Acquiring Knowledge for Business Process Re-Engineering

Hugh Cottam, Nigel Shadbolt and Nick Milton

University of Nottingham, Artificial Intelligence Group,
Department of Psychology, University Park,
Nottingham NG7 2RD. United Kingdom
hdc, nrs, nrm@psychology.nottingham.ac.uk

Abstract

The field of Business Process Re-Engineering (BPR) is aimed at enabling the large scale re-design of processes within organisations. BPR initiatives are by nature highly knowledge intensive activities. In this paper we argue that the knowledge based nature of BPR has not previously received sufficient recognition. We explain how BPR initiatives can be assisted through the use of techniques and tools that have their origins within the knowledge engineering community. In particular, we demonstrate the incorporation of knowledge acquisition (KA) techniques and the use of ontologies within a toolset that supports BPR. This toolset, named the Structured Process Elicitation and Demonstration Environment (SPEDE) has been developed to support both the acquisition and management of knowledge during BPR. SPEDE is currently being applied and validated within the aerospace and automotive industries.

Introduction

The concept of Business Process Re-engineering (BPR) has undoubtedly had immense impact on business thinking in recent years. There are many large organisations that claim to be embracing BPR and achieving significant benefits e.g. AT&T, Ford, Texas Instruments and Rank Xerox (Burke & Peppard, 1995). However, BPR has a high associated failure rate, and is far from being an exact science. It is a vague concept that encompasses a plethora of explanations and definitions. The BPR literature, as a whole, is often contradictory. Individual texts tend to reflect a particular consultant's preference, and the accompanying methodologies are normally expressed at a very high level. The result is a lack of detailed guidance as to how high level re-engineering goals can be achieved. This has often led to high cost, high risk BPR initiatives. It is our claim that this lack of guidance originates in a poor understanding of the knowledge requirements of a re-engineering initiative. This in turn leads to unconstrained and unguided knowledge acquisition (KA), followed by poorly planned and unsupported knowledge management. Because many of the problems associated with doing BPR are knowledge based, the experiences of the knowledge engineering community can be usefully brought to bear in this context.

The Structured Process Elicitation and Demonstration Environment (SPEDE) is a methodologically grounded toolset that provides support for BPR. It guides knowledge acquisition and knowledge management throughout the lifecycle of a BPR exercise using a combination of features. In particular SPEDE concentrates on the use of KA techniques and ontologies to support the acquisition and organisation of process knowledge during BPR.

Business Process Re-Engineering (BPR)

There exists an extensive literature addressing the subject of Business Process Re-engineering (BPR). An analysis of the methods described in the literature (Bach *et al*, 1996) reveals both the number of methods in existence and the extent to which they contradict each other. However, there are general observations that we can make about these methods. Two key concepts stand out as providing a unity to the field and distinguishing it from other business improvement approaches:

- The provision of dramatic improvements in an organisation's performance, as measured by such factors as cost, quality, and time.
- A focus on process (as opposed to, say, the structure of the business).

There is considerable debate within the BPR literature as to what a process is. One definition is that a process is "a sequence of activities performed on one or more inputs to deliver an output to the customer of a process" (Talwar, 1994). Typical examples of business processes are new product design, component manufacture, or sales and distribution. However, simply stating that we are taking a process perspective does not tell us what it is that we must know about these processes. The precise knowledge content of the models that we construct of these business processes will depend upon our goals at a particular stage within a BPR initiative. Uncovering this content, outlining generic aspects of processes, and understanding BPR goals are key research aims of the SPEDE project.

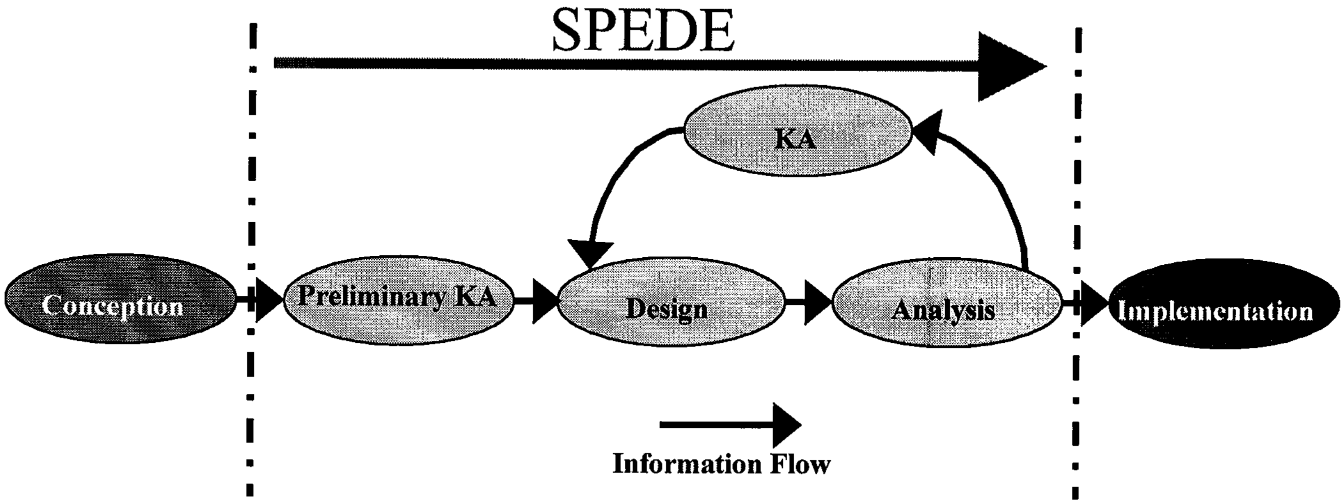

Figure 1: The SPEDE Design-Analysis Loop

SPEDE - Methodological Support

Within SPEDE we find it useful to characterise BPR as a design task. The main goal of BPR is to design (or re-design) a business process in order to improve some aspect of the business's performance. The fact that our design object is a process rather than a product, has important implications for how we provide support for such a design task. Rather than prescribe a rigid methodology for BPR, SPEDE attempts to identify possible sub-tasks within process design and provide tools and techniques to support those subtasks. BPR methodologies commonly make an AS-IS/TO-BE distinction. An AS-IS model captures some aspect of an existing process, whereas a TO-BE model prescribes an aspect of an intended new process. BPR methodologies tend to recommend either an AS-IS or a TO-BE emphasis, such that designing a new process should either commence by, modelling and then analysing the AS-IS or by synthesising the TO-BE from scratch according to certain requirements and constraints. The SPEDE approach discards this traditional AS-IS/ TO-BE process dichotomy, and replaces it with the concept of an iterative design-analysis loop that includes knowledge acquisition.

