BDI Agents: A Soft Model for Organisations?

Adam Jenkins
Bevan Jarvis

University of South Australia
e-mail: adam.jenkins@unisa.edu.au
e-mail: bevan.jarvis@postgrads.unisa.edu.au

Abstract

Soft Systems methodologies provide a means by which systems researchers can understand and work with complex systems. However they rarely attempt to model the cognitive processes that often occur within their models. Artificial Intelligence researchers, on the other hand, do try to develop models of cognitive processes. This paper examines the possibility of combining the two approaches in such a manner as to maintain the strengths of each. To do so it combines Checkland and Holwell's POM (Processes for Organisational Meanings) model with the BDI (Beliefs, Desires and Intentions) architecture from AI. It offers three possibilities for such a mapping, and then briefly explores the possibilities inherent in doing so. Finally, it offers some directions for future research.

Keywords

Soft Systems, POM, Checkland, BDI, Artificial Intelligence

INTRODUCTION

The Soft Systems Methodology (SSM) is a methodology that has evolved over time towards the goal of understanding real situations holistically, employing the concepts and tools related to "systems thinking". Within SSM, organisations are talked about as if they were person-like entities, able to decide, plan, act and learn. One particular tool within SSM, the POM (Processes for Organisation Meanings) model, (presented in Checkland and Holwell, 1998), is used to describe and talk about the processes that go on within organisations. POM has been applied by a number of researchers in different situations, primarily as a tool to assist in the understanding of organisational processes. These researchers include Cook et al (2000), Crawford and Costello (1999), Alatalo et al (2000) and Checkland and Holwell themselves (1998). As such it has been tested and found to have value within the Information Systems discipline.

However, one of the limitations of SSM in general and POM in particular is that they are often unable to model the processes that occur within the models themselves. This may well be because those processes often mirror the complex decision-making processes which occur within intelligent agents, and which have traditionally been viewed as difficult to model accurately within the "soft" side of systems thinking. From a Soft Systems perspective this limitation is not necessarily a bad thing, but nevertheless the ability to model these processes may still prove to be valuable to Information Systems designers and researchers. Especially in regards to providing POM with some degree of predictive power, so that it can be used more generally in predicting the behaviour of some complex systems, or more specifically in the development of Decision Support Systems (DSS).

Artificial Intelligence, as a field of study, is capable of producing architectures which are complex enough to model the behaviour of intelligent agents, and yet the underlying conceptual frameworks may be simple enough to be readily comprehended and practically applied. BDI (Beliefs, Desires, Intentions) in particular is an AI architecture which describes agents that can act purposefully in a dynamic environment. Philosophically, it is based on Bratman's work on human practical reasoning (Bratman 1987). Thus it can be reasonably claimed that BDI represents a cognitive model of an intelligent agent.

Thus there exists a possible opportunity to combine the work of AI researchers, as instanced in BDI, with the Soft Systems Methodology employing a model such as POM. If such an approach were feasible, this mapping could preserve the main features of the POM model, and make available the best features of BDI. This paper is intended to explore the possibilities inherent in such a combination - how they might be combined, why they might be combined, and what directions this presents for future research.

In the first part of this paper, Soft Systems Methodology and POM, Checkland and Holwell's POM model is examined in detail. Belief Desire Intention, the second part of this paper, analyses the BDI framework in three sections; the first looking at why BDI is seen as a valuable model for the behaviour of intelligent agents; the

second looking at the development and properties of BDI; and the third summarising the essential aspects of the BDI architecture. In the third part, Combining POM and BDI, the paper looks at how BDI can be mapped to parts of the POM model. Finally, some conclusions are drawn about the success of the combinations and proposes some directions for future research.

SOFT SYSTEMS METHODOLOGY AND POM

Soft Systems Methodology (SSM) in its current form, as mentioned in the introduction, is the result of an ongoing (three decades long) programme of action research aimed towards a holistic understanding of real-world situations through systems thinking.

