Technologies for Intelligent Workflows: Experiences and Lessons

B.R. Odgers, S.G. Thompson, J. W. Shepherdson, Z. Cui, D. Judge, & P.D. O'Brien

Intelligent Business Systems Research Group,
BT Laboratories
MLB1 / PP 12, Martlesham Heath,
Ipswich, Suffolk, IP5 3RE.
ENGLAND
{brian.odgers,simon.2.thompson}@bt.com
http://www.labs.bt.com/projects/ibsr

From: AAAI Technical Report WS-99-02. Compilation copyright © 1999, AAAI (www.aaai.org). All rights reserved.

Abstract

This paper presents research at BT Labs in intelligent workflow management technologies. We summarise the results of past research, and then describe recent work where we applied range of AI technologies to workflow management systems. In particular, we discuss how we are currently using

- Machine learning (specifically support vector machines and time series analysis)
- Qualitative multi-criteria decision-making
- Ontology management and resolution

Finally, we conclude with three challenges for future research into intelligent workflow management systems.

1. Introduction: Perspectives and Motivation

Virtual and other forms of extended enterprise have become increasingly common as organisations seek new ways of delivering value to their shareholders and customers. They integrate both processes and systems to support an end-to-end value chain.

Virtual enterprises require an open decentralised architecture for workflow management where multiple workflow management systems can support business processes which cross organisations. We have identified three classes of problem which impede the development of such workflow management systems.

1. Organisational issues: product and work information models differ from one business to another, as do business rules, organisational structures and the division of responsibilities. This makes it difficult for organisations to; exchange product information, dictate business (sub) processes to suppliers, flag problems and exceptions with each other, and integrate foreign workflows. For this reason, hand-offs between organisations are costly, as they are difficult to establish, hard to automate and time consuming to monitor.
2. Retaining control whilst decentralising behaviour. Centralised monitoring and control of decentralised workflow systems is problematic when business processes cross organisational boundaries and through different jurisdictions. Ensuring there is sufficient visibility to allow problem management, work item tracking, testing for contract compliance etc. is necessary.
3. Load balancing: workflow tools do not enable the flow of work to be managed in a proactive way. If internal or external forces result in a process being overloaded, work is held in ever lengthening work lists, and may eventually be returned unprocessed to the originator. No attempt to re-distributed this work is applied. In inter-organisational workflows this problem is particularly pressing because elements of the workflow are exterior to an organisation's control structures: the flow of work cannot be shut down.

This paper describes the research motivated by these drivers. In section 2 we describe previous research projects we have participated in and summarise the results of this work. In section 3 we describe an integrated range of AI technologies that we are using in our current work. Finally, section 4 lists the research challenges ahead for intelligent workflow research.

2. Previous Research Projects

2. 1 Agent Based Workflow

The ADEPT system (Jennings et al 1996) was one of the first agent based workflow systems. It was developed as a collaborative effort involving BT, ICI, Queen May & Westfield College, and Loughborough University. ADEPT used intelligent agents that collaborated to manage a real business process. ADEPT demonstrated that multiple autonomous agents could manage a business process consisting of nearly one hundred individual tasks.

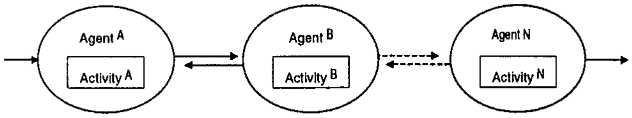

Figure1: Agent Based Workflow

The ADEPT system consisted of multiple software agents which negotiated concurrently with each other in order to reach agreement on how resources are to be assigned to support a business process. The software agents took full responsibility for business process provisioning, execution and compensation, with each agent managing and controlling a given process task or set of tasks.

Agent based workflow systems have been developed by various other research teams (for an example see Harker & Ungar 1996) each offering their own particular enhancements and features.

The overall ADEPT architecture was subsequently developed at BT resulting in the APMS or Agent Based Process Management System, (O'Brien and Wiegand 1998).

The APMS architecture abstracted the most significant elements of the ADEPT system for more generalised use. Agents in the APMS reference model are defined as containing modules for Negotiation, Resource Management and Enactment. Each of these modules can gather and store information for use in process management decisions. A Business Monitoring and Engineering module interfaces with the agents to provide centralised visualisation and monitoring services. The APMS agents manage the tasks which make up a business process. These can be fully automated, semi-automated or manual tasks. Tasks can exchange information during enactment either via the agents or directly with each other.

