerised information system development are implicitly informed by the deliberative theory. On the other hand, we suspect that many non-computerised operational systems make informal use of the situational approach to agency. It should be possible to design computerised operational information systems using the situational theory. As a long-term goal of our research, we propose to develop and test a methodology and set of tools to assist designers to do this.

### DOING BY THINKING: DELIBERATIVE AGENCY

The deliberative theory of agency posits that an agent acts in a goal-directed way by building and maintaining an abstract model or representation of the world and applying deductive processes to this representation to determine what to do. This stance commits the theory to a certain position about the nature of the world that is both 'realist' and 'objectivist'. The world consists of objects with properties. A particular object is of a given kind because it displays certain properties characteristic of that class of object, and differs from other objects of that class by virtue of its particular values for the properties. Objects stand in certain relations to one another. The world is characterised by a set of possible states of the world, which are allowable configurations of objects (including the agent) under these relations.

The representation of the world is of a particular kind, a "symbol/object" representation (Agre, 1997), in which configurations of symbols, their attributes and relations, are isomorphic to objects, properties and relations within states of the world (Figure 1). The agent uses sense data to define the current state within this world model and then applies deductive processes to find a sequence of symbolic "actions" that will transform the model current state to the model goal state. This sequence of symbolic actions is a "plan" that is handed to an effector system that implements it in the real world.



**Figure 1. The agent and its relation to the world as envisioned in the deliberative theory of agency**

We can identify this deliberative theory of agency with the conventional approach to the design of the socio-technical systems of which information systems are a part. Sensing is achieved by recording data about events in the world (transactions). Databases, or more precisely the data models used to design databases, provide the representational scheme. This representational scheme is conventionally of the symbol/object type: files represent object classes, records in files represent instances of objects and the rules of manipulation of records by programs enforce symbol manipulation rules, which mimic the lawful relation among objects in the world. In the strongest version, programs also deduce goal-attaining actions and human actors are only required to define the goal state, execute the actions (according to computer-generated schedules) and record transactions. A good example is Manufacturing Resource Planning (MRP II) (Wight, 1981). MRP II is a computer-based approach to manufacturing operations control. MRP II embodies the deliberative model in a very literal way (Johnston, 1995; Johnston and Brennan, 1996). A computerised master production schedule represents the company-wide production goal. This is processed against a product structure database that describes how raw materials and bought components can be converted into saleable items. The current availability of these items is checked against an inventory database and, by purely deductive means, schedules for buying and assembling parts to meet the master schedule are produced and handed to operators for execution. Such a strong action selection role for computer systems is now not fashionable. More typically, the information system merely *supports* action selection (decision making) by human actors. Nevertheless, the overall architecture of the combined socio-technical systems (human actors + information systems) is assumed to function according to the deliberative model.

In the past 30 years a number of methodologies have been developed to assist information systems designers to produce and document such systems designs, especially the data models that underpin these systems. They can be referred to collectively as information engineering methodologies (IEM), with Structured Systems Analysis and Design Methodology (SSADM, the British government standard) being a typical example. These design methodologies share the ontological assumptions of the deliberative theory, namely, that systems should represent the environment in which the system acts in terms of external, independent and objective entities, properties and relations. Given this focus on symbol/object representation, use of these methodologies encourages designs for the socio-technical systems, of which the information systems form the representational scheme, which mimic the deliberative approach to agency. Therefore, evaluation of these approaches against

philosophical ontologies such as Bunge's that also share the same objectivist assumption is appropriate and may lead to refinement of these tools and methodologies for the task of building systems along these lines.

### ROUTINE ACTION: SITUATIONAL APPROACHES TO AGENCY

In the past decade or so researchers in diverse disciplines have argued against the efficacy of the deliberative theory of agency (Brooks, 1986; Suchman, 1987; Brooks, 1991a; Johnston, 1995; Hendriks-Jansen, 1996; Johnston and Brennan, 1996; Clancey, 1997; Agre, 1997, p392), on the grounds that even if the world can be represented in such an abstract symbolic way – a debatable ontological assumption (Dreyfus, 1992) – it does not provide a manageable engineering solution to the operation of artificial, autonomous agents, particularly in the face of time-constrained repetitive action. For instance, in a series of proof-of-concept robot designs, Rodney Brooks and co-workers (Brooks, 1986; Brooks, 1991b; Mataric, 1992) have shown that goal-achieving behaviour is possible in robots that do not make use of deliberation or centralised symbolic representations of their environments of action. Their robots use more or less reactive behaviours to establish complex interactions between the agent and features of its environment. They have demonstrated that the performance of such agents in complex environments is more robust than that of deliberative agents and degrades more gracefully in the face of the unexpected. In a similar vein, Agre and co-authors (Agre and Chapman, 1987; Chapman, 1991; Agre and Horswill, 1992) have argued for a theory of routine action, based on the interaction of relatively simple action selection rules with the structure of environments and have demonstrated the possibility of robust goal-directed action in this way using computer-simulated agents.