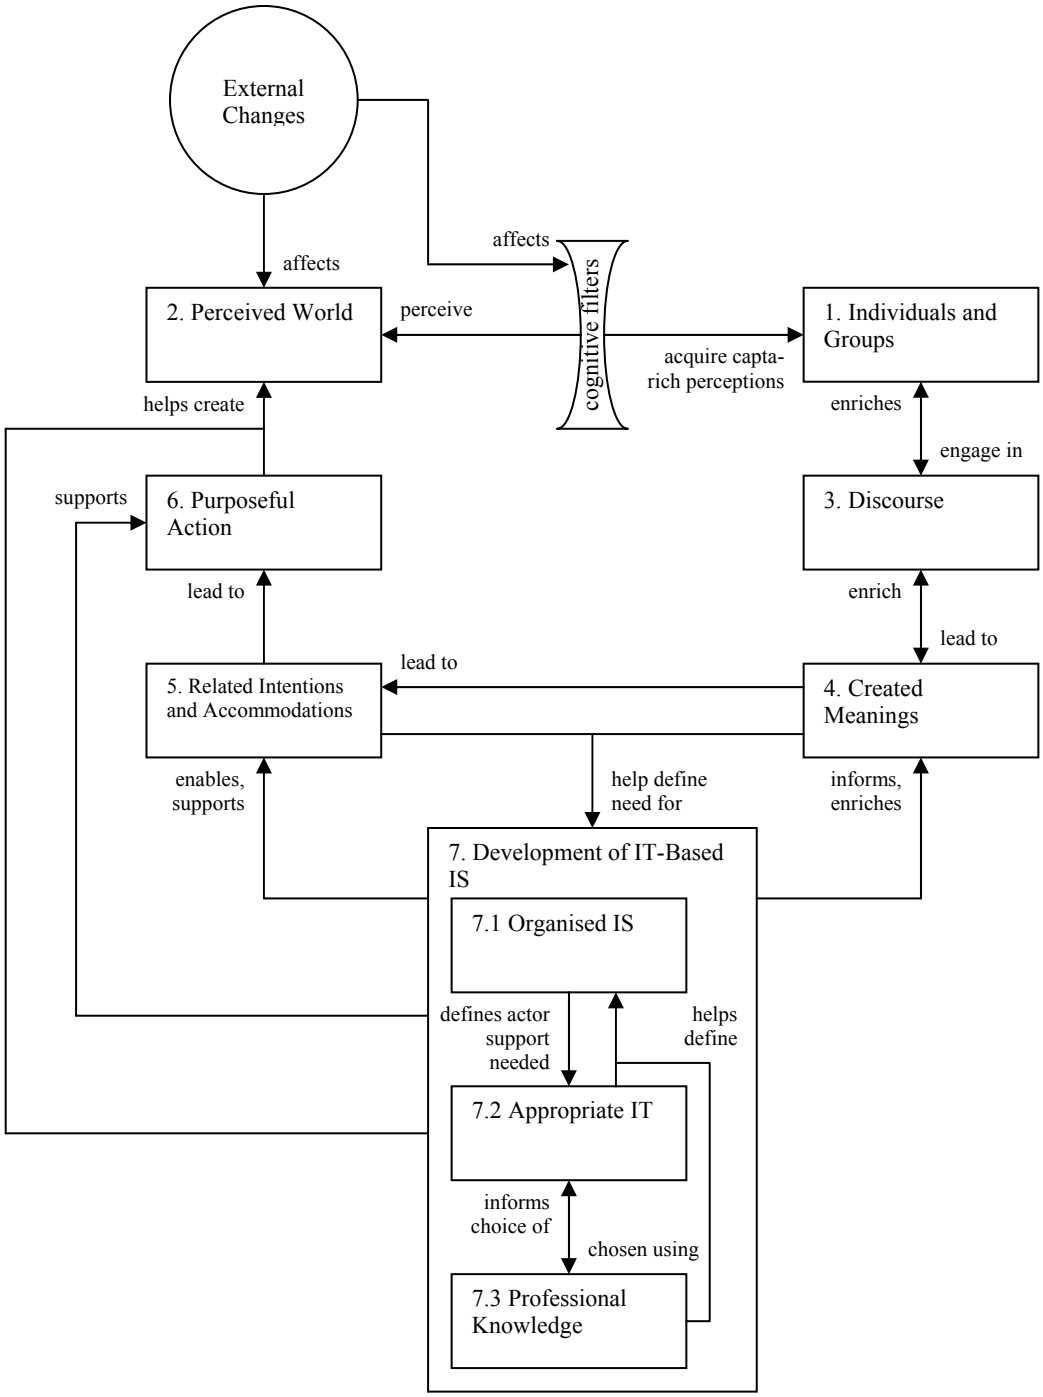

Figure 1: The Processes for Organisation Meanings (POM) model, (Checkland and Holwell 1998, p106).

A fundamental initial observation was that the difficult problems, where the definition of the system itself is problematic, are 'soft': there is in fact no one system that properly describes the situation. The very process of doing systems analysis changes the situation, and thus the 'problem' that one is trying to solve, leading to unpredictable and unrepeatable results. This happens principally because people are involved, who are affected by and *learn from* the process in unique ways. The response to this observation was to emphasise the role of SSM as a *learning process*.