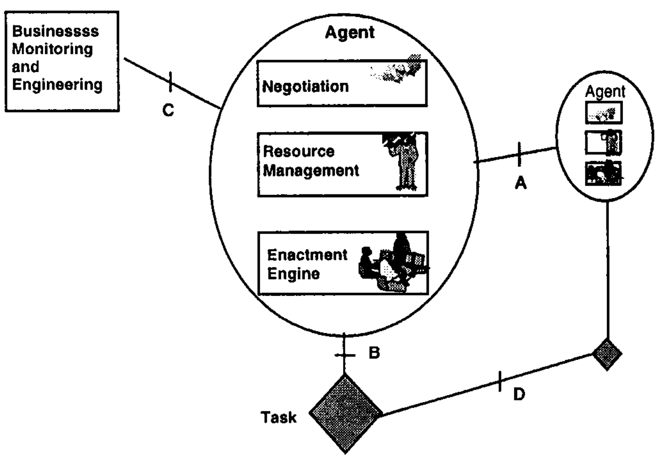

Figure 2: The APMS Reference Model Agent (O'Brien and Wiegand 1998)

Four interfaces, labelled A to D, are defined between the components of the APMS model. Interface A is the agent communication channel between the agents. Interface B is the interface between the agents and the business tasks they manage. Interface C defined the interaction between the agents and the Monitoring and Engineering system. Interface D is the intercommunication system used by the business tasks.

The main outcome of this work was that APMS was a viable technology that could provide significant benefits over existing technologies, in particular, the automation of business process resourcing, its support for decentralised organisational structures, and the automation of compensation or exception handling. However, organisations have already invested in current workflow management technology, adopting APMS would represent a significant cost. Similarly, process tasks were represented as distributed objects, wrapping legacy systems to support such an interface for commercial systems was not addressed.

2.2 BeaT

Project BeaT investigated the engineering issues surrounding the realisation of commercial APMS. It focused on how APMS would access underlying legacy Operational Support Systems (OSS) within an organisation. It identified a number of approaches to interfacing to OSS. Figure 2 shows the different types of interface that were available to agents in a typical enterprise.

The conclusion of this work was twofold. It recognised the need for middleware which provided a layer of abstraction for APMS offering business services. It also highlighted that significant parts of a business process were locked-up within largescale operational support systems and that existing means of interfacing to such systems were inadequate.

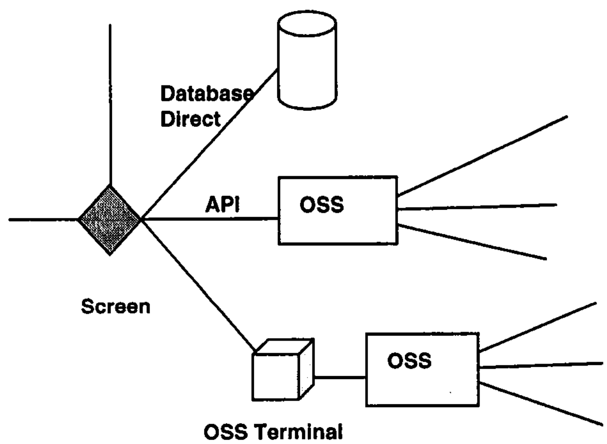

Figure 3: Different Methods of Accessing Information Systems

2.3 Agent Enhanced Workflow

Agent Enhanced Workflow (AEW) (Judge et al 1998) recognised the dependency on legacy workflow management systems as well as the cleaner interface they would provide to underlying process tasks. It investigated the integration of agent based process management with existing commercial workflow management systems. In (Shepherdson et al 1999) the AEW architecture is discussed in full.

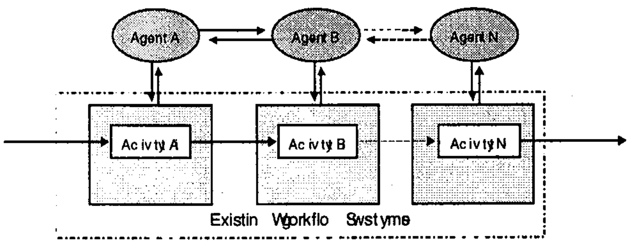

Figure 4: Agent Enhanced Workflow

AEW, in contrast to agent based workflow combines a layer of software agents with a commercial workflow system. The agent layer is given responsibility for the provisioning phase of business process management, whilst the underlying workflow system handles process enactment. In the case of a process failure the agents are employed to renegotiate the flow of work and redistribute the work accordingly, (see Judge et al, 1998, Shepherdson et al 1999). This may be done either reactively are proactively according to the failure type.
In AEW, the workflow engine and associated process model subsume the enactment engine in the APMS architecture. Similarly, the Business Engineering and Monitoring module is replaced with an agency containing agents that can resolve, visualise and verify process models drawn from the process management agents.
This approach provides a number of benefits (including automatic provisioning, interoperability, support for visualisation and verification services), whilst protecting the original investment in workflow technology. At the same time this approach retains the user's confidence in the workflow tools they are familiar with. The use of workflow technology ensures that the business process is explicitly represented and easily maintained.
AEW demonstrated the feasibility of constructing workflow interoperability mechanisms, based on agents as proposed by (Adams, Dworkin 1997). We were able to show that the agent layer could interface workflow systems to other software tools to provide other value-added functionality, such as process monitoring and control.