Emerging from this technical work, is an alternative theory of agency with substantially different representational and ontological commitments from the deliberative theory. The key insight is the essential infeasibility of the continual translation, required by the deliberative theory, from the world of action to the world model and back again to action (Johnston and Brennan, 1996; Agre, 1997), as a principle for time-constrained action in complex environments (this problem is referred to as the “symbol binding problem” in artificial intelligence literature). This translation is required because the world model is an “aerial” or “transcendental” objective view of the world from outside the agent and independent of its goals. Any *real* agent has available only an incomplete, skewed “ground” view from its sensors with which to construct and update its world model. But according to the deliberative theory, it is only by building this aerial world model that its own goal can be expressed and lead to action selection. In addition, the strategy of instantiating the world model independently of the agent's goal means that it is difficult to bound the size of the world model, which typically leads to a computational complexity explosion during action planning (Chapman, 1988).

The solution to these dilemmas proposed in “situational” theories of agency is to provide an agent with more or less reactive responses based on sense data obtainable directly from the agent's *ground view* of the world, and to introduce the agent's intentions and point of view early on in the theory at the level of the ontology and representation schemes. According to this theory, agents respond reactively to “situations” without deliberation. Situations, in contrast to “worlds states”, are accounts of the world from the agent's point of view. They are parsimonious descriptions because they need only include features of the world that are relevant to the agent's purposes (Agre and Chapman, 1987). These features consist of the *relations of things to the agent given its goals*. Actions are selected from a repertoire of more or less reactive rule-like responses to situations. Such a simple approach to action selection will only work (be goal attaining) if the action environment exhibits structure (“affordances”) that makes complex planning unnecessary (Agre and Chapman, 1987), and analysis and exploitation of environmental structure is an important part of designing situated agents (Agre and Horswill, 1992; Hammond et al., 1995; Horswill, 1995; Agre and Horswill, 1997).

We can see the differing roles of representation in the two theories. With the deliberative approach the agent must acquire and make use of a description of the entire universe of discourse from a transcendental point of view. Objects in reality are bound to symbols and states are instantiated. Decisions about action are made taking into account a complete description of reality. Once actions are taken, sensing of the state of reality must continue in order to continue pursuing the goal under consideration. In contrast, because situations are agent-centred and intention-laden, representation of situations with a symbol/object isomorphism is neither possible nor necessary. However, aspects of situations *do* need to be marked in some way in order to fire situation-action responses. Agre and Chapman argue that the appropriate representational scheme should be “indexical” and “functional” in nature (Agre and Chapman, 1987). *Indexical* representations describe the disposition of things relative to the agent and *functional* representations select things according to their relevance for the purposes of the agent.

Agre and Chapman (1987) argue that the ability to rebind indexical/functional representations to situation features that are relevant to the agent's goals on the fly, eliminates the computational complexity explosion inherent in using all-encompassing aerial world models as a basis of action. The reliance of the situational theory on indexical/functional rather than symbol/object representation, indicates it is based on entirely different ontological assumptions and thus systems designed in this way would be difficult to analyse using objectivist ontologies such as Bunge's.

### EXAMPLE OF A SITUATED AGENT: PENGI

In order to make this idea of an agent based on the situational theory more concrete, we now describe the famous “Pengi” simulation, devised by Agre and Chapman to illustrate their ideas about the role of representation and ontology in situated and routine agency.