It is perhaps understandable (especially in hindsight) that as SSM evolved, it would come to describe organisations themselves as learning entities. Indeed, it is argued that SSM can present organisations as cognitive entities - an argument that shall be further developed presently.

Filling a perceived gap, Checkland and Holwell (1998) describe models for an organisation, for knowledge, and for the organisational process that fit better with this "soft" approach to thinking about organisations. This gap is not confined to the literature of SSM: Checkland and Holwell (1998) show that the IS literature in general uses these concepts without rigour or indeed explicit definition.

Checkland and Holwell then discuss the processes that occur within organisations. As a device for making sense of, and talking about, life in organisations and their provision of information systems, they offer the Processes for Organisation Meanings (POM) model, reproduced here as Figure 1. The model contains several key components.

Individuals and Groups

The first element (depending on how the cycle is read, as there is no clear beginning or end), are the individuals and groups of individuals that make up the organisation. Each of these can be seen as intelligent "cognitive" agents, with their own beliefs, desires and intentions, which they express through the process of discourse (Element 3).

Perceived World

The perceived world is not the "real world" per se - rather it is the real world viewed through an individual's "cognitive filters". These cognitive filters shape a person's worldview, or *weltanschauung*, which dictates how a person perceives reality. This view, that we do not see the world directly but through a lens created from our own mental states, is supported by writers such as Dennett (1991) and Brunswik (1952). Brunswik in particular offered the "Lens Model", which parallels Checkland's "cognitive filters".

Discourse

Individuals and groups enter into discourse. It is this discourse that will produce the created meanings (Element 5) and the assemblies of related intentions and accommodations (Element 6). There are some significant points to note about the role of discourse in the POM model. The first is that Checkland and Holwell by no means expect this to be an easy process - they point out that it doesn't just involve rational discourse, but also political battles, coercion, and persuasion (1998, p105). Thus although this process is essential, it is by no means easy, and therefore it may need to be carefully managed. Furthermore, it is not assumed that discourse will produce consensus, as will be seen in Element 6. Instead it is expected that the discourse may result in contradictory aims.

Created Meanings

These are, effectively, the information and knowledge created through the process of discourse (Cook et al 2000).

Assemblies of Related Intentions and Meanings

It is assumed that through discourse, even if complete agreement cannot be reached, a number of related intentions will be derived (or recognized), based upon the various and often conflicting interests of the individuals and groups involved in Element 1. Furthermore, where their intentions are not aligned, accommodations may emerge which will permit continuing progress, in spite of the lack of true consensus. Thus this element is not far removed from being a data store - a collection of intentions that are used to decide what actions (in Element 6) are to be taken. However, some processing is required, as these intentions and accommodations still need to be transformed into practical actions. For example, there may be a need to

determine which intentions should be implemented and which should be ignored, or decisions may have to be made regarding the manner in which the intentions are transformed.

Purposeful Action

This is what many might see as the "end result" of the model, but it is no more the end than the individuals are the beginning. Nevertheless, once an assembly of related intentions emerges, these intentions are used to select actions to perform. And according to Checkland and Holwell's view of POM, these actions affect the perceived world (Element 2).

Development of IT-Based Information System

The IS in this model is both created by and assists in the creation of the other elements. The created meanings and assemblies of related intentions (Elements 4 and 6) define the need for the information system, while the information system then supports those intentions and the actions that result from them. The IS also helps create the perceived world, and it enriches the created meanings. The development of an IT-Based IS is in itself a complex system, and may potentially mirror the rest of the POM model. Checkland and Holwell placed three separate elements within thus one:

Organised IS

This refers to an organised strategy for generating information systems, and the collection of IS's which support the organisation.

Appropriate IT

The "Appropriate IT" element consists of the recommendations as to what IT purchases and developments should be made in order to support and develop the information system.

Professional Knowledge

The body of knowledge possessed by those responsible for IT-related decisions that directly relates to their professional tasks.

It is also interesting from an Information Systems perspective to note that, using Checkland and Holwell's model, the IT-Based IS is peripheral to the main work of the processes being modelled, rather than being an integral part. If the IS was removed from the equation discourse could still occur, meanings could still be created, related intentions could still be derived and purposeful action could still be undertaken. The IS enriches and supports these processes, but is not a key component.

BELIEF DESIRE INTENTION

Why BDI?

In order to model the behaviour of humans or groups of humans, it is necessary to have a model for an intelligent agent. Wooldridge and Jennings (1995) provide a valuable definition of what an intelligent agent actually is: agents are reactive, proactive entities that exhibit social ability. That is, agents:

- Perceive their environment and respond to changes in real time in order to achieve their design objectives;
- Act proactively to achieve their design objectives, so exhibiting goal-directed behaviour;
- Interact with other agents.