The SPEDE design-analysis loop shown in Figure 1 is a top level decomposition of the activities associated with a BPR initiative. SPEDE supports those activities within a BPR initiative that occur after the initial conception of the initiative, and before the actual implementation of the new process. The conception stage supplies requirements and constraints for SPEDE. These initial inputs are used during a preliminary KA activity that involves further requirements elicitation along with more general KA. These act together as inputs to the first pass through the design activity. The resulting first pass process definition then enters the design-analysis loop. Each design step is followed by an analysis step. Every iteration back through the design stage will be preceded by some form of KA activity. The SPEDE method will eventually produce as a final output, a new process description in such a form that it can be implemented.

This high level method description is based on the idea that BPR can be broken down into a number of types of activity. Within SPEDE our initial breakdown includes the activity groupings; preliminary KA, design, analysis and KA. These groupings are used in a very broad sense; e.g . analysis could include checking against original requirements, critical path analysis, computer based process simulation, or presenting a proposal to implement the new process to senior management. It can be seen that SPEDE is not prescriptive in its methodology. Instead, it aims to provide the components with which users can define for themselves the organisation of the activities associated with their own BPR initiative. Advice on how to do this is provided in the form of a hypertext based Principles And Methodology Support system (PAMS).

Assistance in the organisation of BPR activities is also provided in SPEDE through the use of knowledge requirement templates. These templates each consist of a KA activity that is performed in order to enable a design activity, that precedes a particular analysis activity. The templates include a definition of the knowledge that must be gathered during the KA activity. In this manner they define the knowledge requirements of design-analysis activities within BPR. These are explained in greater detail later.

In addition to the high level methodological support we have described, SPEDE also provides tools, techniques and advice that support the actual activities within BPR. Within this paper we shall concentrate on describing those tools within SPEDE that support KA. Our emphasis is upon software-assisted knowledge acquisition, and the tools and methods we describe have their origins within knowledge engineering. The SPEDE KA tools are customisations and additions to the PCPACK knowledge engineering workbench (http://www.isl.co.uk/pc_pack.html) which is a commercial package developed by Epistemics Ltd following research on the VITAL project (Shadbolt, Motta, and Rouge, 1993).

The SPEDE KA tools make particular use of ontologies as both an organising principle for knowledge and as a means of enabling knowledge re-use. The ontologies provide a structure not only for the organisation of knowledge once it is gathered but also for organising and guiding the KA activities themselves.

The Use of Ontologies

It is important that within our discussion of process knowledge re-use we do not lose sight of the overall goal. This goal is to design a better process with less effort. SPEDE has identified the knowledge gathering element of BPR as both a major drain on effort and as being an area where substantial improvement can be achieved. Re-use of process knowledge is one of the main mechanisms by which we believe this improvement in the quality and efficiency of knowledge gathering, modelling and design can be achieved.

Re-use of process knowledge can be achieved in a number of different ways; both in the form of the knowledge that will be re-used, and in the method by which that knowledge is re-used. Process knowledge re-use can occur at both a specific and a generic level. At the specific level, old process designs can be selected and re-used for a new process. This may involve a degree of adaptation of the old design to suit the new situation. At the generic level we may use process descriptions that are considered to be applicable across a class of situations. We can also use generic descriptions that act as templates to constrain the form and content of the knowledge gathered. A template is applicable across a class of situations but must be instantiated/populated with specific knowledge. Within SPEDE we refer to these as knowledge templates. Both generic processes and knowledge templates will require selection and may require some adaptation to the current situation. In summary, SPEDE has identified three forms of re-usable process knowledge: specific process descriptions, generic process descriptions and knowledge templates.

Ontologies have been chosen as a principal mechanism within SPEDE for process knowledge re-use. There are a number of reasons for this. Firstly, ontologies are a structured approach to the representation of generic knowledge. They can be used both to express generic process descriptions and knowledge templates. Secondly, an important issue is the selection and use of any re-usable process knowledge. This becomes problematic as soon as the knowledge repository grows to a reasonable size. Ontologies can be used as an indexing system for both generic and specific process descriptions, as we will show.

The Role of Ontologies

In spite of the increasing interest in ontologies there is still very little agreement on precisely what items they should contain and the manner in which they should be structured. The most commonly used definition of an ontology is that it is an "explicit specification of a shared conceptualisation" (Gruber 1993). This roughly outlines what an ontology is, but doesn't explain what it is for. The simplest answer would be that the purpose of any ontology is to perform some type of knowledge sharing function. This knowledge sharing can take on many different forms: knowledge sharing within software applications, between software applications, knowledge sharing in the form of re-use, knowledge sharing to arrive at a common understanding or to achieve integration. These many forms of knowledge sharing indicate the widely differing tasks and groups of agents that ontologies can support. It is this diversity in the purpose of ontologies that is a major reason for the lack of consensus on their content and structure. In order that they can be utilised in an optimal fashion the SPEDE ontologies have been designed to support specific task and user requirements associated with performing a BPR initiative.