Pengi is a computer simulation of an agent that plays an arcade game called “Pengo” (which is also simulated). In Pengo, a penguin moves in response to commands from a button and a joystick in a two dimensional video screen “world” of ice-blocks, which effectively form a maze. The world is also inhabited by bees whose behaviour is almost entirely random. If a bee gets too close to the penguin, the penguin dies and the game ends with a loss. Both the penguin and the bees can kick the ice-cubes and they slide in the kicked direction until they reach an obstacle. If the obstacle struck is a bee it dies and if all bees die the game ends with a win. If the obstacle is the penguin it dies, and the game ends with a loss. The environment is thus fairly complex, dynamic, and has a high degree of unpredictability. Pengi’s behaviour, in controlling the actions of the penguin, is clearly goal directed. The game environment is illustrated in Figure 2.




Figure 2. A Pengi game in progress (from (Agre, 1997))

Pengi does not use any planning: it is merely equipped with a set of appropriate responses to each type of situation, which have been put in by hand by the designers on the basis of an analysis of the dynamics of the game. In fact Pengi’s action selection is done using a network of binary logic gates that have no internal state. These generate a series of simulated joystick (left, right, up, down) and button (press) movements that are passed to the game simulator. Thus, Pengi’s executor system is essentially reactive and rule-based. However, the visual (sensing) system is quite sophisticated. It is designed to detect characteristic features of the situation on the game board that are relevant to Pengi’s goals of avoiding defeat and winning. While the vision system has access to the data structure representing the positions of the various objects on the game board, this symbolic representation is not made available to the simulated agent. Instead, this data structure is used like a retinal image and various features in it are extracted by “visual routines” which are used to pick out objects of relevance, such as bees threatening the penguin, ice-cubes that can be kicked, and other spatial relationships. The visual system can also place a small number of markers on objects to keep track of them as a simulation of visual attention. Most of the visual operators ask binary questions about the relationships between markers, the answers to which are used as input to the action selection system.

According to Agre, “Pengi plays a pretty decent game of Pengo. In its present state it is a little better than I am, which is to say it wins from time to time and puts up a good fight.” (Agre, 1997, p265). Pengi can clearly be viewed (and was explicitly designed) as an intentional agent that acts in its simulated world to achieve the persistent goal of survival of the penguin and killing all bees. However, Pengi differs from our characterisation of a deliberative agent in the following ways:

1. Pengi does not make use of an objective representation of its world of action. That is to say it does not build or update a model of the Pengo world in terms of symbols that stand in one-to-one correspondence to the ice-blocks, bees and the penguin and their positions on the board. Nor does it make use of an explicit representation of states of this world in its action selection. Rather, to the extent that Pengi has access to its world of action through symbolic means, the symbols stand for *aspects of situations* of concern to the Pengo player. For instance, a marker might be interpreted as representing “the-bee-that-is-chasing-me” or the “the-ice-block-between-me-and-the-bee”. Unlike “bee-29” and “ice-block-9”, such

references are not objective for two reasons. Firstly, they are *indexical* since they point to objects relative to the penguin's position. They are continually rebound to new individuals as the game progresses. They are also essentially intention-laden, or *functional* as Agre and Chapman put it, since they label features that are relevant to the penguin's goal. Features that are too remote from the penguin to be relevant to its survival are not represented at all and thus not available to the player simulation.

2. Pengi does not make use of the notion of "state of the world" in the sense of a configuration of objects in its world. In fact it does not "know" the state of the game at any time. Instead it bases its action selection upon detection of characteristic features of the world viewed from its myopic agent-centred point of view.
3. Pengi is not deliberative in the usual sense of the word, which implies searching the space for future possible states of the world. Pengi reacts in a rule-based way to aspects of situations it encounters in real-time.
4. Pengi's goal-directed behaviour is an interaction between its reactive rule set and recurring structure in its world of action. The structure of its environment is a co-agent, which cooperates with Pengi in goal attainment/non-attainment. This active participation of environmental structure in intentional action is a hallmark of the situational approach to agency. In this simulation it is easy to see that the overall behaviour is dually determined in this way since one set of programs is determining the penguin's action while another is a simulation of the way the game (its world) behaves in response. Thus in a very direct sense the penguin's moves alone are insufficient to account for its goal achievement. They only work because of the non-arbitrariness of the game world (the rules of the game and the physics of the game world, such as two ice blocks cannot occupy the same square).

These points analyse the Pengi simulation in relation to choices of ontology, representation, action selection and the role of environmental structure for goal attainment. We would expect to find similar traits in any agent system informed by the situational theory. Combining this analysis and our earlier discussion of deliberative agent designs, Table 1 sets out systems traits that might allow one to distinguish the theory of agency which informs the design of any particular intentional system.