Note, in particular, that expert systems do not fit these criteria. Expert systems are disembodied from their environment, acting only through a user who is a ‘middleman’; they are not generally capable of reactive, proactive behaviour; and they are not generally equipped with social ability (in the sense of cooperation, coordination and negotiation) (Wooldridge 2002, p27).

As an agent architecture, BDI does fit the criteria for defining intelligent agents. Additionally, BDI:

- Represents a model of cognition, being based on Bratman’s work on human practical reasoning (Bratman 1987);
- Has the rigour of a formal execution model as well as a set of choices and restricting assumptions that define a *practical architecture* (Rao and Georgeff, 1995);
- Has a family of logics that formally describe its behaviour (see, for example, Rao and Georgeff (1995) and more recently Wooldridge (2002));

- As represented by the PRS model is a mature agent architecture (PRS was first implemented in the mid-1980s) that has proved to be the most durable agent architecture developed to date (Wooldridge 2002, p82);
- Is available commercially in at least one implementation (JACK).

Thus BDI can be seen as being, at the very least, a useful model of cognitive, decision-making processes within intelligent agents.

Its potential value within IS then follows directly from this. Information Systems as a discipline, and in particular the "soft" approaches within the field, is intrinsically concerned with the actions of intelligent agents - either as users of systems, parts of systems, or as the systems themselves. If BDI can be shown to be a useful model of these agents, then BDI may prove to be valuable within the field.

The BDI Framework – A Short History and Relevant Principles

The Belief Desire Intention (BDI) model of intelligence is a product of the Rational Agency project at CSLI / Stanford Research Institute in the mid-1980s. The project researched an area of particular interest at that time: building autonomous intelligent agents that can both work towards goals and respond to changes in their environment in real time (Georgeff and Pollack, 1986). In tackling this, the team identified two separate problems:

1. The need to recognise that agents have bounded resources (Simon 1982);
2. The problem of balancing the amount of time an agent spends deliberating (deciding what to do), versus doing means-end analysis (planning how to do it).

Real agents cannot do arbitrarily large computations in limited time. This fact impacts the task of explaining the behaviour of, and indeed designing, agents that are situated in the world and must operate effectively in real time. Consequent on this is that both deliberation and means-end analysis take time. During this time, the situation may change (Georgeff and Pollack, 1986).

Bratman provided the philosophical groundwork for the team's approach, describing intention and showing it to be crucial to understanding (his theory of) human practical reasoning (Bratman 1987). Bratman's work relies on previous work on practical reasoning and ultimately on folk psychology: the concept that our mental models of the world are, in essence, theories. First discarding and then radically extending the Desire-Belief theory, Bratman shows the importance of *intention* as a distinct mental attitude, different in character from both desire and belief. His thesis is that intentions are vital for understanding how limited agents, such as human beings, plan ahead and coordinate their actions.

Crucially to the problem of describing limited resource agents, Bratman showed that the *commitment* implicit in an agent's intention to act limits the amount of practical reasoning that agent has to perform (Georgeff and Pollack, 1986).

In Bratman's model then, agents have:

1. *Beliefs* about their environment, about other agents and about themselves. It should be noted that beliefs can and do change over time.
2. *Desires* which they wish to satisfy.
3. *Intentions*, which are *commitments* to act towards the fulfilment of selected desires. These intentions must be consistent both with the agent's beliefs and with its other intentions. (That is to say, to intend to bring something about, an agent must believe that it is possible to do so, and that doing so will not make other current intentions impossible.)
4. Means-end analysis results in partially specified *plans*, which are (see below) sets of intentions which together serve to fulfil desires in a particular way. They are in fact "fleshed out" intentions. These plans are typically only partially specified at the time of conception, being filled out at the time of their eventual execution (when conditions may have changed).

The results of the Rational Agency project comprise three principal components:

1. The Intelligent Resource-bounded Machine Architecture (IRMA), developed by Bratman, Israel and Pollack (Bratman et al, 1987).
2. The Procedural Reasoning System (PRS), developed by Georgeff and Lansky (1986).
3. The formalisation of a logic of intention, developed by Cohen and Levesque (1990). This puts previous work on speech acts into the context of bounded agents.

After the Rational Agency project was wound up, Georgeff continued to work on BDI. At the Australian Artificial Intelligence Institute (AAII), he and Rao worked on logics that formally describe the behaviour of BDI agents (eg Rao and Georgeff, 1991), and proposed a set of choices and restricting assumptions (Rao and Georgeff, 1995) to obtain a practical architecture. (The latter had already been implemented in PRS and dMARS).

