PLANET Workflow Management R&D RoadMap

Members of the PLANET Technical Coordinatation Unit
on Workflow Management[1]

25th February 2003

[1]People that have contributed so far are: Paul Kearney, Daniel Borrajo, Amedeo Cesta, Nicola Matino, and Nikolay Mehandjiev

Contents

Chapter 1 Introduction

1.1 Preface

This is the second version of an R&D Road Map for AI Planning and Scheduling (AI P&S) applied to business process management (BPM) produced by the Workflow Management Technical Coordination Unit (TCU) of PLANET. PLANET is the European Network of Excellence on AI P&S. The purpose of a Road Map is to coordinate R&D by establishing end-user requirements on short medium and long time scales and proposing research and technology transfer goals and activities that will enable the requirements to be satisfied. The current version is only a first step towards such a Road Map, which in any case should be a living document updated regularly. BPM and AI P&S are two disciplines with many parallels, but which have largely been pursued by disjoint communities. A necessary precursor to producing a Road Map is to align the two disciplines so that specialists in each can understand each other. One of the main achievements to date has been to develop an understanding of how the "world view", vocabulary, challenges, etc. of Business Process / Workflow Management relate to AI Planning and Scheduling. This has been possible because of the active participation of a number of workflow and process management experts from end-user organisations and consultancy companies.

1.1.1 TCU Presentation

Currently, there is a growing interest in the application of Artificial Intelligence (AI) Planning and Scheduling (P&S) techniques to real world problems. We have recently seen impressive applications of P&S in space, robot, elevator control, military missions planning, etc. However, there are still many open tasks that can be (semi-)automated using AI P&S technology. One of such tasks is the production, and execution of models of organisations (workflow management). A popular way to model how organisations work is to focus on their internal processes (the ways they do business and the activities and business rules they follow). However, this is not a simple task; organisations operating in the current economy, especially those doing business via the Internet, have processes which are constantly changing in response to the needs of their customers and the business environment. This raises two distinct requirements. Firstly, organisations need powerful tools to automatically model, simulate and optimise their processes in such a way that the generated models comply with business rules. Secondly, organisations need computational tools, usually called workflow tools, for executing (enacting), monitoring, and dynamically adapting to changes those processes. So, the output of the first type of tools, the process models, become the input to the second type of tools (see Figure 1.1). From the point of view of the organisation, it can be seen as a set of agents (human and machine resources) executing processes (in parallel or in sequence) where each process is composed of a set of activities linked by constraints and dependencies.

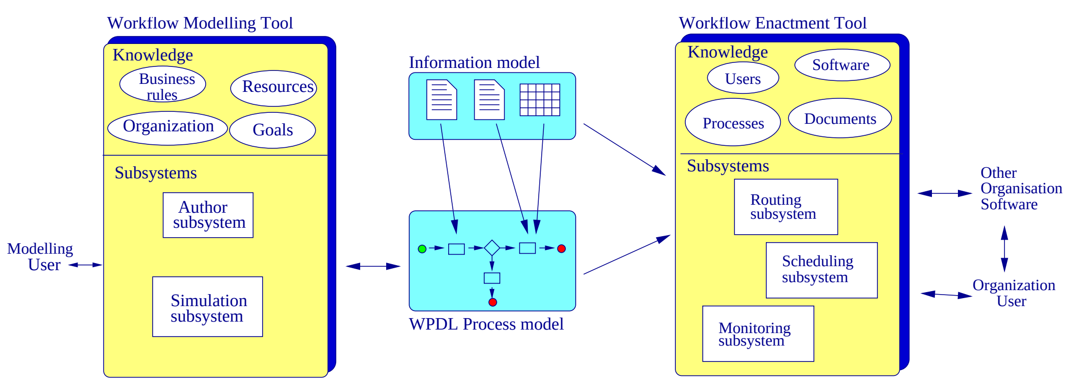

Figure 1.1: High level view of an architecture for workflow applications.