Trait	Deliberative	Situational
Type of representation	Symbol/object	Indexical/functional
What represented	Objects (Entities and properties) represented as though external to and independent of the system	Situations and aspects of situations relevant to goal-attainment
Persistence of representation	Persistent account of the state of the world from a dispassionate view point	Transient representations intended only for initiating immediate action
Underlying ontology	Objectivist	Non-objectivist
Action selection	Involving deduction on representations of the state of the world and resulting in plans and schedules for future execution	Immediate (reactive) rule-like responses to aspects of situations
Significance placed upon environmental structure	Structuring of the social and physical environment of work marginalised	Structuring of social and physical environment of work manifest as essential to system functioning

**Table 1. Traits that might allow the underlying theory of agency to be discerned in any given intentional system design**

Pengi also illustrates our main point about ontology. In designing Pengi the designers could have dissected Pengi's world using a realist objectivist ontology that would talk of objects, relations, possible states and state transitions but such an analysis is not pertinent to the design of the agent, since *Pengi does not know its world in these terms*. To Pengi the world consists of situations, detectable aspects of which determine its action. Pengi never needs to make the difficult translation from this ground view of the world to the objective aerial view of individuated objects. Such an objectivist ontology is not part of the theory of agency that Pengi implements.

#### **KANBAN: A SITUATIONAL INVENTORY REPLENISHMENT SYSTEM?**

Can we find examples of socio-technical systems that can be viewed as intentional agents operating on the situational theory of agency? We suspect there are many examples and they should be most often found in environments where the situational theory is known to be most appropriate, namely, routine interaction in time critical environments. Given our suggestion that we have a predilection to designing systems using the deliberative approach, we also suspect they will have evolved rather than been explicitly designed. As a result they may often be unofficial systems and may well be manual information systems. Nevertheless, we would

argue that these systems contain important lessons for formal design of computerised systems that support routine action.

We will present one example, which is a manual operational system that is widely used in the automotive industry for replenishing parts for production, namely the Japanese Kanban system (Schonberger, 1987; Womack et al., 1990). Kanban is the Japanese word for “card” and in this system, stock levels and replenishments are controlled by the movement of cards associated with containers of parts. Each part is stored in a fixed size container and in each container is placed a Kanban card that has printed on it minimal information about the item, usually product-ID, primary supplier and location (see Figure 3). There are a fixed number of Kanbans in existence for a given item and in normal operation they are neither created nor destroyed.



Figure 3. A typical Kanban card

Goods are taken from the container (which is usually stored right at the work-centre) until the container is emptied. The Kanban in that container becomes free and it then becomes a signal that the items need replenishment. These free Kanbans are taken to a Kanban-board near the goods receiving area and are placed on hooks that are sequenced by supplier-ID. The board traditionally has the Kanban system operation rules (Kanban rules) clearly displayed. When a given supplier arrives with shipments of items to deliver, they check the Kanban board and take the Kanbans on their hook back to the supplier site to authorise replenishment of these items next time round. When the items are subsequently supplied the Kanbans are returned in the appropriate containers with the items and they are taken to the work-centre where they are used. The process repeats.

Unlike Pengi, which is a computational intentional agent, if we are to discuss the Kanban system in terms of agency we must recognise that the goal-achieving agent is socio-technical. It consists of the Kanban cards, the Kanban board, the Kanban Rules and the operators (those using the materials and taking the Kanbans to the Kanban-board according to the rules, and those taking the Kanbans as an authorisation to supply and returning them with the goods). Together these elements implement an intentional system that achieves the goal of keeping the parts available for production. The Kanbans are the representational part of this system.