More recently Wooldridge (2000) has extended Rao and Georgeff's work on formal logics, developing the Logic Of Rational Agents (LORA), which comprises four components:

1. First-order logic;
2. A Belief-Desire-Intention modal logic component;
3. A temporal logic component;
4. An action (dynamic logic) component.

Implementations of BDI have in the main followed the PRS model. UM-PRS is the University of Michigan's implementation. Developed at the AAII, dMARS is another "implementation of the PRS architecture" (d'Inverno et al 1998). JAM is a hybrid language that draws from PRS (particularly UM-PRS), Structured Circuit Semantics and Act plan interlingua (Huber, 1999). JACK is distinguished through being designed with a focus on extensibility and with the objective of supporting non-BDI as well as BDI architectures (Busetta et al, 2000).

BDI Summarised

Given the history and strengths of BDI, the key elements in terms of IS can be summarised as follows:

1. BDI has been founded on a well-developed and respected theory of human rational action (Wooldridge 2000).
2. BDI is a mature, practical architecture for modelling intelligent agents that are able to interact with their environments and each other in a meaningful way.
3. BDI requires three essential components - Beliefs, Desires and Intentions. The combination of the three permits the modelling of the decision-making process that results in the undertaking of practical action.
4. A BDI agent can be modelled through the use of a formalised and rigorous family of logics.

COMBINING POM AND BDI

At first glance the possibility of combining POM and BDI may well seem to be difficult at best. Nevertheless, there are reasons for believing that such a move is both possible and valuable. POM does appear to represent a cognitive system in that there are parts of the system that make decisions which lead to purposeful actions, and the system is able to change, react to and learn from the external world (Element 2 on Figure 1). This is exactly the sort of thing that BDI is able to model. However, three questions are of paramount importance:

1. Why model aspects of POM with BDI in the first place?
2. What parts of POM can be modelled by BDI, if any?
3. How can BDI be mapped against these parts?

As will be demonstrated, it appears that three parts of POM are open to modelling with BDI. The "fit" will vary with each of these parts, and the value will also vary, but overall the approach is viable.

Why Combine POM with BDI

As argued by Crawford and Costello (1999), models of human activity in SSM are useful for a number of reasons. They present four:

1. To explore real-world situations;
2. To provide a process which can be followed;
3. To 'structure a debate';
4. To find accommodation between the stakeholders involved.

These four can be summarised as explanatory/descriptive (option 1) and instructive (options 2, 3 and 4). Checkland and Holwell (1998, p107) emphasise the first of these options, with the proviso that the POM model does not purport to describe a real-world situation. They argue that the complexity of real-world situations is such that no model could hope to encompass it all. Thus they view POM as an explanatory model used to make sense of real-world situations, not as a descriptive model (Checkland and Holwell, 1998, p107 and p155).

The POM model, as it stands, is a very broad overview of the processes which an organisation undertakes in order to generate meaning. What it does not claim to do is to model the processes themselves. BDI is ideal for modelling certain cognitive processes which mirror some of those occurring within POM. As such, BDI may offer an opportunity to produce a more fine-grained explanatory model. However, it should be noted at this point that a) such a fine-grained model would be unlikely to assist in the instructional roles described by Crawford and Costello, and b) this does not entail the further claim that the model produced is descriptive.

Furthermore, should BDI be applied to POM in a meaningful way, it provides the additional advantages of being a practical, well-tested model with a rigorous formalised logical basis. As such, not only may it prove to be a model that can be readily applied within IS research, but, when combined with the formalised logical system, it may be able to serve a predictive role, being able to predict possible behaviours within an organisational setting. Such a role is not open to POM on its own, but when combined with BDI such an option may become available. This predictive role may lead to a number of advantages – most notably the creation of Decision Support Systems which would assist managers and others either understand the results of an action or in the formation of the intended action in the first place.

WHAT PARTS OF POM CAN BE MODELLED BY BDI, AND HOW?

Element 1: BDI + Individuals and Groups

Element 1, the individuals and groups, seem to be the most obvious candidate for BDI modelling. Individuals can certainly claim to be intelligent agents, and to possess desires, beliefs and intentions. They react to their environment and each other, and they enter into discourse. As has been demonstrated, BDI can model all of these aspects. The question then becomes not whether BDI can model individuals within the POM model, but a) can it also handle the groups, and b) can doing so be informative.

The answer to whether or not BDI can be used to model the groups within Element 1 is dependent upon how those groups are viewed. There are three possibilities:

1. The groups are homogonous collections of individuals with identical (or at least shared and defined) desires, beliefs and intentions. If this is the case the groups can be treated as intelligent agents in their own right, and modelled accordingly.
2. The groups mirror the structure of the POM model, in which case they can be modelled if and only if the POM model as a whole or significant part can be mapped to BDI.
3. The groups are collections of individuals who do not share desires, beliefs or intentions, and the group as a whole cannot be described in terms of POM.

