Modelling knowledge worker behaviour in business process studies

Tony Elliman and Julie Eatock
Department of Information Systems and Computing, Brunel University, Uxbridge, UK, and
Nicky Spencer
ROOM Solutions Ltd, London, UK

Abstract

Purpose Aims to describe a successful use of simulated knowledge worker behaviour used in the developing online procedures and software for arbitration – the E-Arbitration-T project

Design/methodology/approach Presents four common factors – deadline, length of task, importance of customer, importance to business – that need to be incorporated within any business process model of knowledge worker behaviour.

Findings A richer model of knowledge worker behaviour is postulated and elements not necessary for the E-Arbitration-T model are identified. The knowledge worker's day was defined as being made up of Scheduled, On-demand and At-will tasks, only some of which may relate to the business process being modelled. A particular question that must be addressed in this extended model is how to model the choices knowledge workers make between competing at-will tasks.

Originality/value The two pieces of work reported here have generated a rich model of knowledge worker behaviour ready for application and refinement in further business process modelling studies.

Keywords Information officers, Simulation, Business process re-engineering, Arbitration, Information systems

Paper type Research paper

Introduction

One of the main catalysts of business process change has been investment in information and communications technology (ICT). Companies are finding themselves under immense pressure to radically improve their performance, either in terms of services provided, or productivity and "information technologies are important enablers of this change" (Grover *et al.*, 1994). Despite this the returns on these investments have often proved disappointing with some research studies showing failure statistics as high as 60 or 70 per cent (Hochstrasser, 1993; Hide, 2000; Taylor, 2000). The reasons that investments made in ICT fail to achieve the expected outcomes or rewards are often related to management of change or expectations rather than shortcomings of the technology installed (Walsham and Waema, 1994; Pinsonneault and Rivard, 1998; Kanellis *et al.*, 1999; Irani and Love, 2000).

This research was performed as part of an EU funded project E-Arbitration-T (IST-2000-25464). The consortium consisted of Brunel University; TIGA Technologies (Paris); Cuatrecasas SRL (Madrid, Spain); Catania University and the Catania Chamber of Arbitration (Italy); and Comercio Electrónico Global and AEDED (Zaragosa, Spain).

The Journal of Enterprise Information Management
Vol. 18 No. 1, 2005
pp. 79-94
© Emerald Group Publishing Limited
1741-0398
DOI 10.1108/17410390510571501

Understanding how changes to the underlying technology will affect the business processes is the key to reducing the disappointment often experienced with investments. It is apparent that one of the major problems contributing to the failure of business process change projects is a lack of some predictive dynamic model for evaluating the effects of designed solutions before implementation. If there is a thorough understanding of the effects of an investment ICT then the benefits can be accurately identified and realised, and the drawbacks can be identified and managed or used to prevent an investment that would otherwise produce disappointment. However, without some form of exploratory evaluation mistakes can only be realised once the redesigned processes are implemented, when it is too late to correct wrong decisions.

Computer-animated simulation models of business processes offer a mechanism for bringing organisational structures alive and arriving at informed recommendations for change (Tumay, 1995; Pinsonneault and Rivard, 1998; Paul *et al.*, 1999; Giaglis, 2001). Although simulation modelling has been applied successfully to manufacturing, process industries and services like hospital clinics, there is much less experience with knowledge-based activities. The purpose of this paper is to examine the development of process models for such knowledge based services. In particular, it looks at the modelling of the key knowledge worker's activity and resource constraints. In the following section the definition of this role is examined together with the implications for modelling their behaviour.

The modelling of human activity systems using discrete event simulation is a non-trivial task and confidence in any results depends on the integrity of the model design. However, it is a well-defined technique for examining constrained forms of human behaviour (Macarthur *et al.*, 1994; Paradi *et al.*, 1995). Variability is introduced in to discrete event models using random number – Monte Carlo methods. However, the model creates variation in their behaviour by setting random start times for cases and randomising, within limits, the time taken to perform various tasks.

The next section of the paper presents the use of business process models in a European Union project – E-Arbitration-T – to design an online arbitration service. Here the knowledge workers are the representatives of the parties and the arbitrators themselves. Within the project simulation models were developed alongside a more conventional prototyping exercise and the relative effectiveness of the two techniques is discussed. Both make essential, but complementary contributions to understanding the evolving business activity.

However, the success of the E-Arbitration-T project, which used very simple models, could be attributed to specific characteristics of the arbitration process The second study presents the result of a survey activity and discusses the implications for building more sophisticated knowledge worker models.

Modelling knowledge workers

Before continuing, we must be clear about our understanding of the term "knowledge worker". It appears to have been introduced by Peter Drucker in the late 1950s. (Drucker, 1959). He described the knowledge worker as somebody:

> . . . who puts to work what he has learned in systematic education, that is, concepts, ideas and theories, rather than the man who puts to work manual skill or muscle.