We can make the following observations about this system, which show the connection between its design and the situational theory of agency:

1. Although the Kanbans form the representational scheme in the system, Kanbans are not associated with objects (the parts or the suppliers) in any simple way. When a Kanban is in the container it stands in one-to-one correspondence with the objects in that container. However, since the system is a visual manual system these Kanbans provide no information to the operators. Only free Kanbans have a symbolic significance. When a Kanban is freed by the exhaustion of these items, it represents a shortage of a particular quantity of items and when it is on the Kanban-board, it represents an authorisation to replenish a certain quantity of these items. The physicality of the Kanban as a symbol is crucial to the functioning of the system. Firstly, the constancy of the number of Kanbans, ensured by their physical robustness, effectively controls the item stock level, according to the desired system goal. Secondly, the multiple representational roles of the Kanban are enabled by its changing association with physical places and things.
2. The socio-technical system does not function by acquiring and manipulating an account of the state of the objects in the domain of action. The only information that is apparent to the operators is that conveyed by the free Kanbans. The state of most parts is not available through this visual system, nor is it of any interest to the operators. The free Kanbans are best thought of as marking situations of concern to the operators, namely, shortages of certain parts. This contrasts to conventional computerised systems where the state of stock of any item is available through on-line enquiries made possible by elaborate transaction processing. Just like the markers in Pengi's visual system, the Kanbans are finite in number and constantly rebound during system operation. They are indexical/functional because they mark only aspects of situations relevant to the socio-technical agent with its specific stock replenishment goal, and do not provide an all-purpose account of parts appropriate to a “disinterested” observer such as an accountant.
3. Kanban is a purely reactive and rule-based system. Actions are triggered by situations of shortage marked by free Kanbans. The operators cannot and do not need to plan in the sense of searching through possible future states of the system. Materials Requirements Planning mentioned above is an example of

a conventional transaction-based computerised inventory system that does use extensive state-space search.

4. Kanban will only work in a structured environment. It is well known that Kanban will not tolerate fluctuations in daily item demand greater than about 10% (Schonberger 1987). This environmental constraint is achieved in Japanese practice by smoothing the release of assembly work to the shop floor in such a way as to have a constant mix of variants each day (so-called *Hijunka*) (Monden, 1983). This practice protects the Kanban system from external demand fluctuations, and contrasts markedly with Fordist practice of batching production to achieve economy-of-scale. Because of the manual, visual and socio-technical nature of the system many other simplifications of the physical and social work environment (including cellular machine layouts, teams, quick tooling changeovers, etc.) are required for its smooth running (Monden 1983). The Kanban rules alone are insufficient to ensure the goal state, but interact with these affordances of the environment to achieve the goal as an emergent outcome.

This analysis shows that the Kanban part replenishment system has marked similarities in its functioning to *Pengi*. The characteristics of the Kanban system enumerated above indicate particular commitments in relation to ontology, representation, action selection and the role of environmental structure for goal attainment, which are quite characteristic of situational agency in general. We hypothesise that many other manual card or token systems, particularly those that control routine action, function by implementing the situational theory of agency. Many have evolved from practice and have survived because of the inherent robustness and cognitive efficiency that has been documented for the situational approach in other areas concerned with intentional action. However, because of their origins in praxis they are often marginalized as “practical wisdom” rather than manifestations of a genuine theory of “how to get things done” efficiently with minimum fuss. The Kanban system is perhaps an exception because it is now widely implemented in Western and Japanese car companies (often replacing deliberative approaches) (Muffatto, 1999; Spears and Bowen, 1999) because the phenomenal success of Japanese manufacturing since the 1973 oil crisis can be largely attributed to its use of Just-In-Time, of which Kanban is a part (Womack et al., 1990).

It is interesting to speculate on why the situational approach has had so little impact on information system design. It is probably partly due to the deep influence in our intellectual tradition of the Cartesian duality of body and mind, which encourages system approaches that begin with an objective representation of the environment of action. But robotics researchers operate in the same tradition. In a sense it is easier to detect the limitation of the deliberative approach when designing a totally autonomous system. Bad performance of socio-technical systems can always be repaired (if not totally worked around) by the tacit knowledge of their human components. A second major reason is that information system design is so intimately linked to database design. Existing database designs almost force a symbol/object (even set-theoretical) view of the world on designers. Notice that the Kanban system does not implement database principles at all. Although the Kanban system evolved out of practice, it might equally be reinvented by systematically applying situational agency concepts to the inventory replenishment domain. There is no reason to suppose that computerised systems cannot also be devised, possibly using information technologies other than databases, in this way, that would provide similarly novel designs for control of operations in other domains of routine, time-constrained action, such as retail distribution, emergency dispatch and hospital wards, to name but a few.