The Structure of Ontologies

A survey of the recent ontological literature (Noy and Hafner, 1997) outlines the space of possibilities for the structure of ontologies. Ontologies can be taxonomically or axiomatically based. They can be constructed using one large taxonomy (e.g. CYC, Lenat and Guha, 1990), or they can be structured around a number of smaller taxonomies (e.g. TOVE, Gruninger and Fox, 1995). The concepts in these smaller taxonomies can then be linked by relations. The nature of the taxonomies themselves can also vary. The links in the hierarchy can be limited to just "is-a-subtype-of" relationships, in which case the division of concepts into subconcepts is disjoint (multiple inheritance will normally also be allowed). Alternatively the taxonomic breakdown can occur on a number of parallel dimensions. This latter approach means that the resulting categories cannot be disjoint unless a category is created for every possible combination of parallel distinctions. Any of the taxonomy based approaches can lead to taxonomies that are either sparse or dense at their top level, and can have different degrees of tangledness (though taxonomies using many parallel dimensions are likely to become very tangled). Ontologies may also vary in the level to which they are task dependent. CYC is a good example of an ontology that is task independent, where as the PHYSSYS ontology (Borst, Akkermans, and Top, 1997) is strongly linked to engineering tasks. There exists strong debate within the ontological engineering community as to the task dependent and independent character of ontologies.

The decision to commit to a particular structure for ontologies will depend heavily on the purpose to which the ontologies are applied. Within the SPEDE project this relationship between the structuring of ontologies and their goals has been explored more thoroughly.

A Process Ontology to Support KA in SPEDE

Below is the SPEDE general process ontology. The ontology given lists basic process object types along with pre-defined attributes and relations for that object type. All basic objects types can be organised in a taxonomy for that particular type using the "is-a-subtype-of" relation. The local taxonomies are then linked via relations, e.g. the "Has Data Input" relation links a task concept with a data concept. This is a base ontology that is designed to support process KA. We use the term KA here in a broad sense that embraces both knowledge elicitation and knowledge modelling. The ontology is designed to force the distinction during KA between basic process object types: task, data, result, role, organisational group, agent, product item, cost, location, resource, requirement. These basic process object types have been decided based on an extensive analysis of existing ontologies, enterprise and process modelling formats, and the knowledge requirements of the analysis activities conducted during BPR. We have listed the ontology (though not in its entirety) to illustrate its nature and scope. There are also additional concepts that are being considered for inclusion in the base ontology such as: skill, communication link, goal, metric, risk, quality. It is envisaged that the base ontology will undergo modifications following further experimentation.

The base ontology is closest in structure to the virtual enterprise ontologies from the TOVE project. It should however be noted that the SPEDE ontology is primarily aimed at supporting KA, whereas the TOVE ontology is aimed at the support of enterprise databases and deductive capability over such databases. The formal logic definitions have been omitted from our listing for the sake of brevity (These formal definitions are required in order to remove ambiguities).

Process

Concept ID:	
Name:	
Attributes:	Start Time, Finish Time, Duration
Hierarchy Relations:	Has Sub-Process, Has Parent Process
Sequence Relations:	Ends Before Starts, Starts Before Ends, Starts Before Starts, Ends Before Ends, Starts After Ends, Ends After Starts, Starts After Starts, Ends After Ends, Meets, Contains
Resource Relations:	Uses, Produces, Consumes, Releases
Data Relations:	Has Data Input, Has Data Output
Other Relations:	Is Performed By, Has Location, Has Result

Role

Concept ID:	
Name:	
Relations:	Has Members, Has Constraint, Has Group, Has Authority Link, Is Empowered, Has Goal, Has Agent, Requires Skill

Organisational Group

Concept ID:	
Name:	
Relations:	Has Role, Has Constraint, Has Group, Has Authority Link, Is Empowered, Has Goal, Has Location, Has Members

Agent

Concept ID:	
Name:	
Relations:	Has Role, Has Constraint, Has Group, Has Authority Link, Is Empowered, Has Goal, Has Location, Has Skill

Location

Concept ID:	
Name:	

Result

Concept ID:	
Name:	
Type Attribute:	Fail, Pass, True, False, Continue, Return

Requirement

Concept ID:	
Name:	

Data Item

Concept ID:	
Name:	
Type Attribute:	Electronic, Paper

Product Item

Concept ID:	
Name:	
Type Attribute:	Complete Product, Assembly, Component, Feature
Relations:	Has Sub-Part, Has Requirement, Has Constraint, By-Product of, Product of, Has Location

Resource

Concept ID:	
Name:	
Type Attribute:	Raw Materials, Facility, Tool, Operator, Space
Motility Attribute:	Mobile, Stationary
Consumption:	Continuous, Discrete
Unit of Measure:	Length, Area, Volume, Weight, Time, Manhours

Cost Per Unit:	
Relations:	Has Location

Table 1: SPEDE General Process Ontology

Knowledge Acquisition and Tool Support

In this section we describe how new KA tools have been developed to support the acquisition of process knowledge and to facilitate re-use. We then go on to show how this has allowed the utilisation of previously developed KA tools within a process context

Ontologies are used in SPEDE as the building blocks for the construction of well structured process descriptions. There are two levels of use for the ontology based tools in SPEDE:

- the **developer** accesses a library of ontologies using a purpose built SPEDE ontology editor in order to assemble ontologies that are more specific both in terms of task and domain. These are then handed to the modelling level user.
- the **modeller** browses the ontologies handed down to them from the developer in order to construct a knowledge description of some aspect of a process. As discussed earlier this will normally be part of a KA activity that fulfils the requirements of design and analysis activities at a particular stage within BPR. The ontology will reflect this intention.

This distinction between the two user levels corresponds with a marked difference in the level of ontological engineering competency expected of the developer level user and the modelling level user. This is reflected in the design of the tools that support each level of ontology use.