However, "knowledge worker" has since been defined in a number of different ways, which means that there is no set definition for the term.

Kidd (1994) provides one of the clearer and more useful definitions. She says:

> … the defining characteristic of knowledge workers is that they are themselves changed by the information they process.

Although this is true of all human beings, Kidd says that this is the primary motivation of knowledge workers and the job that they are paid to do. Kidd offers further clarification of her definition by mentioning two groups of people, clerical and communication workers, who she does not consider to be knowledge workers. She argues that in both of these cases the information is external to the person and does not change/inform them. Communication workers arrange information from other sources in order to change other people's understandings and beliefs, rather than their own. The job of the clerical worker is to know what information to use in different situations and apply it in a way that will produce consistent outcomes.

Kidd's distinctions are important when we turn to the field of business process modelling (BPM). Malhotra (1998) defines a business process as "a set of logically related tasks performed to achieve a defined business outcome". Traditionally, simulations have been used in the manufacturing industry and generic components have been developed for conveyor belts and other machine parts (for example Law and McComas, 1998; McClean and Shao, 2001; Rehn, 2001). Simulation models engender improved understanding in other fields as well. For example, Eldabi *et al.* (1999) and Swisher *et al.* (1999) show that it improves understanding of healthcare procedures, and Hegarty and Bloch (2002), describe the use of simulations to train intensive care unit staff.

However, despite the availability of a variety of tools for BPM, companies still face problems when trying to model in detail the way in which knowledge workers operate. Giaglis *et al.* (1996) list several reasons for this, including the complexity of most real-world business processes, the different perceptions of users regarding the way in which work is done and the interdependencies between tasks. Simulation models rely on being able to predict the time taken for an activity and to clearly identify the resources (workers and machines) needed to carry out a task. All resources are assumed to belong to pools where any machine or worker within the pool has the ability to carry out the task. They also make the assumption that; in general, once a task is initiated it will be completed. While this may be true of clerical tasks it would not be true of something like a lawyer preparing a brief or an architect preparing plans for a house. This, Kidd's definition of knowledge work identifies precisely the roles that present problems in BPM.

To handle the nature of knowledge worker behaviour in our models they are represented as full-blown entities rather then ubiquitous resource pools. The basic activity life-cycle for our knowledge workers is to await the arrival of tasks (represented by project files) within an "in tray" (see Figure 1). Once a knowledge worker has at least one such task available, they evaluate the job priorities and select a task to work on. They then work on this for a "session" before returning it to the "in-tray" and re-evaluating the priorities.

Given this infrastructure it is now possible to incorporate the knowledge worker within a conventional discrete-event simulation model. This passes work across the

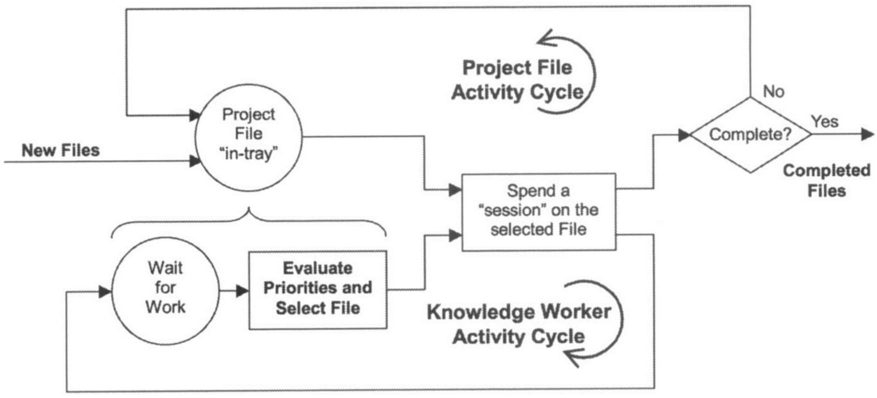

Figure 1.
Core knowledge worker activity cycle

knowledge worker's desk accumulating information about the level of activity, costs and delays as tasks are carried our. Once a knowledge work activity has been completed the project files more through a conventional clerical, workflow and distribution model until the project again needs input from a knowledge worker.

The critical process in this cycle is the evaluation and selection in the grey box. This needs to identify files pending attention for the particular person and reflect the way they prioritise and allocate time to the task. Although most simulation systems have building blocks reflecting the other components in the diagram this element is likely to need additional programming. Our first implementation of this model used a very crude strategy of taking files on a shortest time to deadline basis and working in roughly two-hour sessions on them until they were completed. Thus the behaviour was to work on the most urgent task until the next coffee, tea or lunch break.

This model formed the basis for our evaluation of arbitration service in the project – E-Arbitration-T.