3. Ongoing work: Integrating AI Technologies in a Workflow Management System

In this section we discuss more recent research building on the results of previous projects. It describes how several different AI technologies have been applied to Business Process Management. The three AI technologies that we have utilized are:

- Machine Learning
- Multi-Criteria Decision Making
- Ontology Management

3. 1 Machine Learning and Time Series Analysis

A crucial element in any intelligent system is that it should have the ability to learn from and adapt to changes in its environment.
Both internal and external forces can influence workflow systems. These forces jeopardise the integrity of the system, causing the build-up of unprocessed items in work lists. To counteract this problem, each agent monitors queue lengths on a per-activity basis. If a queue length surpasses a previously defined threshold, the agent enters a compensation phase, negotiating with others in an attempt to off-load unprocessed work.
Agents can also proactively re-negotiate work allocation plans before a failure has occurred, by analysing trends in activity queue lengths. Thus new allocation plans can be formulated before a failure has even transpired. This analysis can also differentiate between a cyclic flow of work that crosses a failure threshold but will not cause a failure in the long term, and an actual trend that will eventually lead to a process failure. Thus an agent enhanced workflow system can act both reactively and proactively to alleviate the causes of process failure.
Queue-trend learning is a time series analysis task, but we do not know, a-priori what type of concept it is that we are trying to learn, nor the quantity of information that the concept will contain, or what time base the concept will extend over. We have investigated using Support Vector Learning techniques (Vapnik, V.N. 1995) to allow an agent to predict when a failure is likely to happen. Because Support Vector machines are kernel based methods of learning, we are able to select from a library of different kernels each of which is suited to a different type of concept (cyclic, linear etc). Also, because Support Vector machines store the learned classifier as a vector of 'co-efficient/example' pairs, it is possible to time stamp examples and discard individual pairs as they grow older. This enables an agent to avoid the pitfalls of fitting on one particular cycle early on in its management period, namely being unable to adapt to later changes. We have been able to demonstrate that agents can learn complicated demand cycles and predict failure by separating the trend of the demand from the periodic fluctuations of the cycles.

3.2 Qualitative Multi-Criteria Decision Making

As part of the AEW project we developed a qualitative multi-criteria decision-making system for bid ranking by agents.
By assigning belief functions associating bid attribute values with normalised qualitative terms it is possible to aggregate several bid attributes together to provide a more flexible response to bids.
We can illustrate this idea with a short example:

An agent receives two bids for some work from two service providers A and B as shown in .
A simple cost-based decision making process would select bid B, because it is the cheapest. A more sophisticated approach might normalise the values of the two bid attributes Quality and Price, but because the price of Bid B is 10% lower, and the quantity only 2% lower this approach would probably select Bid B as well.
However, our multi-criteria decision making process, based on (Brown and Mamdani 1993), is able to take into account the relative importance of the various attributes, as well as their values.
For example, if we associate 99% Quality with the qualitative value "Very Good", 97% Quality with "Good", $9 price per unit with the qualitative value "Very Good" and $10 price per unit with "Good", and weight the attributes such that Quality is more important than Price, then Bid A will be preferred to bid B.

Quality = 99%
Price = $10 per Unit
Quantity = 1000
Bid A

Quality = 97%
Price = $9 per Unit
Quantity = 1000
Bid B

Figure 5: Two bids for the same work parcel

3.3 Ontology Management and Resolution

One key problem with transferring information between organisations is that the semantics of the information are frequently lost. The problems are less acute when dealing with domain-specific ontologies relating to the same domain, but synonyms, hypernyms, hononyms and other relationships specific to the domain must be approached with care (Sheth 1998)..
Agent-based abstractions offer a solution to this problem by demanding that the members of the agent layer subscribe to a common ontology. The common ontology acts as a shared semantic, which is then mapped to the local semantics of each of the inter-operating systems.
Although the Workflow Management Coalition (WfMC) has defined a Workflow Process Description Language (WPDL) (Workflow Management Coalition 1998), we used PIF, the Process Interchange Format (Lee et al 1997) as the shared ontology that our agents use to communicate process definitions. We made this choice for two reasons, namely that PIF is extensible (and hence can support domain-specific constructs), and that WPDL was not available when we started work on this aspect of the project.
The use of standard interaction protocols (FIPA 1997), as defined by the Foundation for Intelligent Physical Agents (FIPA), between the various workflow agency brokers provides a mechanism for managing inter-organisational hand-offs. Once an agreement has been reached between two workflow agencies, an electronic contract is exchanged which describes the responsibilities of both parties (in terms of volume and type of work, price, timescales etc.). The subsequent performance of each agency is compared with the terms of the various contracts that it is a party to, and the findings are used to support learning and predictive behaviour in itself and/or other agencies.