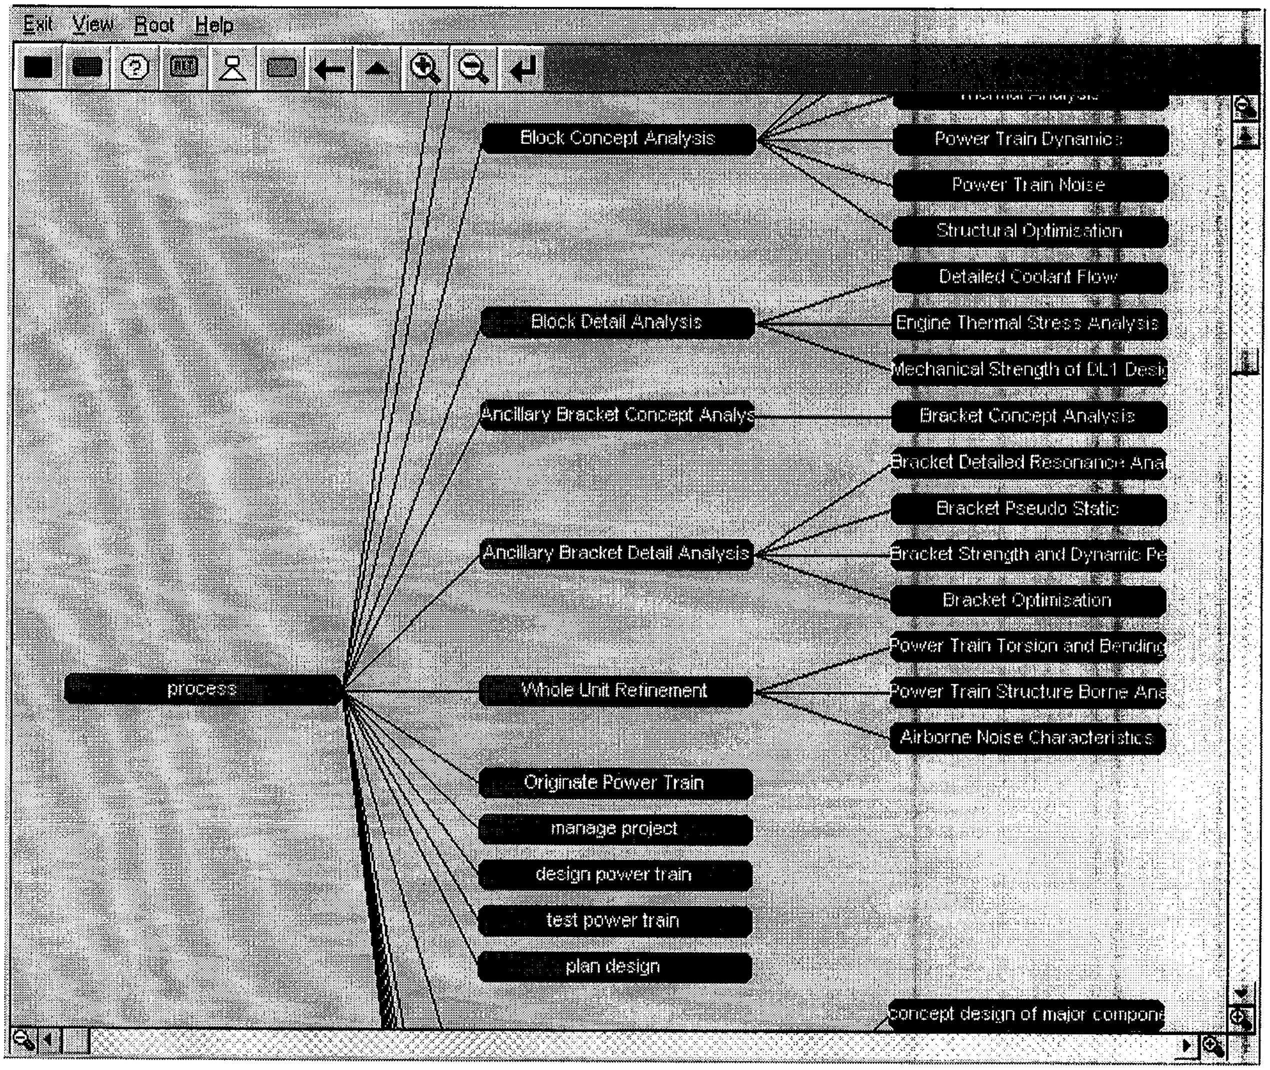

Figure 2: Ontology Editor

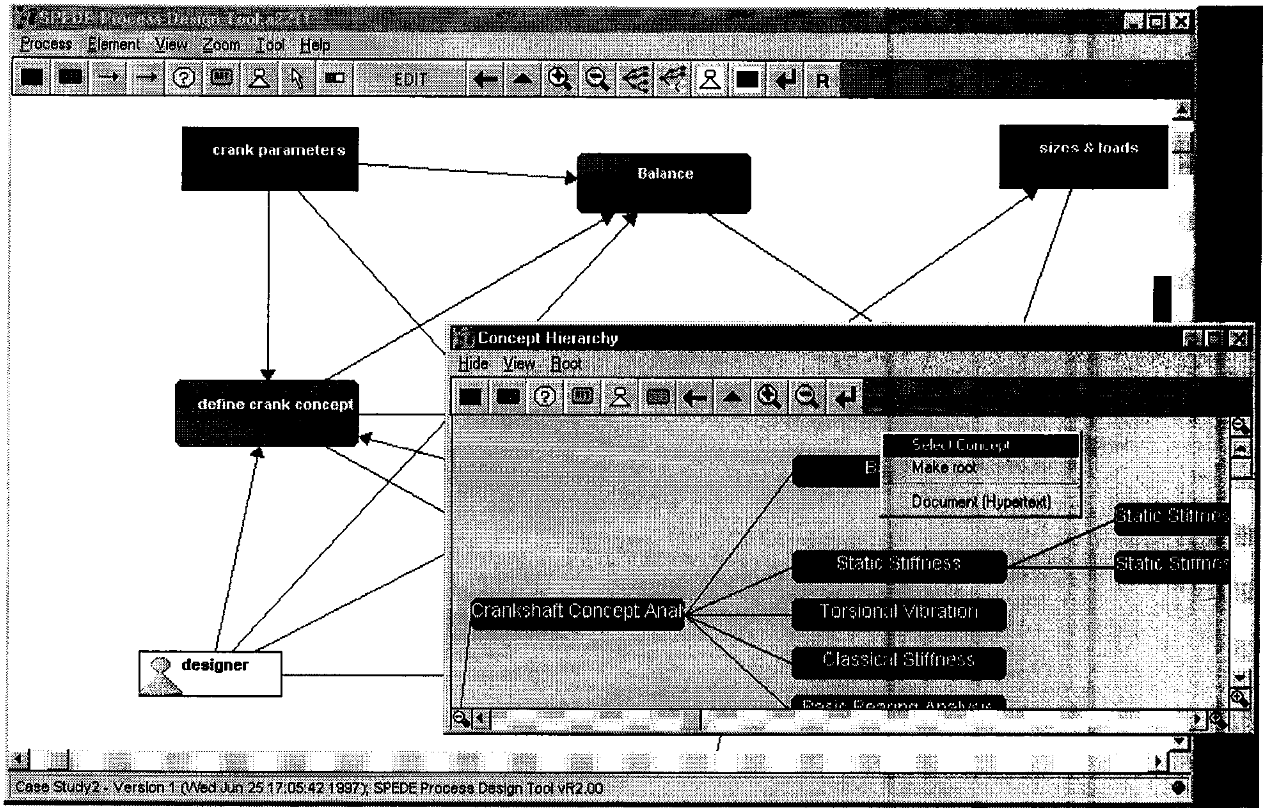

Figure 3: Process Design Tool

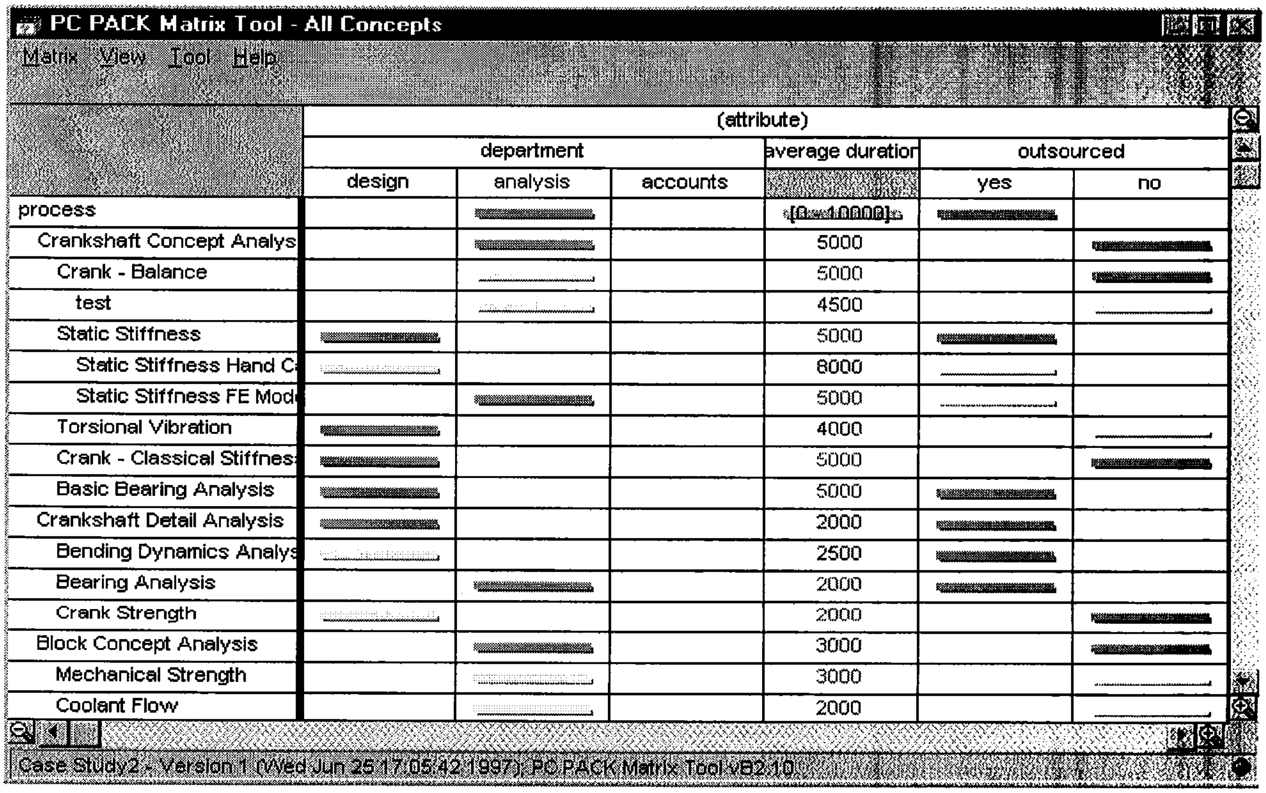

Figure 4: The Matrix Tool

SPEDE ontologies that support diagrammatic process descriptions make certain high level distinctions between concepts. At the highest level the primary distinction made within the ontology is the decision as to the basic object type: e.g. task, data, agent, result, e.t.c. This is the fundamental modelling commitment that must be made when adding any object to a process model, and is reflected in the design of the ontology editor pictured in Figure 2. The SPEDE ontologies supporting diagrammatic process modelling are a set of local taxonomies (for basic object types) that are linked by relations. This base ontology was listed in table 1. Figure 2 shows a task taxonomy selected within an ontology for modelling the analysis of designs within new product introduction processes. Other basic object taxonomies can be selected, including those representing data, agents, results, e.t.c. The rounded rectangles represent task concepts and the links represent "is-a-subtype-of" relationships. Numerical or categorical attribute values can be assigned to concepts, and these are inherited down the "is-a-subtype-of" relationship.