Online arbitration

However good we are at planning and setting up clear contracts with our business partners, the time will come when unexpected events lead to a dispute. If we cannot agree about what to do then we turn to someone else for help. An alternative to the state's courts is for the disputing parties to agree that some trusted third party – an arbitrator – can resolve the dispute for them.

With electronic commerce continuing to grow, companies are turning to international trade; and one of the main aims of the E-Arbitration-T project was to develop an online system that could give them access to fair dispute resolution. In addition to delivering a software design the project also needed to provide guidance on adapting the current business processes (arbitration) to online working and providing prompt service at minimum cost. Thus the project needed to address business process change and ensure that the general principles of arbitration in the New York Convention (UN, 1958) were met.

A business process model for arbitration

The essence of an arbitration case is that each side of the dispute makes a careful argument as to the merits of its claim in front of the arbitrator. Having heard both sides, considered the evidence and allowed each party to comment on the other's argument the arbitrator (or a panel of arbitrators) rules on how the dispute is to be resolved. In practice most of this takes place through documentation although formal hearings may occur to wind up the process.

The workflow in any case is well defined and usually runs to strict published deadlines. Constructing a detailed workflow model covering the clerical and distribution processes presents no problem and this can be integrated with the knowledge worker elements as shown in Figure 2. The basic model for both the arbitrators' offices and the representatives' offices is as described above. However, there is one modification to the representatives' offices that needs to be commented on. When a file is completed and documents are ready for return to the workflow system an artificial delay is added so that it is released just before the deadline. When asked why they don't send replies before the deadline one of the lawyers in the project replied, "Well if we sent the documents any later they would miss the deadline." The philosophy is that if the timetable gives 30 days to prepare a response then to take significantly less time considering the response would be negligent.

Integrating real-time meetings into the system does not present a significant problem. The people involved simply leave their offices (the normal knowledge worker cycle) and go to the meeting place. Travelling time and the duration of the meeting are modelled in the normal way reflecting an appropriate interruption to the desk bound tasks.

There are hundreds of arbitration service providers each with their own procedural rules. The BPM needed to be applicable to a range of providers and E-Arbitration-T therefore combined validation of the adaptation algorithms with exploration of the different business processes as shown in Figure 3. The information system prototype (bottom left) receives signals of human events and actions just as it would when deployed as a full implementation. It responds with status and guidance information and thus influences the behaviour of the people and business processes components represented in the discrete event model (top left).

A collation process occurs at the knowledge worker's office when no immediate action is required. The documents are merged into the case file but it is not placed in the knowledge worker's in-tray for action until all the documents required for the next

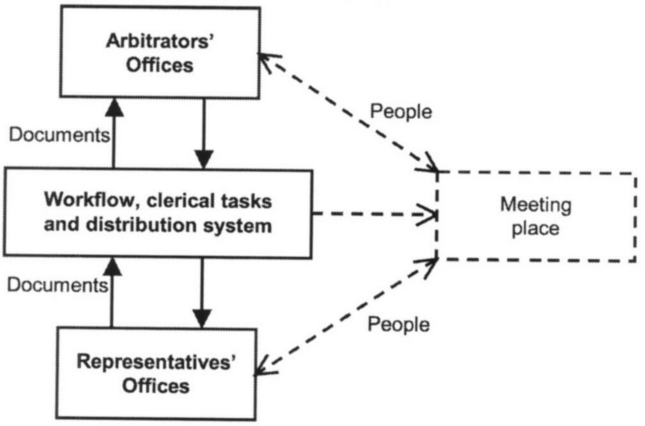

Figure 2.
A human activity system for arbitration

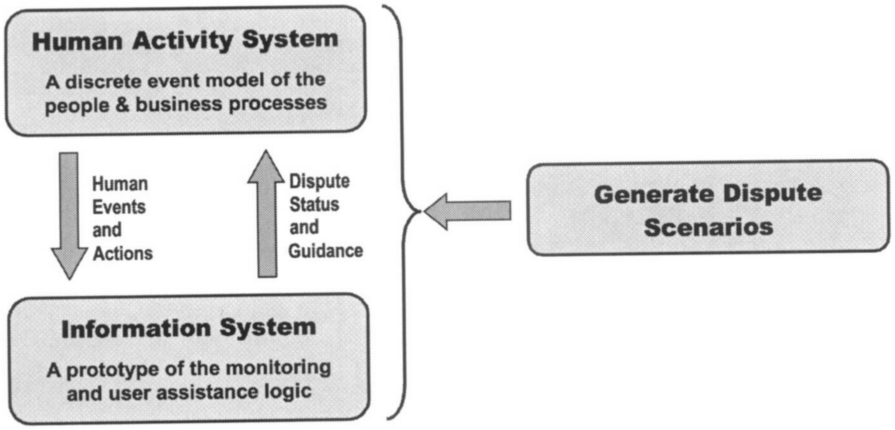

Figure 3.
Combined IS and human activity test architecture

action are to hand. Each document (or piece of information) flowing through the system is assigned a type code that defines the copying, collation and addressing rules as it is sent round the system. The type code also indicates which information is generated by the next actions of an appropriate knowledge worker.