From an AI perspective, P&S tools can be effectively applied for the first type of tools, since they provide a declarative representation of the knowledge within the activities of processes, as well as means to generate only valid process models. This is crucial for people to understand how the organisation are really behaving. Also, P&S tools are able to obtain good process models according to one or several criteria (time to enact process, user satisfaction of the process, cost of the process, ...). In relation to the second type of required tool (enacting processes), AI P&S allow for efficient monitoring (through the explicit reasoning about predicted situations after applying plans or parts of plans), as well as re-planning when problems arise during enactment of processes.

Workflow management is a field in which few applications exist of these techniques. Workflow management has two interesting properties: a potential big impact in modern organisations (specially in the e-business context); and a close connection of the way AI P&S techniques work and how problem solving occurs in workflow management. The current state of the art in workflow shows that very few commercial tools incorporate AI P&S techniques into them. Current Workflow Management systems (WfMS) essentially automate the routing of documents between workers or teams according to pre-defined processes definitions. At the same time, they also handle the sets of tasks to be performed by the workers. WfMS and AI P&S are two disciplines with many parallels, which have largely been pursued by disjoint communities.

The purpose of the work of this TCU is to bring together researchers, practitioners, and software vendors. During phase I of research network, this TCU produced a Road map in which the commonalities were identified and described in some detail. The main conclusion is that they deal with common problems, so that AI P&S can greatly help on the automation of processes within organisations. This is an updated version of that document.

Therefore, the main purpose of the TCU on Workflow Management is to promote the effective application of AI Planning and Scheduling (AI P&S) techniques to Workflow Management.

1.1.2 Workflow and its role in business process management

A business process is the chain of activities involved in delivering a product or service to a customer (within or outside the organisation). Designing business processes is a knowledge-intensive human activity supported by software modelling and simulation tools, and is closely tied in with matters such as business policy and enterprise organisation and culture. An instance of a business process created, for example, to deliver a particular service to a particular customer is analogous to a plan in AI. In BPM terminology, however, a plan also includes allocation of resources (e.g. workers) and target start and completion times. In terms of AI, this would be the equivalent of generating a plan with resource and temporal information, that is, the integration of planning and scheduling techniques [Drabble, 1999].

In some application domains, for example military logistics, generating a plan and instantiating it with appropriate resources and time windows is complicated, and AI planning techniques are being applied successfully in such areas [Tate *et al.*, 2000]. However, following the way in which business processes are handled currently, planning only involves selecting from a set of pre-defined templates. The main technical challenges in this setup arise because an organisation is a distributed system that executes many process instances concurrently in an uncertain environment. Furthermore, failures and other exceptions occur frequently, and re-planning must be integrated with execution. In the next future, automated planning tools should not only instantiate processes templates, but also be able to generate dynamically the executable processes templates.

A workflow management system (WfMS) automates the coordination of activities and transfer of documents within a business process[Georgakopoulos *et al.*, 1995]. It delivers the work to the "in-tray" of the appropriate software component or human worker or team according to pre-defined rules (a process or workflow definition). Current WfMS do not (generally) perform planning, scheduling or resource allocation. Any such considerations must be built in to the process definition or else handled by the productive resources owning the in-trays. Specifying this low level process or workflow definition is again primarily a human design activity performed with the assistance of software tools (often specific to the WfMS).

1.1.3 Requirements

Within this document, requirements have been classified as short, medium and long term as follows:

short term address short-comings in current generation process management software. The most important items in this category are: integration of temporal reasoning and resource allocation/management algorithms into workflow management software; and incorporation of a planning capability to enable a WfMS to modify the process instance automatically during execution, to cope with failure, changed objectives, and other exceptions.

medium term current generation workflow software handles high volume routine processes, typically involving low-skill workers. The medium term requirements concern extending this support to high-skill knowledge workers. This may involve, for example, building process knowledge awareness into software tools.

long term more radical (e.g. adaptive self-organising) approaches addressing the need for organisations to function in a business environment that is increasingly uncertain and subject to change.