The tool supports the hypertext documentation of each concept, which enables the provision of advice and explanation to the users of the ontology. The ontology shown in Figure 2 represents a theory of design analysis that is specialised for the automotive industry. This includes tasks that perform design analysis, the data associated with those tasks, the agents that perform design analysis, the decision points within design analysis, and possible return values for design analysis processes. Figure 3 shows how the process design tool can be used to browse and access the ontology in order to build an instance model of a process that utilises design analysis. The figure shows the ontology browser being used to select a task concept from the ontology. An instance of this concept is then created in the process instance model.

The Link to Existing KA Tools. The ontology based KA tools that we have described, enable existing KA tools (i.e. those that are part of the PC Pack commercial toolset) to be applied in the process context. Three tools have proved particularly useful when applied in this manner. Firstly, the matrix tool shown in Figure 4 allows the user to

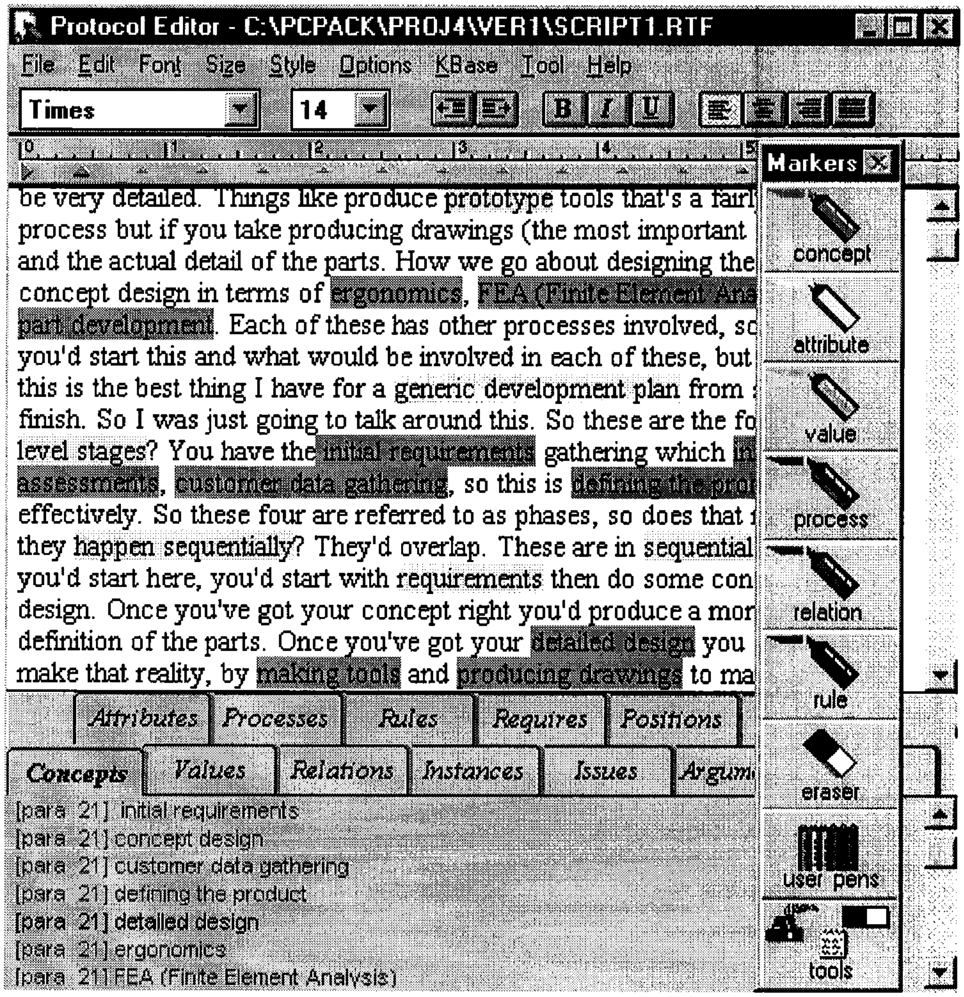

Figure 5: The Protocol Editor

assign attributes and values of those attributes, to both concepts and instances. Thus concepts that are present in the ontologies constructed using the ontology editor (shown in Figure 2) can have attributes and values associated with them.

In Figure 4 process concepts are listed top to bottom along the left hand side. The indentation of these reflects the hierarchy depicted in Figure 2. Instances can also be shown in this hierarchy, though in the example given they are not. Attributes are arranged in Figure 4 from left to right along the top of the figure. The attributes shown in figure 4 indicate the department that a process is usually performed in, the average duration of a process and whether a process can be outsourced or not. As can be seen the first and last of these have associated categorical values, whereas the middle one has a numerical value. The attributes can be arranged in a deep hierarchy though the ones depicted in Figure 4 are merely a flat set of attributes. The toolset also supports the inheritance of attribute values down the concept hierarchy. Thus, it should be noted that in Figure 4, a dark bar in a box indicates that a particular concept has that value for a categorical attribute. A light bar indicates that a concept has inherited that value from a parent concept. Numerical values are inherited in a similar manner though the colour coding is not easily distinguished in Figure 4. It can be seen that for the top level process concept in Figure 4 no value has been defined for "average duration". In this situation, the matrix tool displays the numerical range associated with that attribute. The toolset also supports multiple inheritance, and conflicting values are flagged in the matrix tool when a concept inherits different values from different parent concepts.

The second existing KA tool that can now be applied more easily within a process context is the protocol editor. This is shown in Figure 5 being used to highlight a transcript that was used to populate the tools shown in figures 2, 3, 4 and 6. The transcript shown has been elicited from an expert on the application of design analysis within processes. The tool allows the user to highlight the text with various pens. Different colour pens indicate that a particular section of text is the label for a concept, an attribute, an attribute value, e.t.c. Having labelled the text in this manner the tool can automatically generate these objects in an underlying database that can then be accessed by the other tools that we have discussed.