The third major element of the system architecture (Figure 3) is the scenario generator. This is a piece of software capable of generating different dispute scenarios at random and injecting the details into the model. This enabled us to either generate a small number of coherent cases to evaluate a particular problem or issue or to run stress tests by generating a significant number of mixed complexity cases. The scenario generator also controls a number of switches within the simulation so that the workflow could be routed through physical channels or handled on line at will. Making this dynamic allowed us to explored scenarios where only part of the workload has been moved to an online environment.

Exploring the arbitration process

This use of the test architecture offered several advantages over field trials. The arbitration process is lengthy, a case with no major delaying issues runs for 200 to 350 days, and involves different users of the system taking independent action on an unfolding view of a case. Validation of software and new business processes through beta testing with real cases would take years and it would be difficult to get committed users. By the time organisations take a dispute to arbitration a level of distrust has grown between the parties and they will be risk averse in selecting any resolution procedure. Artificial testing in some form is the only realistic option.

If we want to evaluate the new rules and procedures we need to recreate this independence of view among the participants. This gives us a simple test of adequacy for the new business process rules and any supporting IS: if the workers only react to the information coming from the IS and the case unfolds to a successful conclusion we can be confident that the design meets its purpose. On the other hand if the process stalls or participants get confused and respond incorrectly then we can conclude that the process is inadequately specified or implemented.

In a pilot study with real people this would require several players each able to interact with and respond to the software without being located in the same place at the same time or being able to discuss the case. Nor should the participants in such a study have prior knowledge of the outcome in the dispute. It is important to explore the procedure with participants who have not been exposed to the design discussion within the project team. However, there is another aspect of arbitration that makes a "real user" pilot difficult or impracticable as a way to validate the system. To assess its ability to adapt it needs to be used with a wide range of different rules and case structures. Getting together enough different cases, all with unknowing users would need a major investment in time and manpower.

A simulation experiment is no substitute for real user trials if the questions under investigation are related to GUI specifications or other details of user-computer interaction. Clarity of screen design, comprehension and ability to find information quickly are not assessed because interaction is at an abstract level. Although the simulation models can explore the consequences of user mistakes they cannot assess whether a GUI design is more or less likely to induce such mistakes. The E-Arbitration-T project used a two-pronged approach to validating the system concepts. While the Brunel team developed the simulated BPR models the French partners produced a pilot system for hands on trials with the lawyers.

The pilot system experiments enabled us to explore many issues that arose at the point of interaction with the computer and to develop the interfaces so that the case and system state was transparent. It was very noticeable in these experimental sessions that the target users quickly became involved in the detail of the interface and lost sight of the overall case development. The easily lost sight of the fact that for certain participants in the process there would be gaps of weeks between their needing to become re-involved with a particular dispute.

Working with the simulated model we were able to stand back and take an overview of the complete arbitration process. The modelling software was designed not only to produce statistical data but also to print case descriptions and diaries in a form readily understandable to the target users. From these we could study how the procedure unfolded and identify long-term problems instead of the immediate ones. The simulated environment thus allowed us to devise mutually supportive procedural rules and software to improved the process in ways which were not possible using the pilot study. The model offered other specific advantages:

(1) The experiments were not constrained by real time delays and a many different scenarios could be presented. Trial runs with mixed complexity cases (typically 400 cases spread over a two-year period) could be completed in a day or so.

(2) The procedures under test could be changed or corrected and re-evaluated with exactly the same set of scenarios.

(3) The simulated "users" could be genuinely naive in both their response to guidance and their foreknowledge of the outcome in a particular case.

The simulated naivety of users (point 3 above) is particularly powerful in unveiling inappropriate or missing procedural guidance. As human users our memory and common sense often steps in so that we do the right thing despite pointers to the contrary. At one point in early testing the software miscalculated the due dates for a document – instead of ignoring the clearly wrong information the simulated user duly took the whole two years to contemplate the contents of the document before submitting it to the tribunal! Had this happened in a real user trial the participants, having become accustomed to the sequence of events in earlier experiments, may well react on the basis of earlier experience rather than notice the bizarre timing in the new system.

As well as supporting the development of new procedural rules the modelling activity also provided insights into the process as a whole. We were able to demonstrate that some advantages from new business processes were achievable even without the employment of ICT and that the potential value for arbitration service providers depended on their case management and charging strategies.

Modelling knowledge workers in other scenarios

Given the success of this approach to modelling knowledge work in the E-Arbitration-T project it is important to consider how effective the technique might be in other circumstances.