These requirements describe research and development goals that could be fulfilled with integrated projects within FP6.

1.1.4 Conclusions and recommendations

AI planning and scheduling and business process management are complementary disciplines with much to gain from collaboration. The ability to invoke AI components flexibly and dynamically from within a workflow framework would considerably enhance business productivity and give the European software industry a competitive advantage. The Workflow Management TCU has a valuable continuing role to play in bringing together researchers, software developers and end-users from the two communities and promoting joint work between them. The active participation of workflow and process management experts from end-user organisations and consultancy companies in discussion with planning researchers has enabled considerable progress to be made on the R&D Road Map. The TCU must make every effort to involve more end-user representatives from a spectrum of industries. A number of commercial software vendors are registered on the TCU mailing list but have not as yet participated actively. It is important bring such organisations fully into the fold. The concluding section of the report makes some proposals for future activities.

1.1.5 Further reading and links

Perhaps, the current best text for an overview of the relation between AI P&S and workflow management is the paper by Myers and Berry [Myers and Berry, 1999]. With respect to the field of workflow management and business process re-engineering, some overviews can be found in [Alonso *et al.*, 1997, Diimitrios Georgakopoulos, 1995, Jacobson *et al.*, 1995, Mowshowitz, 1994]. In the Web page of the TCU,[1] one can find more links to Workflow Management related Web pages and documents.

[1] http://scalab.uc3m.es/∼dborrajo/planet/wm-tcu/

1.2 Introduction

The two communities, Workflow Management and AI, are separated by a common language, with terms such as planning being used in both with different meanings. A necessary precursor to producing a Road Map is to align the two disciplines so that specialists in each can understand each other. One of the major successes of the TCU to date has been to bring AI P&S researchers into contact with BPM specialists and so further mutual understanding. This progress is documented in the next chapter, which gives and overview of BPM and how it aligns with AI P&S. This is followed by an outline of requirements on short, medium and long time scales as perceived by BPM end-users.

The remainder of the document expands on these requirements and looks at ways in which they could be met by existing and future results from AI P&S and related disciplines. The discussion is divided into a number of themes (see Figure 1.2).

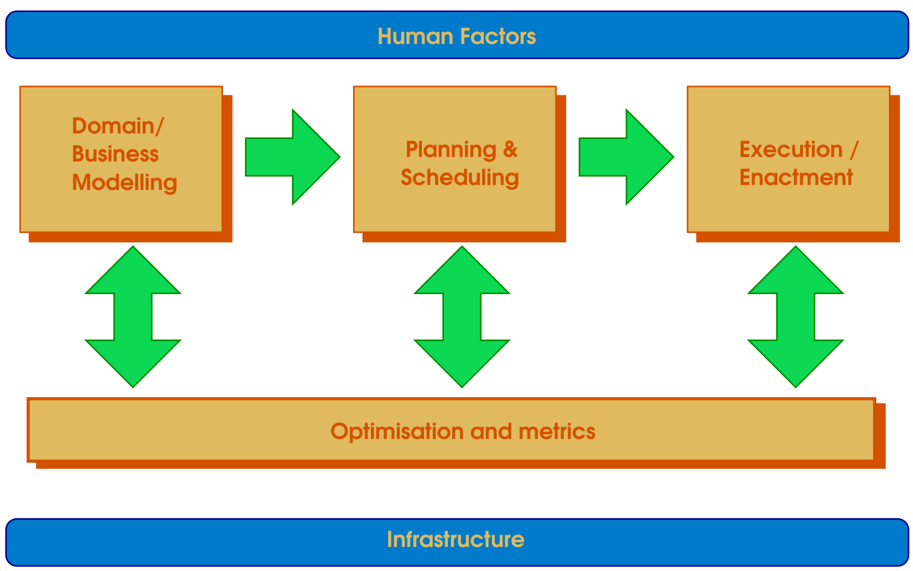

Figure 1.2: Themes that are discussed in this document.