It should be stated that the question as to whether or there is value in modelling the groups with BDI is a separate concern as to whether or not the mapping is a true and accurate depiction of reality. It may well be the case that modelling the behaviour of the groups with a BDI framework will have explanatory power, even if it cannot be seen as providing a descriptive model.

The question as to whether or not modelling the individuals and groups is informative is one that is somewhat more difficult to answer. This will depend on three factors:

- The difficulty of generating the model;
- The complexity of the model;
- The accuracy of the model in terms of description, prediction and/or explanation.

Further research will be required to determine the outcome of these factors.

Element 7: The Development of an IT-Based IS

At first glance Element 7 (see Figure 1, and the Checkland/Holwell version in Figure 2) does not appear to be an ideal candidate for BDI modelling, based, at least in part, on the lack of anything resembling desires. Indeed the only aspect of the element that mirrors any part of BDI would have to be Element 7.3 (Professional Knowledge), which would be an obvious candidate for beliefs. However the development of an IT-based IS does seem to involve some decision making, as decisions have to be reached regarding what will constitute "Appropriate IT" (Element 7.2). These decisions can be reasonably mapped to Intentions.

The aspect that appears to be missing is that of Desires. There is no part of this element that can be directly mapped to such a concept. However, a close examination of the POM model, as presented by Checkland and Holwell, reveals that the development of an IT-based IS is informed by both "Created Meanings" (Element 4) and "Related Intentions and Accommodations" (Element 5). These two elements "help define the need for" the IS - in other words, they assist in determining the requirements for the IS. These "requirements" can easily be

mapped to desires, as they constitute the reasons for developing or changing the information system. This leads to Figure 2 - a revised version of the development of the information system. The difference is in the acknowledgement that a) the requirements for the IS, when combined with the professional knowledge, lead to decisions being made about what constitutes appropriate IT, and b) that professional knowledge, in an of itself, is not sufficient to generate the purchase or development of IT - there must also be the "intention" to make the purchase or to develop the IT.

To summarise the revised model of the Development of an IT-Based IS (Figure 2):

1. The organisation develops a set of requirements (Desires) for the IS.
2. Those desires are combined with professional knowledge (Beliefs) about IT.
3. The combination of beliefs and desires leads to Intentions being formed regarding what will constitute appropriate choices of IT for the development of the organisation's information systems.
4. These Intentions are then implemented in the purchase or development of systems (Purposeful Actions).
5. The results of the actions performed then inform the professional knowledge base, resulting in a learning loop.

It is clear that this element of the POM model is open to mapping against a BDI architecture, as it has the necessary core features - beliefs, desires, intentions, practical actions, and is both proactive and reactive. Furthermore, from a practical perspective, this element is easier to model than the individuals and groups were, as a) there is only one agent to model, resulting in dramatically reduced complexity, and b) as a result of point a, there is no need to model the interaction between multiple intelligent agents.