The party representatives (lawyers) in the arbitration system are very likely to be handling one or two arbitration cases as part of a broader workload involving other types of case. The same will be true of most arbitrators. They are frequently professionals in other disciplines (relevant to the dispute) with additional training and experience as arbitrators. We suspect that many knowledge workers have this characteristic of contributing to a variety of business processes and that in any business process model there will be a need to reflect the proportion of time devoted to the target business. This was not developed significantly within the E-Arbitration-T model because:

- For a party representative to over run a deadline would amount to professional negligence.
- Arbitrators cannot be compelled to act and are offered cases, which they may accept or decline. It is considered unethical to accept cases where it could not be given priority and dealt with expeditiously.

Many business processes will not be governed by such strict rules. In these cases the model will need enhancement to reflect interaction with "outside" tasks.

When the modelled business process is not at the top of the agenda for knowledge workers their strategies for prioritisation and scheduling become more critical and could have a significant impact on the new process being modelled. Many professional workers have a mix of tasks that can be divided into three different categories:

(1) *Scheduled*. These are tasks that are designated to take place at a particular time; such things as meetings, presentations and surgery. Provided the times are set in advance the model can simply capture them as periods of unavailability for the target process.

(2) *On-demand*. These are tasks that occur with no prior notice and require immediate attention; such things as help-desk support, unscheduled visits and on call services. These are again periods of unavailability. Provided the frequency and duration of such events are known they can be modelled in much the same way as a manufacturing model handles machine breakdowns.

(3) *At-will*. These are tasks of the form being considered here. They are characteristically individual activities where the knowledge worker engages with the business process for a significant length of time; such things as

drafting, designing, planning and analysing. Individual task selection strategies become critical here and this is explored in more depth in the next section.

In effect the mix of responsibilities will mean many knowledge workers only contribute "part-time" to the business process being modelled. However, unlike the part-time manual worker these periods of availability do not follow a regular pattern but are mediated by the individual knowledge worker's behaviour and perception of priorities.

The other factor that distinguishes knowledge work from more mechanistic tasks is the distinction between elapsed time and handling time. As indicated above the handling of documents (or information) from the party representatives involved placing a delay after the work was "complete" before it was returned to the workflow process. Thus for each task the elapsed time, while the knowledge worker task was in progress, could be significantly greater than the hours input by the knowledge worker, This was modelled very crudely by the delay until deadline in the E-Arbitration-T model but that type of completion logic is not necessarily typical of knowledge work in general. Instead of the single time to complete, as in a manual task, a knowledge work task has separate, but related, working and elapsed times.

Accurate modelling of all of these phenomena will be important if the model has to show the consequences of overloading the knowledge worker or overrunning deadlines.

A survey of knowledge worker strategies

To gain some further insight into the variation between knowledge worker behaviour one of the authors conducted a series of interviews. Fourteen people were interviewed, including researchers, academics, software engineers and solicitors, in order to provide as general a picture as possible. The interviewees also varied in age group and sex (four females and ten males). All the interviews followed the same agenda (see Appendix) and were recorded. The transcripts were then analysed under the following subject headings:

- Structure of a typical day or week.
- How workload is prioritised and how frequently it is done.
- Types of activities undertaken and how they relate to each other.

Structure of a typical day or week

Most interviewees believe that each working day is different and dependent on the daily workload. Many have project-related jobs, which means that their workload varies week by week. Additionally, the support-related aspects of some jobs meant that urgent bug fixes might be required at a moment's notice. Students also provided an unpredictable aspect for the academics because students may request help *ad hoc*. To cope with this, some academics dedicate weekly timeslots to students. In contrast, interviewee 14 (a solicitor) considered every day to be similar reflecting the procedure-driven nature of the job.

The interviewees are roughly evenly divided on the ways they mix their work during each day. Many want to work on one task at a time, but the nature of their job means they are often disrupted or given other urgent work to complete, e.g. one interviewee employed to develop software also provided systems support. In practice six of the interviewees tended to see a task through to completion and the others switch between two or three tasks each day (some switching deliberately to keep the day interesting).

The only daily task identified by most interviewees is the checking of e-mails (or post) at the start of each day. E-mail is seen as a key way to communicate and receive tasks and information. After checking their e-mails, the interviewees then looked at what they need to accomplish that day. One interviewee tried to allocate one "e-mail free" day every week, highlighting the fact that e-mail processing is a significant task. Four non-academics also fill out daily timesheets so that their managers can monitor their time usage. Certain individuals hold or attend weekly team meetings to discuss their progress, forming part of the structure of a typical week. Others attend courses and client meetings, sometimes off-site, as and when needed.

The academics interviewed have more structure to their weeks than non-academics, because they must fit their other work around a teaching timetable. One interviewee also described an annual cycle of activities that need to be done. Therefore there is a structure to their year, rather than just a daily or weekly pattern.