Each theme chapter includes sections on the state of the art (including trends and current projects), research goals and open issues, and recommended actions. The first theme deals with human issues. A feature of business processes is that much of actual work is performed by people. There is a tendency in BPM to pursue automation and to treat human actors in the process as if they were machines. This is often counter-productive resulting in de-motivation and a failure to utilise human qualities. This chapter explores these issues and examines how they might be addressed. The following chapter looks at software infrastructure. It is important to appreciate that for AI P&S techniques to be applied in practice they need to be integrated with / interfaced to commercial software packages. This chapter looks at issues such as reference architectures and interface standards. A common understanding of architecture would also facilitate collaborative research and demonstrations. Chapters covering life-cycle oriented technical themes then follow:

- business / process modelling and knowledge engineering: generating a computer usable representation of processes is the first task to be solved in both fields. Here, common problems arise such as how much the process representation corresponds to reality, or what language to use to represent the model.
- planning, scheduling and resourcing: both fields require the generation of sequences of activities to be executed (enacted). These activities need also information with respect to resources, as well as time frames implied in their execution.
- enactment/execution and monitoring: less studied in the field of AI P&S and with more software tools available in the field of BPM, the execution (enactment) of the plans is a key component of the cycle. Specially important in the enactment of business process are failures (or predicting them) and how to handle them.
- adaptation, optimisation and metrics: usually, processes (or plans) have been generated without optimisation goals in mind. Organisations in current very restricted and competitive markets need gradually more emphasis on metrics and finding better processes. Also, a related aspect is how to dynamically change the processes to adapt to those markets.

Finally, there is a chapter summarising conclusions and recommendations for future actions.

1.3 Overview chapter

This chapter gives an overview of the “standard model” of process management and attempts to position AI P&S within this context.

1.3.1 Process management and workflow

A core business process is the end-end chain of activities involved in delivering a product or service to a customer (they can be external or internal to the organisation). "End-end" means that a business process starts with an initial contact with the customer and runs through to completion of the contract. In fact, since the customer's satisfaction with one service influences requirements for future services, a business process is best seen as a closed loop. In addition to core business processes, there are management processes (including processes concerned with designing the core processes) and support processes that facilitate the other types of processes. The set of business processes for an organisation comprises the organisation's working practices. Organisations differ in how explicitly the processes are defined, and in the form they are represented. In some cases the processes are implicit, in others they are recorded in textual codes of practice, in others they are documented in (semi-) formal representations and/or software modelling tools. A set of business processes is highly analogous to a set of stored plan templates or a hierarchical task network (HTN) and could readily be represented in this way (this will be explained further later). Business process management can be presented as having the following aspects (as also shown in Figure 1.3):

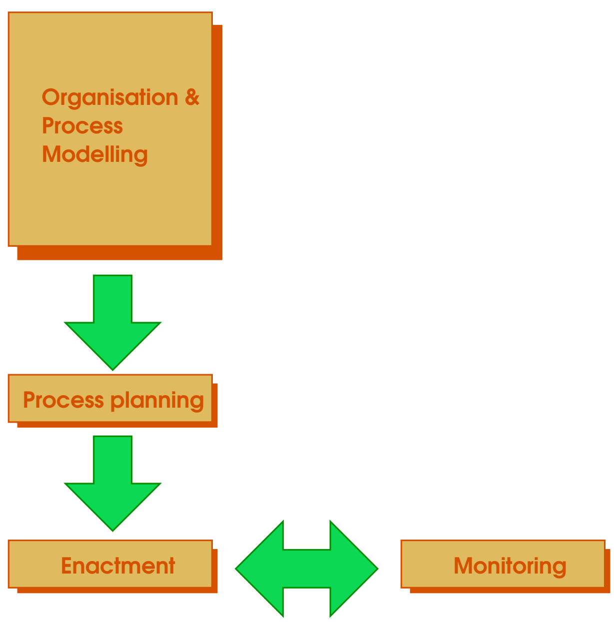

Figure 1.3: Business process management steps.