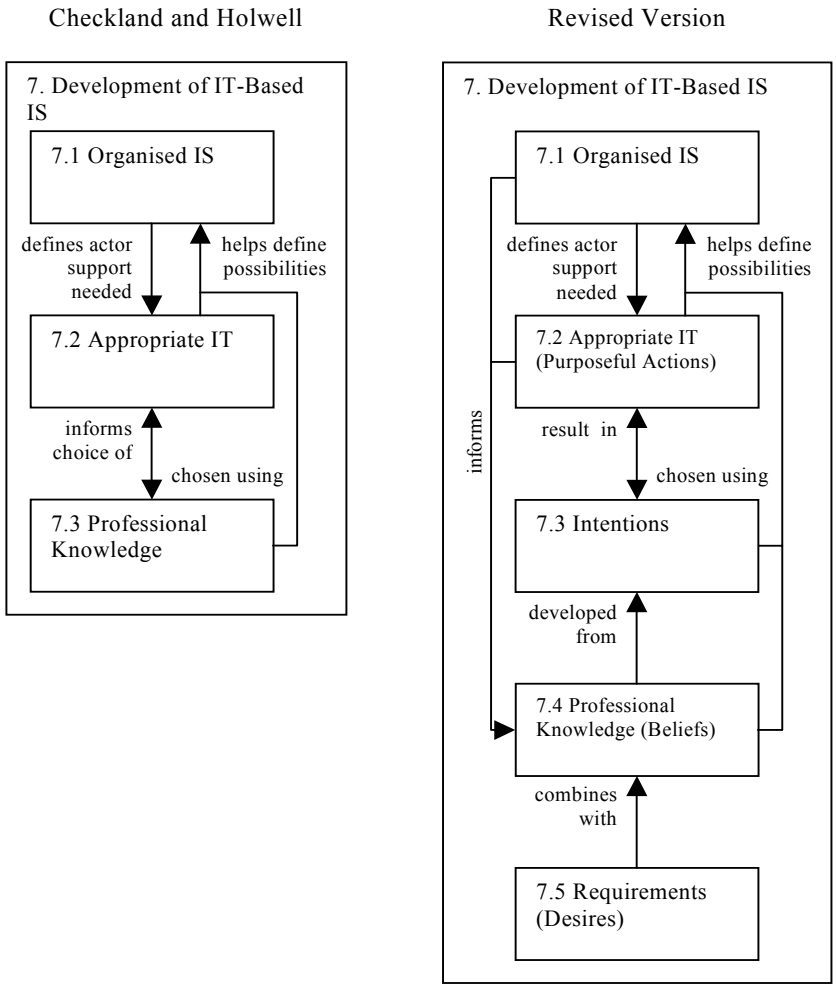

Figure 2: Revised model of the Development of an IT-Based IS

Elements 4-6: The Generation of Purposeful Actions from Created Meanings

There is one more possibility for the application of BDI to POM. Within POM there is a cognitive decision-making process occurring in both the development of related intentions (Element 5 in Figure 1) and in the conversion of some of those intentions into purposeful actions (Element 6). This, of course, depends almost entirely on the interpretation of the relation between Elements 4 and 5. In the Checkland and Holwell model, it is suggested that "Created Meanings" (Element 4) *lead to* "Related Intentions and Accommodations" (Element 5). This suggests that the related intentions are *not* taken directly from "Discourse" (Element 3), which entails that there is, somewhere in the system, a cognitive process that converts one to the other. If this is indeed the case, then clearly there must be room for BDI, for BDI can model this conversion process. However, missing from this equation are desires, which, as has been discussed, are a core component of the BDI architecture. So how can desires be incorporated into the POM model?

Figure 3 shows a possible solution. It suggests that desires emerge through the process of discourse. As individuals and groups enter into discourse, their own, individual, sets of desires merge together. What is produced is Element 9 - a collection of related and unrelated desires, which can be combined with the beliefs contained in Element 4, to bring about a set of related intentions and accommodations.

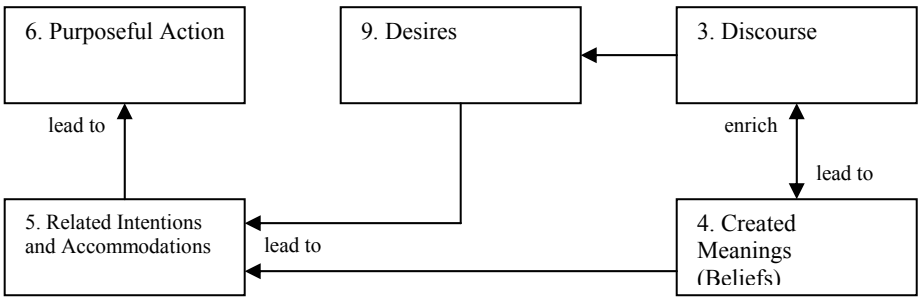

Figure 3: An alternative model of the development of actions from discourse.

The mapping of this model to BDI is reasonably straightforward. "Created Meanings" (Element 4) are mapped to Beliefs. Desires emerge directly from "Discourse" (Element 3), and the combination of Beliefs and Desires brings about Intentions (Element 5). These intentions are then used to guide "Purposeful Action" (Element 6).

CONCLUSIONS AND FURTHER RESEARCH

There is clearly room within the POM for the application of a BDI framework. Exactly what form such a framework would take is yet to be decided, as is the question of which of the possible mappings would be the most valuable. But the mapping is certainly possible, as the key components are there - cognitive systems that use beliefs and desires to develop intentions, and which act both proactively and reactively. Naturally, the claim here is not that organisations *are* intelligent in the same way that individuals are, but that modelling as if they were intelligent is a valuable move that can assist in explaining their behaviour.

Possible future research projects will depend on which of the three mappings is to be examined. In the case of modelling individuals and groups with BDI the aims have already been outlined - it needs to be proven that BDI is a) practical; b) not overly complex; and c) capable of explaining and/or predicting their behaviour.