The third existing KA tool is shown in Figure 6. This is a repertory grid tool that enables the user to assess the similarity of both concepts and the attributes that are used to describe them. This is based on a technique known as

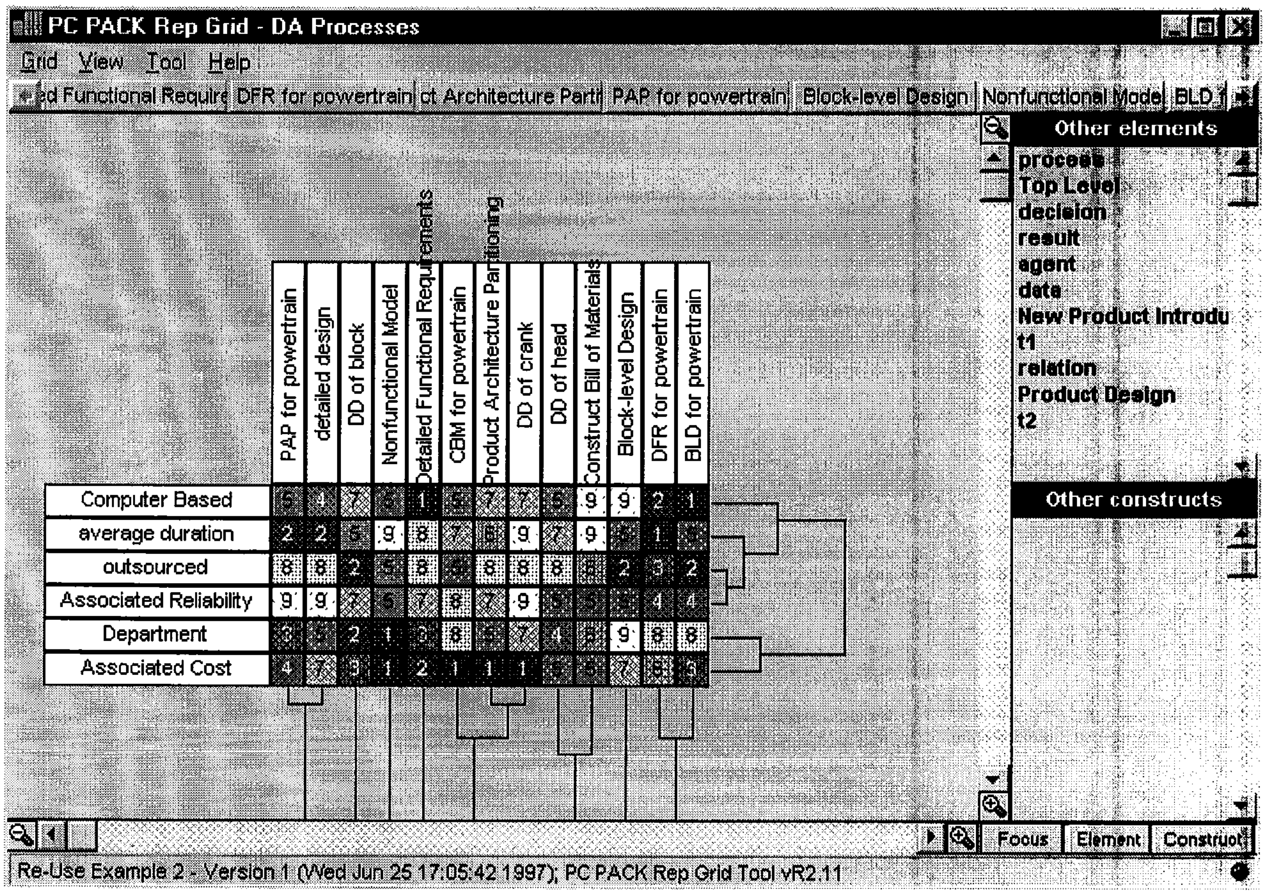

Figure 6: The Repertory Grid Tool

cluster analysis. Attributes must be chosen such that they have two poles representing the limits of a value on a scale from 1 to 9. The various concepts are then rated for each attribute using these scales. In the example shown in Figure 6, a number of process concepts are being compared, and these are arranged along the top of the figure. The attributes that are being used to compare the concepts are arranged along the left of the figure. The attributes used are as follows:-

Computer Based -- (Poles: not to heavily)
Average Duration -- (Poles: short to long)
Outsourced -- (Poles: never to often)
Associated Reliability -- (Poles: low to high)
Department -- (Poles: always analysis to always design)
Associated Cost -- (Poles: cheap to expensive)

The numbers shown in the grid in Figure 6 indicate the rating that concept has been given for a particular attribute. The lines that are shown at the side of the grid represent a dendrogram. A closer linking of the lines that are attached to either concepts or attributes indicates a higher degree of similarity between either those attributes or the concepts (based on the information given). This is a powerful KA technique that can be used as a prompt to elicit:

- new higher level attributes or concepts, to describe groups that are similar.
- new attributes to differentiate between concepts that have been analysed as similar based on the information so far given.
- correlations between different attribute values, in the form of rules.

Discoveries such as a stable super-ordinate class of concepts or attributes will then often lead to reformulation of the ontology concerned.

Knowledge Templates

The process KA within SPEDE is performed to meet the knowledge requirements of a particular stage within a BPR initiative. The content and structure of SPEDE process models will depend on a number of interdependent factors:

- The BPR goals may be to minimise cost, minimise time, maximise quality, or more likely to combine a number of these goals within more specific requirements and constraints.
- The analyses that will be performed on the process models also impose certain knowledge requirements.
- The knowledge that the process models must contain in order for implementation of the re-engineered process to take place.