Process modelling This involves designing, modelling, evaluating (simulating), modifying, optimising, etc. the organisation's processes . For each basic product or service the organisation offers to its customers, the activities involved, the relationships between them, resource requirements, etc. must be defined. It is basically a human activity, though supported by computer-based tools to record and display the process model, run simulations, etc. Design decisions are made based on experience and analogy to previous designs. Choices are tied closely to other aspects of the enterprise and business environment such as: the nature of the business, business goals, organisations standards or norms, organisational structure (of the enterprise), enterprise culture, legacy infrastructure, etc. Although process design is often presented as happening top-down, the practical constraints imposed by the current state of the enterprise mean that there is a strong bottom-up behaviour. Design of the processes and activities typically go hand-in-hand, so that although the analogy between process and AI plan is strong, the analogy between the activities of process design and classical AI planning is much weaker.

Process planning (elaboration, resourcing and scheduling) A process definition is basically a template. This phase involves identifying the appropriate template to use, elaborating and filling in an instance of that template in sufficient detail for it to be executed. The first step is normally to gather information from the customer on the product or service required. This allows the tree of possible processes to be pruned considerably, but a number of alternative branches may still remain. The next step is to produce a schedule based on a target end date required by the customer, dependencies between tasks, and knowledge about how long tasks take (e.g. typical and minimum times). If it is not possible to achieve the target end date, negotiation with the customer takes place. Then, the people and other resources required for each task are identified and "reserved" for the appropriate time slots ("resourcing" or "provisioning").

The resourcing and scheduling problems are coupled by virtue of finite capacity and/or non-sharable resources. If the required resources are not available, then the earlier steps must be revisited. In process management the result is referred to as the plan, and the process of producing it as planning, which is a source of confusion as the usage is different from that in AI. Note that further detail will often be decided at execution time, and the balance between design time and execution time decisions varies considerably. Again, these activities are often performed by people assisted by relatively dumb software tools. The nature of the tools and the form of the output depend on context. For example, MS Project (or the equivalent from another supplier) could be used to create a "production plan" to be carried out by a human organisation. Alternatively a proprietary tool could be used to generate a process description for enactment by a workflow engine. In each case, the tool and the representation is often different from those used earlier in the modelling phase.

Enactment The production plan is carried out, with detail being elaborated during enactment. The boundary between planning and enactment is context dependent. Process planning is essentially the first part of enactment of a core process. Furthermore, at the start of enactment, the plan may still contain alternative branches that are pruned as information is gathered and decisions made during enactment. Almost always, execution is distributed, with different production resources, computer programs, or people carrying out the constituent activities. The activities have to be coordinated to ensure correct sequencing and that compatible variants of the activities are performed. Coordination takes place via mechanisms such as: events, transfer of documents, existence checks on documents, etc.

A workflow management system uses information contained in a low-level process plan definition to route work items to the appropriate production resource and provide the necessary coordination signals. Note that workflow systems (generally) do not plan work, and workflow also assumes resources will be available. Production resources will be involved in enacting multiple processes and instances of the same process in a time-sharing manner. A production resource (or rather a component encapsulating one or more resources) sees the processes in which it participates as a queue of work items (or tasks) waiting to be acted upon. Depending on how the system is organised it may simply work on the next task whose preconditions are satisfied, or it may have rules for prioritising tasks. Either way, different processes can interfere with each other due to the finite capacity of a shared production resource.

Monitoring As execution proceeds, information on progress (e.g. notification of completion of tasks, delays and other problems) is fed up to a management function. This compares actual progress with the production plan. Minor differences between the plan and actual progress may simply require updating of the plan (for example with slightly different commencement times for tasks). These changes need to be propagated to the resources executing the plan. More significant differences may require the planned activities to be altered during execution. This may include some back-tracking, for example to remove some item of equipment that was installed following the earlier plan, but is now no longer required. More drastic problems may require all the effects of the plan to be undone and a new plan created. The monitoring function may try to anticipate future problems and modify the plan in advance to avoid the problems. This is sometimes known as jeopardy management.