Of more practical interest is the mapping of BDI against Element 7 - the development of an IT-based IS. One possible line of research would be to study a specific organisation, look at the needs that organisation has for IT/IS development (the desires), combine that with an understanding of the professional IT knowledge possessed by members of that organisation (the beliefs), and employ a BDI-based model to understand how the organisation uses these beliefs and desires to make choices for appropriate IT. What would be of particular interest is to then examine how the choices made reflect on future decisions, again through the use of the BDI architecture. The questions would focus around two issues - first, can BDI explain the choices made by the organisation; and second, can the use of a BDI architecture provide a predictive model. Such a predictive model could then be employed to develop a Decision Support System, where an individual could "plug in" the "desires" for an IT-based IS, and have the model, based on the contained beliefs, output an intention recommending a course of action.

A similar approach is also available for the third option provided, only instead of examining the selection of IT the focus will be on broader organisational decisions. By examining the desires and beliefs that emerge through discourse, it will be possible to place them within a BDI framework, and to study how those beliefs and desires bring about purposeful action. The difficulty here is that discourse is an ongoing process, so it may prove necessary to limit the problem domain by examining a subset of the organisational decisions. Nevertheless, should such a tool be developed, it may prove possible to assist managers in making decisions by allowing them

to input the "desires" and "beliefs" of the component parts of the organisation, and to then see how those elements will respond to changing circumstances.

Naturally, these applications are currently sketchy at best – but any applications beyond the explanatory role offered by Checkland and Holwell are likely to require some sort of modelling of the internal processes, and thus BDI is offered as a potential path to take.

REFERENCES

Alatalo, T, Heikkinen, T. Kallinen, H. and Pihlajamaa, P. (2001) Mobile Information Systems, Seminar work, Theory of Information Systems, Dept. of Information Processing Science.

Bratman, M., Israel D. and Pollack, M. (1987) Plans and resource-bounded practical reasoning, Technical Report CSLI-87-104, CSLI / SRI / Stanford University, 1987.

Bratman, M. (1987) Intention, Plans, and Practical Reasoning, Harvard University Press, Cambridge, MA USA.

Brunswik, E. (1952) The Conceptual Framework of Psychology, International Encyclopedia of Unified Science, (1:10)

Busetta, P., Howden, N., Rönnquist, R. and Hodgson, A. (2000) Structuring BDI Agents in Functional Clusters, Intelligent Agents VI: Proceedings of the 6th International Workshop, ATAL 1999 (LNAI volume 1757). Springer-Verlag: Berlin, Germany.

Checkland, P. and Holwell, S. (1998) Information, Systems and Information Systems, John Wiley and Sons, England.

Cohen, P. and Levesque, H. (1990) Intention is Choice with Commitment, Artificial Intelligence (42, pp213-261).

Cook S., Kasser J. and Burke M. (2000) Assessing the C4ISR Architecture Framework for the Military Enterprise, The 5th International Command and Control Research and Technology Symposium, Australia War Memorial, Canberra ACT, Australia.

Crawford, L. and Costello, K. (1999) Towards A Transferable Methodology For Managing Strategic Change by Projects, IRNOP IV Conference - Paradoxes of Project Collaboration in the Global Economy: Interdependence, Complexity and Ambiguity, Sydney, Australia: University of Technology, Sydney

Dennett, D. (1991) Consciousness Explained, Penguin Books, London, England.

Georgeff, M. and Lansky, A. (1986) "Procedural knowledge", in Proceedings of the IEEE, (74, pp1383-1398).

Georgeff M. and Pollack M. (1986) "Rational Agency (Project Report)", in CSLI Monthly (2:3).

Huber, M. (1999) "JAM: A BDI-theoretic mobile agent architecture", in Proceedings of the Third International Conference on Autonomous Agents (Agents'99).

d'Inverno, M. Kinny, D. Luck, M. and Wooldridge, M. (1998) "A formal specification of dMARS", in Intelligent Agents IV: Proceedings of the International Workshop, ATAL 1997 (LNAI volume 1365). Springer-Verlag: Berlin, Germany.

Rao A and Georgeff M. (1991) "Modeling Rational Agents within a BDI-Architecture", in Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning, Morgan Kaufmann, San Mateo, CA, USA.

Rao A and Georgeff M. (1995) "BDI-agents: from theory to practice", in Proceedings of the First International Conference on Multiagent Systems, San Francisco, USA.

Simon H. (1982) Models of Bounded Rationality. Vol. 2. Cambridge, MIT Press.

Wooldridge, M and Jennings, N. (1995) "Intelligent agents: Theory and practice" The Knowledge Engineering Review, (10:2).

Wooldridge, M. (2000) Reasoning About Rational Agents, MIT Press, Cambridge, MA (USA).

Wooldridge, M. (2002) An Introduction to Multiagent Systems, John Wiley and Sons, Chichester (UK).

COPYRIGHT