## SIGNIFICANCE AND CONCLUSIONS

If we are right in our contention that many routine operational systems, such as the Kanban system, are best understood as implementing the situational theory of agency, it will generally not be useful to analyse these systems using objectivist ontologies such as Bunge’s. We would need an ontology that included at least the non-objective (that is agent-centred and intention-laden) notions of “situations” and “aspects of situations”, such as the ontology that has at least been partly worked out by Agre (1997). Furthermore, it is unlikely that systems design methodologies such as those from the information engineering paradigm would be useful in designing such systems, since they would encourage the designer to decompose the world of action for these systems in a way that is totally inappropriate to the theory of agency they implement. That is to say, most systems designers trained with these tools would come up with system designs similar to familiar transaction-based database systems, which are known to be difficult to make routine in operational environments (Johnston, 1995).

This suggests that by using objectivist ontologies to benchmark systems design methodologies and tools we may be perpetuating, in our information systems practice, a particular style of systems design implicitly informed by the deliberative approach to agency, and that the general conformance of these methods to the ontologies may be a self-fulfilling prophecy. On the other hand, the robustness and efficiency of the situational approach, which has been exploited in other disciplines concerned with automating goal-directed behaviour, could have potential to solve many of the performance and acceptance problems that have dogged attempts to use information systems to automate operational activity. To test this hypothesis, we need to study robust systems that have grown out of practice in operational domains to see if their functioning can be attributed to the situational theory of agency. At the same time we need to develop system design methods which would allow a designer to analyse the design problem in ways that might encourage the discovery of solutions (computerised as well as



manual) exploiting the situational theory of agency. Such an approach to designing a socio-technical system using the situational theory of agency would involve:

1. An analysis and specification of *situations* (relationship between the system and its environment) important to goal attainment. This contrasts with the emphasis of the deliberative approach upon objects and properties.
2. Defining a scheme for representing key *aspects of situations* that trigger action. These schemes will be indexical/functional in nature.
3. Defining a rule-set for actions to be taken in *response* to the occurrence of situations of importance. This emphasis upon reactivity contrasts with the emphasis on planning and scheduling in deliberative theories.
4. Analysing and defining *structuring features* of the system environment that will enable goal-attainment through interaction with such reactive responses.

Our contention is that employing such an approach will produce systems designs that are paradigmatically different from those produced using conventional methods. As with conventional systems, the information system will form the representational part of the socio-technical whole, but now using different representational principles. Although the Kanban system evolved out of practice, it could equally be reinvented by applying the above method to the inventory replenishment domain. There is no reason to suppose that computerised systems cannot also be devised in this way that would provide similarly novel designs for control of operations in other domains of routine time-constrained action, such as retail distribution, emergency dispatch and hospital wards, to name only a few.