1.3.2 AI planning and scheduling

AI planning and scheduling (AIP&S) is concerned with determining a sequence of actions that when executed by one or more agents with the world in some initial state satisfying given conditions, results in world state satisfying given goal conditions. A process is a description of a set of activities. A plan is a description of activity for a given objective; it is an instantiated process. AI planning provides many different techniques to generate plans, but there are two main ways of specifying the domain. On one hand, in STRIPS style planning, the operators consist of individual activities [Fikes and Nilsson, 1971]. A planner combines instantiation of these for a given objective to form a plan. On the other hand, HTN planning domain descriptions are essentially process descriptions. They let you specify parameterised descriptions of processes that can be automatically assembled and instantiated to form a plan for a given objective. Classically, this overall problem is divided into a number of stages:

- Modelling (or knowledge engineering): this concerns finding the right way to represent the world and the problem so that planning and scheduling may be performed. Classically, this representation consists of a definition of some space of states that the world and its constituents may be in, and a set of primitive operators that can be applied to cause (constituents of) the world to change states.

- Planning: this concerns finding one or more sequences of actions that should cause the world to change from the initial state to a state satisfying the goal conditions. The ordering of these sequences is not necessarily completely determined. Planning is concerned with logical dependency of actions in the sequence, e.g. that if action A is necessary to bring about the pre-conditions for action B, then A is performed before B. Planning may be performed bottom up by chaining together actions until the gap between initial and final states is spanned (e.g. STRIPS model). Alternatively, it may be performed top-down by recursively refining generic plans until they are expressed entirely in terms of executable actions (e.g. HTN model).

- Scheduling: a plan expresses the orderings of actions that should be able to bring about the goal. Scheduling determines which of the orderings of actions consistent with the plan will actually be used and on what time frame each activity will be executed. Often, this choice is based on some form of efficiency measure, for example overall time taken to execute (makespan). Also, scheduling handles the assignment of resources to individual actions (activities) so that resources are not over-allocated.

  Traditionally, planning and scheduling have been separated fields. However, recently there is a strong interest on performing them in an integrated way. Either by following a scheduling step after a planning step, by providing the adequate interfaces, or by developing integrated tools [Ghallab and Laruelle, 1994, Muscettola, 1994], a growing number of researchers are focusing on this issue, as it can be seen from the latest workshops and conferences papers [Drabble *et al.*, 2002, Ghallab *et al.*, 2002, Cesta, 2001].

- Execution: as such has not been a major concern of AIP&S except for some domains, such as robotics, or space missions. However, particular branches of AIP&S are concerned with execution-time issues. For example, it is recognised that an action does not always achieve its intended result. Thus monitoring must take place to compare anticipated events with actual ones, and if deviation is significant plan repair is initiated. If deviation is excessive or repair impossible, then the plan is abandoned and a new plan generated. At the extreme end of the scale are so-called reactive planners in which planning and scheduling take place at execution time, with planning, scheduling and monitoring actions interleaved with the goal-achieving actions. There are also hybrid approaches that lie between the two extremes and allow to efficiently go from one extreme to the other.

1.3.3 A comparison

It is clear from the above descriptions that process management and AIP&S address similar issues, and there are many parallels between the two disciplines. Figure 1.4 compares the two at a coarse level, aligning phases that are roughly equivalent. However, there is no direct equivalent to AI planning on the process management side - although a process (model) is approximately equivalent to an AI plan, it is generated by people supported by software drawing and modelling tools rather than by an analogue of AI planning. Usually, in the case of BPM, processes have very few knowledge describing each activity. They usually have information on issues such as who is responsible of the activity, or the time and cost of the activity. Very rarely, one has to provide information on their pre- and post-conditions, as it is the case for AI P&S. However, from a knowledge-rich perspective of an organisation,[2] those conditions should be specified, so that reasoning about itself can take place as was studied during the SHAMASH EU-funded project [Aler *et al.*, 2002a].