As can be seen, there are a number of factors that impose knowledge requirements upon our SPEDE process knowledge acquisition. We can view the cycle of stages in a BPR initiative as representing a knowledge flow. By better understanding this knowledge flow we can better direct and constrain the knowledge acquisition throughout the BPR initiative. A simple example is the difference between the knowledge requirements of an analysis stage in BPR and those of the implementation stage. Clearly, the analysis is likely to require a much smaller amount of knowledge than the implementation. In addition to this the analysis might well result in extensive modification being necessary to the process design, therefore requiring additional KA. It would seem sensible to suggest that the KA associated with implementation not be performed until the process design is stable. However, this judgement must be combined with an estimate of the additional costs associated with performing KA in separate stages. The linking of high level requirements to the type of knowledge to be acquired can be explicitly supported within SPEDE using knowledge templates. Examples of ontology configuration provided in the form of knowledge templates would be the following.

A Knowledge Template for Critical Path Analysis. Table 2 shows a subset of the general process ontology. It provides a template for the knowledge required to perform critical path analysis on a process.

Process

Concept ID:	
Name:	
Attributes:	Start Time, Finish Time
Hierarchy Relations:	Has Sub-Process
Sequence Relations:	Ends Before Starts, Starts Before Ends, Starts Before Starts, Ends Before Ends

Table 2: Knowledge Template for Critical Path Analysis

A Knowledge Template for Role Definitions. Table 3 shows a template for the knowledge required to make role definitions associated with the tasks within a process. These role definitions also make use of knowledge about agents, organisational groups and the locations, constraints and authority links that are associated with these. Task sequencing knowledge is only required to ensure that too many roles are not assigned to concurrent tasks.

Process

Concept ID:	
Name:	
Attributes:	Start Time, Finish Time
Hierarchy Relations:	Has Sub-Process
Sequence Relations:	Ends Before Starts, Starts Before Ends, Starts Before Starts, Ends Before Ends
Other Relations:	Is Performed By, Has Location

Role

Concept ID:	
Name:	
Relations:	Has Members, Has Constraint, Has Group, Has Authority Link, Is Empowered, Has Goal, Has Agent, Requires Skill

Organisational Group

Concept ID:	
Name:	
Relations:	Has Role, Has Constraint, Has Group, Has Authority Link, Is Empowered, Has Goal, Has Location, Has Members

Agent

Concept ID:	
Name:	
Relations:	Has Role, Has Constraint, Has Group, Has Authority Link, Is Empowered, Has Goal, Has Location, Has Skill

Location

Concept ID:	
Name:	

Table 3: Knowledge Template for Role Definition

The SPEDE tools and methodology are being put through continuing assessment in large scale industrial enterprises in the UK, such as Rover and Rolls Royce Aerospace. This is aimed at improving the quality and design of the toolset and accompanying methodology, but also at populating the toolset with real-life complex examples.

Considerable work has also been done within the SPEDE project on developing techniques and methods for acquiring process knowledge in the form of structured natural language descriptions. A grammar and natural language process ontology have been constructed to assist in this. Further work is required to provide tool support for these descriptions and to establish mappings to the ontologies discussed in this paper.

SPEDE reflects a growing trend within the field of knowledge engineering; that the lessons learnt about the acquisition, organisation and use of knowledge within the area of knowledge based systems are widely applicable elsewhere.

SPEDE facilitates the application of knowledge acquisition tools, techniques and methodology to the activity of acquiring and managing process knowledge within the context of BPR. A principal mechanism for delivering knowledge re-use and enabling knowledge organisation within SPEDE is the use of ontologies. There are great advantages to using ontologies in this context:

- They provide a means of generally facilitating knowledge re-use. In particular this is possible within groups and organisations, but may also be viable across organisations.
- They provide links from common clearly defined terms to process models gathered from a potentially wide number of sources across the organisation. This has great benefits for enabling the meaningful interpretation and analysis of the models.
- They can be used to express knowledge requirement templates that guide KA according to the activities carried out during BPR .

In summary, SPEDE views BPR as a knowledge based design activity. It provides users with the tools, techniques and methodology to support the acquisition, organisation, management and re-use of the knowledge associated with performing BPR. In so doing it greatly reduces the costs and risks of doing BPR.

Acknowledgements

The work described was carried out under the EPSRC Innovative Manufacturing Initiative (IMI). The SPEDE project is being jointly undertaken by the Artificial Intelligence Group at the University of Nottingham, the Warwick Manufacturing Group, the Keyworth Institute at the University of Leeds, Rolls Royce, Rover and Computervision.

References

Bach, V., Brecht, L., Hess, T. & Osterle, H. 1996. Enabling Systematic Business Change: Integrated methods and software tools for Business Process Redesign. Wiesbaden: Verlag Vieweg.

Borst, W.N., Akkermans, J.M., and Top, J.L. 1997. Engineering Ontologies. *International Journal of Human Computer Studies* 46, 365-406. Special Issue on Using Explicit Ontologies in KBS Development

Burke, G. & Peppard, J. 1995. Business Process Re-engineering: Research Directions. In G. Burke & J. Peppard (eds), *Examining Business Process Re-engineering: current perspectives and research directions.* London: Kogan Pages 25-37.

Gruber, T.R. 1992. ONTOLINGUA: A Mechanism to Support Portable Ontologies, KSL-91-66, Knowledge Systems Laboratory, Stanford University.

Gruber, T.R. 1993. Towards Principles for the Design of Ontologies Used for Knowledge Sharing, KSL-93-04, Knowledge Systems Laboratory, Stanford University.

Gruninger, M. and Fox,M.S. 1995. The Logic of Enterprise Modelling. In J. Brown and D.O.Sullivan (eds) *Reengineering the Enterprise* pp. 83-98. Chapman and Hall.

Lenat, D. B. & Guha, R. V. 1990. *Building Large Knowledge-Based Systems.* Reading, MA: Addison-Wesley.

Noy, F. N. & Hafner, C. D. 1997. The State of the Art in Ontology Design, *AI Magazine,* Fall 1997, 18-3.

Shadbolt, N., Motta, E., & Rouge, A. 1993. Constructing knowledge based systems. *IEEE Software*, November, 34-38.

Talwar, R. 1994. Re-engineering - a Wonder Drug for the 90s? In C. Coulson-Thomas (ed), *Business Process Re-engineering: Myth and Reality.* London: Kogan Page.