Prioritising workload

The interviews show that the way work is prioritised depends on the individual and what may be deemed as important to one person could be viewed as less important to another, as one interviewee said:

> You need to consider the person receiving the document. Their priorities may be different to yours and so although you may view your document as very important, they may view your document as a low priority, when compared to the other tasks they must complete.

Similar statements are made by other interviewees: "Another thing to consider is how important is the task to you" and "We have different producers, who give us work. They don't realise that they are not the only one".

Table I lists the factors each interviewee uses to prioritise their work. This shows considerable variation with some individuals balancing a range of factors and others with a clear-cut method for prioritising or even only single factor. However, it is possible to identify some common factors:

(1) *Deadline driven*: this strategy gives the highest priority to tasks with the shortest time left to the deadline. Choose the task with the nearest deadline.

(2) *Importance to customer*: this assigns priority to the task by assessing the importance of the recipient rather than the task. Choose the task for the most important customer.

(3) *Task length*: by the choosing the shortest task first this strategy seeks to maximise the number completed.

(4) *Importance to business*: this strategy takes a view of the importance or urgency within the business as a whole rather than the individual. Choose a task on the critical path or a task that is holding up someone else.

(5) *Priority set by others*: this is only a strategy in so far as the individual knowledge worker accepts the prioritisation strategy of another person – usually their manager.

Person	Prioritising factors	Frequency
1	Time it will take Date due Who it is for	Every morning
2	Choose closest deadlines first Higher priority to tasks that people are waiting for Length of task (plan mornings/days for long tasks)	Weekly
3 and 4	Choose tasks due soon Driven by pressure	Infrequently – tend to know in mind what needs to be done when
5	Experience What they deem to be important	Plan in mind – are aware of what deadlines are coming up
6	Project manager assigns the priorities	N/A
7	Current level of interest in the task – choose interesting tasks first Pressure from the customer – more pressure means higher priority Importance to the project – if on the critical path the task gets done first	Weekly
8	Deadline Importance of the customer Length of the task – try to do the shortest tasks first to get them out of the way	Daily
9	Critical task priorities are dictated by project manager For non-critical tasks Who it is for How difficult How critical it is	Daily
10	Choose simplest task first Things that affect others Project manager dictates most priorities	Daily
11	Needs of business Importance of customer	Daily
12	How difficult the task is Length (choose quickest first) Deadlines Pressure created by the task Importance of the customer	Weekly
13	Assigned by manager	N/A
14	Urgency of the task	*Ad hoc*

Table I.
Prioritisation strategies

The first three factors – deadline, importance of the customer and length of the task – are the most popular with academics and non-academics alike. Several interviewees combine two or three together to prioritise the tasks. Sometimes this is a simple cascade where another factor is introduced simply to resolve ties but in other cases the way they combine is more complex and less transparent.

Some academics used the length of task factor in a different way. They face a particular problem because significant timetabled commitments will mean some days are heavily broken up by lectures or other activities. By considering the length of the tasks, they can allocate longer pieces of work to areas of the week when they have more free time, e.g. interviewee 2: "It all depends on whether I have lectures that day. I will probably try to do a fairly large piece of work in the morning". This might be considered as a strategy to maximise efficiency by minimising task switching.

One main difference between academic and non-academic interviewees is that some non-academics do not prioritise their own work; their managers dictate what needs to be done. The non-academics also frequently consider the needs of the business to be an important prioritising factor, reflecting the unpredictable nature of their jobs, e.g. supporting systems.

Only nine of the interviewees prioritise their own work regularly. Most re-assess their work on a daily basis (at the start of each day), which highlights the fact that they can receive new work every day. Some also mentally re-prioritise their workload when new tasks arrive. Some interviewees prioritise weekly "to get a picture of the week ... because the days are cut up by the lectures" (interviewee 2) and because "I tend to get work on a weekly basis" (interviewee 7).

Activity types and relationships

The types of tasks identified by the interviewees can be divided into three categories:

(1) one-off tasks, e.g. review a journal article;

(2) project-related tasks; and

(3) on-going or seasonal tasks, e.g. maintaining systems and planning exams.

Some academics identified tasks that have to be performed at the same time each year – for example writing exam papers and reviewing course material. For them other projects, such as research, get slotted in as and when they arise. In contrast, most of the work undertaken by non-academics takes place in a project context with much less of an annual cycle. The types of tasks include specification, design and testing documentation. Other tasks, such as bug fixes, need to be classified as on-demand tasks, which must be dealt with as and when they arise. Only one interviewee had a simple task structure that just required delivery of one main document, detailing all of their findings, at the end of each study.

The interviewees work individually on tasks and so the nature of collaborative work could not he explored in the interviews. Interviewees sometimes send out documents to he reviewed and the feedback determines whether the documents are acceptable or require amending but most other task dependencies identified by the interviewees were related to the project lifecycle. Most deadlines are set by the academic timetable, by other people (for example the project manager) or are commitments made in the contract with the customer or research funding body.