Also, both business processes and AI plans can occur on multiple levels. Thus process management can itself be seen as an enactment of a meta process; enactment of strategic processes may involve definition of tactical processes and so on. Similarly execution of strategic AI plans may involve planning at a

[2]There is clearly now a trend towards an explicit and declarative representation of organisation knowledge through knowledge management, competencies modelling, ontologies, etc.

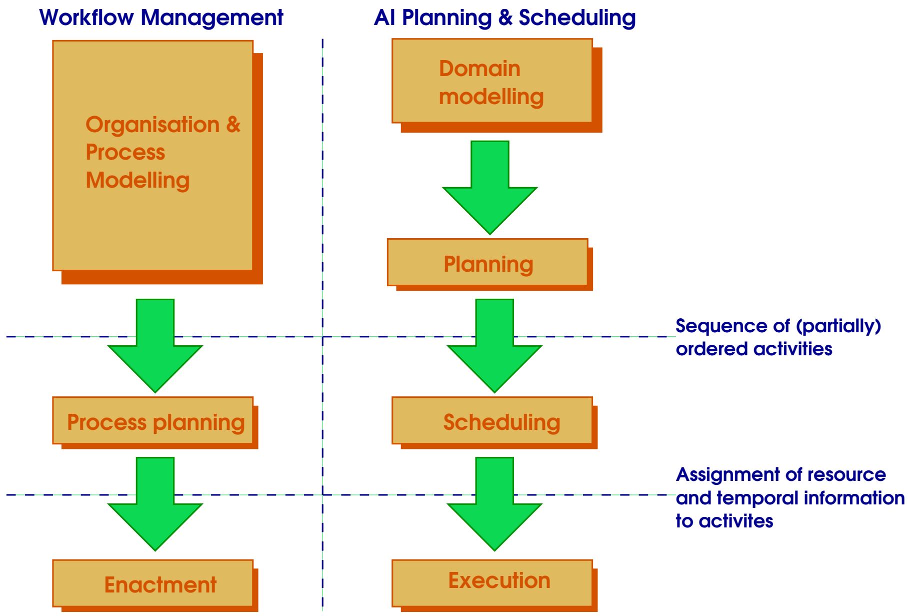

Figure 1.4: A comparison of process management and AIP&S.

tactical level. In consequence, it is important to consider applications of AIP&S within the activities taking place during enactment a business process. For example, an early step in a process may involve detailed planning or scheduling of activities occurring later in the process. Despite the similarities, there are also significant differences:

- Terminology - the word "plan" itself has a different meaning in the two disciplines as it has been specified before;
- Most of the design-time (as opposed to execution-time) activities in process management are performed by people assisted by relatively simple software tools. In contrast, the emphasis in AIP&S is on producing intelligent software that can perform planning and scheduling largely automatically, with occasional assistance from a person;
- AIP&S representations tend to be mathematically formal and semantically precise, though this often means they are difficult (for a non-expert) to understand. In process management, the opposite is true: the representations are domain-oriented and easy to understand, though the semantics are often somewhat vague.
- Languages for defining processes as input to workflow engines are basically scripting languages for coordination of activities and are at a lower level than AI plan languages.

- Classical AI planning techniques focus on difficult combinatorial problems - many different combinations of operators and states are possible, only a few of which constitute viable plans. In process management, activities are fairly specific to processes, and there is much less scope for combining them in different ways to form different processes. These differences present opportunities for synergy as well as barriers to be overcome.

1.4 Requirements

This section deals with the requirements that the workflow management world imposes on AI P&S in order to have an impact on a short, medium, or long term.