#### REFERENCES

- Agre, P. E. (1997) **Computation and Human Experience**, New York, Cambridge University Press
- Agre, P. E. & Chapman, D. (1987) "Pengi: An Implementation of a Theory of Activity", **Proceedings of AAAI-87. The Sixth National Conference on Artificial Intelligence**, Menlo Park, CA., Morgan Kaufman, Los Altos, CA., pp 268-272
- Agre, P. E. & Horswill, I. (1992) "Cultural Support for Improvisation", **Proceedings of AAAI-92. The Tenth National Conference on Artificial Intelligence**, Menlo Park, CA., AAAI Press, pp 363-368
- Agre, P. E. & Horswill, I. (1997) "Lifeworld Analysis", **Journal of Artificial Intelligence Research** 6, pp 111-145
- Brooks, R. A. (1986) "A Robust Layered Control System For A Mobile Robot", **IEEE Journal of Robotics and Automation** 2(1), pp 14-23
- Brooks, R. A. (1991a) "Intelligence Without Representation", **Artificial Intelligence** 47(1-3), pp 139-159
- Brooks, R. A. (1991b) "New Approaches to Robotics", **Science** 253(9), pp 1227-1232
- Bunge, M. (1977) **Treatise on Basic Philosophy: Vol. 3: Ontology I: The Furniture of the World**, Boston, Reidel
- Bunge, M. (1979) **Treatise on Basic Philosophy: Vol. 4: Ontology II: A World of Systems** Boston, Reidel
- Chapman, D. (1988) "Planning for Conjunctive Goals", **Artificial Intelligence** 32(3), pp 333-377
- Chapman, D. (1991) **Vision, Instruction, and Action**, Cambridge, MA., The MIT Press
- Chisholm, R. M. (1996) **A Realistic Theory of Categories-An Essay on Ontology**, Cambridge, Cambridge University Press
- Clancey, W. J. (1997) **Situated Cognition: On Human Knowledge and Computer Representations**, Cambridge, Cambridge University Press
- Dennett, D. C. (1978) **Brainstorms: Philosophical Essays on Mind and Psychology**, Brighton, Harvester
- Dreyfus, H. L. (1992) **What Computers Still Can't Do**, Cambridge, Massachusetts, The MIT Press
- Hammond, K. J., Converse, T. M. & Grass, J. W. (1995) "The Stabilization of Environments", **Artificial Intelligence** 72(1), pp 305-328
- Hendriks-Jansen, H. (1996) **Catching Ourselves in the Act: Situated Activity, Interactive Emergence, Evolution, and Human Thought**, Cambridge, MA, The MIT Press
- Honderich, T. (ed.) (1995) **The Oxford Companion to Philosophy**, Oxford, Oxford University Press
- Horswill, I. (1995) "Analysis of Adaption and Environment", **Artificial Intelligence** 73(1-2), pp 1-30.
- Johnston, R. B. (1995) "Making Manufacturing Practices Tacit: A Case Study of Computer Aided Production Management and Lean Production", **Journal of the Operational Research Society** 46(10), pp 1174 - 1183
- Johnston, R. B., & Milton, S. K. (2001) "The Significance of Intentionality for the Ontological Evaluation of Information Systems", **Proceeding of the Americas Conference on Information Systems**, Boston
- Johnston, R. B. & Brennan, M. (1996) "Planning or Organising: The Significance of Theories of Activity for the Management of Operations.", **OMEGA, International Journal of Management Science** 24(4), pp 367-384
- Lakoff, G. (1987) **Women, Fire and Dangerous Things: What Categories Reveal about the Mind**, Chicago, The University of Chicago Press
- Mataric, M. J. (1992) "Integration of Representation into Goal-Driven Behavior-Based Robots", **IEEE Journal of Robotics and Automation** 8(3), pp 304-312

- Milton, S. K. & Kazmierczak, E. (2002) "On the Study of Data Modelling Languages using Chisholm's Ontology", **Information Modelling and Knowledge Bases XIII**. Kangassalo, H. et al (eds.). Amsterdam, IOS Press: pp 19-35
- Milton, S. K., Kazmierczak, E. & Keen, C. D. (2001) "Data Modelling Languages: An Ontological Study", **Proceedings of the 9th European Conference on Information Systems**, Bled, Slovenia: pp 304-315
- Monden, Y. (1983) **Toyota Production System**, Norcross, Institute of Industrial Engineers
- Muffatto, M. (1999) "Evolution of Production Paradigms: the Toyota and Volvo cases", **British Journal of Management** 10(1), pp 15-25
- Schonberger, R. J. (1987) "The Kanban System", **Just-In-Time Manufacture**. Voss, C.A. (Ed.) London, IFS Ltd. UK: pp 59-71
- Spears, S. & Bowen, H. K. (1999) "Decoding the DNA of the Toyota Production System", **Harvard Business Review**, pp 1-11
- Suchman, L. A. (1987) **Plans and Situated Actions**, Cambridge, Cambridge University Press
- Wand, Y. (1989) "A Proposal for a Formal Model of Objects", **Object-Oriented Concepts, Applications, and Databases**, Kim, W. & Lochovsky, F. (eds.) Reading, Addison-Wesley, pp 537-559
- Wand, Y., Monarchi, D. E., Parsons, J. & Woo, C. (1995) "Theoretical Foundations for Conceptual Modelling in Information Systems Development", **Decision Support Systems** 15(4), pp 285-304
- Wand, Y. & Weber, R. (1993) "On the Ontological Expressiveness of Information Systems Analysis and Design Grammars", **Journal of Information Systems**, pp 217-237
- Wand, Y. & Weber, R. (1995) "On the Deep Structure of Information Systems", **Information Systems Journal** 5(3), pp 203-223
- Weber, R. (1997) **Ontological Foundations of Information System**, Blackburn Australia, Buscombe Vicprint
- Wight, O. W. (1981) **Manufacturing Resource Planning: MRP II**, Essex Junction, VT., Oliver Wight Publications, Inc
- Womack, J. P., Jones, D. T. & Roos, D. (1990). **The Machine That Changed The World**, New York, Rawson Associates