In general, the interviewees did not take on more than two large projects at a time. However, although academics will often only work on one major research project at a time they are frequently involved in planning and preparation of bids for their next piece of research. In the non-academic group this self-determined work pattern was less evident with acceptance of new contracts or projects being undertaken at a higher level of management.

Based on the survey findings a new knowledge worker model was created and evaluated as an isolated component. This was used with a random input of on-demand

and at-will tasks to create a demand on a pool of knowledge workers. Each worker dealt with on-demand tasks on a first-come-first-served basis before employed a single strategy – deadline driven, importance to customer, or task length – to select their next at-will task. This was used to generate daily diaries for the knowledge workers and the original interviewees were asked to comment on the diaries. In broad terms the diaries were credible but, using on one factor, they did not capture the sorts of choices most interviewees would have made.

In broad terms the diaries were credible but, just using one factor, they did not always capture the sorts of choices most interviewees would have made. For example in one case an employee (A) had three tasks to do and another (B) had just one task, which could not be started until A had finished one of theirs. However A's strategy did not prioritise the task B needed and it was left for three days while A worked on other things. The general opinion of the interviewees was that, regardless of how they say they prioritise, they would have tried to ensure that their colleague was not waiting around. Clearly task prioritisation is a multifaceted phenomenon that needs to be modelled in some depth if it is to be effective.

Conclusions

The success of the E-Arbitration-T project, both in developing new business process rules and in showing where the benefits could be expected, demonstrates the effectiveness in modelling knowledge workers as part of the investigation. However, the key characteristics of knowledge workers, as defined by Kidd (1994) and Elliman and Hayman (1999), mean that they cannot be modelled in the same way as production and clerical workers. The discussion above shows that the success of the E-Arbitration-T model is, in part, due to specific elements of the business process, which cannot be generalised to all knowledge work.

A richer model of knowledge worker behaviour is postulated and elements not necessary for the E-Arbitration-T model are identified. The knowledge worker's day was defined as being made up of scheduled, on-demand and at-will tasks, only some of which may relate to the business process being modelled. A particular question that must be addressed in this extended model is how to model the choices knowledge workers make between competing at-will tasks.

These choice strategies were examined in a small survey of academics and other knowledge workers. This shows that most knowledge workers, with a free choice of task selection, weigh a number of factors together in choosing where to place their efforts next. Four common factors emerge from the survey. Two factors – deadline and length of task – are properties of the task itself, but the other two could be called "environmental" factors. These latter two place a priority on the task according to the customer (or recipient of the task) and according to its value to the organisation as a whole. Several of the knowledge workers interviewed claimed to combine these factors into more complex strategies. This was confirmed by some simple experiments where it was demonstrated that applying one factor alone couldn't generate a credible daily work pattern.

Knowledge workers enjoy significant autonomy and frequently negotiate or determine the tasks they later have to perform. This affects two elements of our model in ways we cannot yet predict. All meetings with colleagues are subsumed under the heading of scheduled tasks – that is those where the time, and probably the duration, is fixed in advance. However, the worker may well have been involved in the negotiation and has the opportunity to influenced the time and duration of a meeting to reduce its impact on other tasks in their workload. Some responses in the survey also exposed the fact that, within limits, knowledge workers can choose the length of time they spend on a task. Thus they may extend or reduce the time spent trading off quality of product on one project for more time spent on another.

The two pieces of work reported here have generated a rich model of knowledge worker behaviour ready for application and refinement in further business process modelling studies. However, research remains to be done, particularly in the areas where the knowledge worker is free to modify the business process in the light of their other commitments.

References

Drucker, P.F. (1959), *Landmarks of Tomorrow*, Harper, New York, NY.

Eldabi, T., Paul, R.J. and Taylor, S.J.E. (1999), "Computer simulation in healthcare decision making", *Computers & Industrial Engineering*, Vol. 37 No. 1-2, pp. 235-8.

Elliman, T. and Hayman, A. (1999), "A comment on Kidd's characterisation of knowledge workers", *Cognition. Technology and Work*, Vol. 1 No. 3, pp. 162-8.

Giaglis, G.M. (2001), "A taxonomy of business process modelling and information systems modelling techniques", *International Journal of Flexible Manufacturing Systems*, Vol. 13 No. 2, pp. 209-28.

Giaglis, G.M., Paul, R.J. and Doukidis, G.I. (1996), "Simulation for intra- and inter-organisational business process modelling", *Winter Simulation Conference Proceedings, 1996*, IEEE, New York, NY, pp. 1297-304.

Grover, V., Fiedler, K.D. and Teng, J.T.C. (1994), "Exploring the success of information technology enabled business process reengineering", *IEEE Transactions On Engineering Management*, Vol. 41 No. 3, pp. 276-84.