4. Challenges in Intelligent Workflow Research

This section describes three challenges for the next generation of workflow management systems that have arisen from our research.

1. *Dynamic process creation.* Personalisation of services for individual customer needs is increasingly a key differentiator in today's competitive markets. This has to be reflected in an organisations business processes which need to be tailored for individual requirements. Business process management systems are required which support the dynamic creation and adaptation of business processes which are customised for particular customers.
2. *User centred workflow.* A key ingredient of customer service is a motivated workforce. Workflow management systems are invariably viewed a "big brother" figures dictating work and measuring performance. User centred workflow which facilitates working practices which are tailored for the local workforce and environment would help improve the acceptability and usability of workflow management systems.
3. *Knowledge management tools for business process management.* The ability of an organisation to learn from the performance and effectiveness of its' business processes is vital. Capturing and propagating best-practice through an organisation is key to business process engineering. Knowledge management tools are required which facilitate this exchange of ideas across organisational boundaries.

5. Concluding Remarks

This paper has reviewed a number of successful research projects that have examined the use of new technologies for intelligent workflow management. We have described how we are applying machine learning, multi-criteria decision making and ontology management and resolution technologies to workflow management in our current investigations. Finally we posed three challenges for future work in intelligent workflow management.
The adaptability of an organisations IT infrastructure is key

to its commercial success. Workflow management systems offer a way of automatically managing business processes as well as a means of separating business logic from underlying system. Extending this capability with intelligent technologies which automate activities traditionally performed by middle managers and quicken an organisations response to a changing market can provide significant competitive advantage.

6. References

Adams, T., and Dworkin, S. 1997. Workflow Interoperability Between Businesses. In Lawerence P. ed. 1997. *WfMC Workflow Handbook*, 211-222. New York: John Wiley & Sons.

Brown, E. H., and Mamdani, T. N. 1993. A New Fuzzy Weighted Average Algorithm. In Proceedings of CIMNE. Barcelona, Spain.

FIPA 1997. FIPA '97 Specification Part 2, Agent Communication Language. The Foundation for Physical Agents, Geneva, Switzerland.

Harker, P. T., and Ungar, L. H. 1996. A market-based approach to workflow automation. In Proceedings of NSF Workshop on Workflow and Process Automation in Information Systems: State of the Art and Future Directions. Athens GA.

Jennings, N. R., Faratin, P., Johnson, M. J., O'Brien, P., and Wiegand, M. E. 1996. Using intelligent agents to manage business processes. In Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM96), 345-360. London, UK: The Practical Application Company Ltd.

Judge, D. W., Odgers, B. R., Shepherdson, J. W., and Cui Z. 1998. Agent Enhanced Workflow. *BT Technical Journal,* 16:3, 79-85.

Lee, J., Gruninger, M., Jin, Y., Malone, T., Tate, A., and Yost, G. 1997. PIF: Process Interchange Format v.1.2. PIF Working Group.

Milner, R. 1995. *Communication and Concurrency.* Prentice Hall.

Nwana, H. S. 1996. Software agents: an overview. *The Knowledge Engineering Review*, 11(3), 205-244.

O'Brien, P. D., and Wiegand, M. E. 1998. Agent based process management: applying intelligent agents to workflow. *The Knowledge Engineering Review* 13:2, 161-174.

Shepherdson J.W, Thompson, S. G. and Odgers B.R. 1999. Cross Organisational Workflow Co-ordination by Software Agents. Workshop on Cross-Organisational Workflow Management and Co-ordination, WACC '99 San Francisco, February.

Sheth A. 1998. Changing focus on interoperability in information systems: from system, syntax, structure to semantics. *Interoperating Geographic Information Systems* Goodchild, M., Egenhofer M., Fegeas R., Kottman C. (eds.), Kluwer.

Vapnik, V.N. 1995. *The Nature of Statistical Learning Theory*, Springer Verlag, New York, 1995

Workflow Management Coalition, 1998. Interface 1: Process Definition Interchange Process Model. Document No. WfMC TC-1016-P.