Hegarty, M. and Bloch, M. (2002), "The use of simulators in intensive care training", *Current Anaesthesia and Critical Care*, Vol. 13 No. 4, pp. 194-200.

Hide, A. (2000), "Failure? Who Says?", *Computer Bulletin*, July, pp. 26-7.

Hochstrasser, B. (1993), "Quality engineering: a new framework applied to justifying and prioritising IT investments", *European Journal of Information Systems*, Vol. 2 No. 3, pp. 211-23.

Irani, Z. and Love, P.E.D. (2000), "Investment evaluation: drawing lessons from a failed IT/IS implementation", *Journal of Management Information System*.

Kanellis, P., Lycett, M. and Paul, R.J. (1999), "Evaluating business information systems fit: from concept to practical application", *European Journal of Information Systems*, Vol. 8 No. 1, pp. 65-76.

Kidd, A. (1994), "The marks are on the knowledge worker, CHI'94: celebrating independence", *Proceedings of the Conference of Human Factors in Computer Systems*, Association for Computing Machinery (New York), Boston, MA, pp. 186-91.

Law, A.M. and McComas, M.G. (1998), "Simulation of manufacturing systems", *1998 Winter Simulation Conference Proceedings*, Vols 1 and 2, IEEE, New York, NY, pp. 49-52.

Macarthur, P.J., Crosslin, R.L. and Warren, J.R. (1994), "A strategy for evaluating alternative information-system designs for business process reengineering", *International Journal of Information Management*, Vol. 14 No. 4, pp. 237-51.

McClean, C. and Shao, G. (2001), "Simulation of ship building operations", *2001 Winter Simulation Conference*, IEEE, New York, NY, pp. 870-6.

Malhotra, Y. (1998), "Business process redesign: an overview", *IEEE Engineering Management Review*, Vol. 26 No. 3, pp. 27-31.

Paradi, J.C., Marr, C., Wong, M. and Yan, G. (1995), "Simulation-based alternative strategies in risk and service guarantee tradeoffs in the international funds transfer system at a large Canadian Bank", *IEEE Transactions on Engineering Management*, Vol. 42 No. 4, pp. 387-96.

Paul, R.J., Giaglis, G.M. and Hlupic, V. (1999), "Integrating simulation in organisational design studies", *International Journal of Information Management*, Vol. 19 No. 3, pp. 219-36.

Pinsonneault, A. and Rivard, S. (1998), "Information technology and the nature of managerial work: from the productivity paradox to the Icarus paradox?", *MIS Quarterly*, Vol. 22 No. 3, pp. 287-311.

Rehn, G.D. (2001), "Simulation in daily factory operation: setting the line boggey in Augusta", *2001 Winter Simulation Conference*, pp. 1380-4.

Swisher, J.R., Jacobson, S.H., Jun, J.B. and Balci, O. (1999), "Modelling and analysing a physician clinic environment using discrete-event (visual) simulation", *Computers & Operations Research*, Vol. 28 No. 2, pp. 105-25.

Taylor, A. (2000), "IT projects: sink or swim", *Computer Bulletin*, January, pp. 24-6.

Tumay, K. (1995), "Business process simulation", *Winter Simulation Conference (WSC'95)*, SCS, Washington DC, pp. 55-60.

UN (1958), *Convention on the Recognition and Enforcement of Foreign Arbitral Awards"*, No. 8, March 2003, United Nations, New York, NY.

Walsham, G. and Waema, T. (1994), "Information-systems strategy and implementation – a case-study of a building society", *ACM Transactions On Information Systems*, Vol. 12 No. 2, pp. 150-73.

Appendix. Interview agenda

(1) What is your job title?

(2) Do you manage other people or departments?

(3) What does your job involve – types of broad activities or one-off tasks?

(4) What does each activity consist of → break down into smaller tasks?

(5) Roughly what percentage of your time is spent on the activities you have identified?

(6) Who assigns your activities?

(7) What does a typical day or week consist of? Is every day the same or different?

(8) Do certain activities have to be done each day?

(9) How many long-term activities (weeks or months) do you have on the go at any one time? What is the maximum number that you could take on?

(10) How do you keep track of your progress on your long-term activities?

(11) How much time do you put aside for unplanned events?

(12) How do you prioritise the work you have to do? Is this done on daily/weekly basis or *ad hoc*?

(13) How do you determine if you can take on new work?

(14) Please could you talk through the stages that were involved in a recent long-term activity, e.g. a recent project? How long did the stages last, how many man-hours were involved, what deliverables were required, did the stages involve interaction with other parties, e.g. meeting with users to collect requirements?

Note: An activity is a piece of work that can be broken down into stages and individual tasks. An example may be the design of a new report, which would involve spec, design, testing phases etc. A task is an indivisible piece of work, e.g. clearing down a database for the test team or writing a document.