1.4.1 Current state of the art in workflow and process management

The idea of process management is still fairly new. In the past, organisation processes were implicit in each organisation structure and culture. Departmental procedures and practices would be known within the department, but no individual had a clear end-end view of a process. A similar statement could be made about the software systems that support the enterprise operations. These were, and still are, often large monolithic applications in which the business processes are implicit. Consequently they are difficult to change and tend to tie the organisation into the processes encoded in the software. However, the importance of the day to day operation of the organisation and the expense and disruption involved in replacing them mean that many of these so-called legacy systems are still in active use. The current trend in both organisations and their operational support software is to represent the business processes in an explicit and distinct manner. As a result, it is easier to study how to improve a process and also easier to implement the improvement. In the case of the software, the need for modifiability and software re-use has led to a component-based philosophy. Instead of monolithic applications, functionality is encapsulated in re-usable modules that can be combined in different ways to construct new "virtual" applications rapidly. One way to view workflow management systems is as the architectural glue that links the components together to form the application. At least in theory, the process definition can be changed independently of the components, and functionally equivalent components substituted without changing the process definition. Often these components do not replace the legacy applications. Rather the components use them as servers in providing their functionality.

Of course, much of the work in a business process is performed by people. A workflow management system treats people in much the same way as the computation components. Typically, an interface is provided that presents the user with an in-tray and out-tray of work items. This interface encapsulates the user in a similar manner to that in which the component interface encapsulates the software functionality. This approach is suitable for partial automation of well-understood routine processes.

Sometimes groupware software systems (such as Lotus Notes and Microsoft Exchange) are described as workflow systems. These systems primarily provide a messaging and information sharing environment that can be used by participants in business processes. However facilities such as document routing scripts and forms can be used to define workflows to some degree. Industry is currently in transition from the old-style monolithic support applications and paper based office processes to workflow-based systems. Legacy applications certainly will not disappear overnight. Rather, components and workflow systems will gradually diminish their role. The legacy software may never disappear entirely, however, especially where the products, services and associated processes are relatively mature. AI scheduling (and to a much lesser extent planning) techniques have certainly been used in special purpose business support applications. However there has been little or no influence by AIP&S on process management as a discipline or on the methods and tools through which it is applied. Similarities (for example between plan and process description languages) are due more to convergent evolution than to direct influence.

1.4.2 Enumeration of requirements

There now follow descriptions of areas in which current business process management is recognised to be deficient. The list is not exhaustive, and we invite proposals for additions to the list. The requirements are into short, medium and long time scale categories. The short term requirements concern ways in which current practice and tools can be improved. The medium term requirements concern extension of workflow-related support into classes of processes and users that are not catered for by current workflow systems. The long term requirements concern the need for a more radical re-think of how organisations and their software support infrastructure are organised.

Short term

The following are seen as short-comings in current-generation process management software. They are presented in approximate order of importance, though the first two are of comparable ranking.

- Integration of scheduling and resource allocation/management algorithms into workflow management software. Current workflow management software automates the flow of work items between work queues according to pre-determined rules. It does not deal with allocation of resources to tasks or take resource availability into account in prioritising or scheduling the work.

- Re-planning. There is a requirement for incorporating an ability to modify the process instance automatically during execution, to cope with failure, changed objectives, and other exceptions. This could be done by altering the process instance plan being executed (inserting and deleting steps) or by creating and executing an ancillary plan (conditional plan) containing the additional process steps.

- Generation of workflow definitions from high-level process models. Process modelling tools work with relatively high level process definitions, whereas workflow management systems require low level definitions. Current generation tools do not do a good job of bridging this gap. Tools are required that automatically generate low-level definitions that can be input directly to workflow management systems. The ability to do this in reverse is also desirable.

- The ability to feed data captured in the workflow engine back into the modelling and simulation tool to improve modelling at that level. Workflow engines capture a great deal of data in the course of enacting process instances. This contains useful information latent within it, but it is rare that data mining techniques are used on it.

Medium term

- Process support for intermediate level and knowledge workers: current workflow and groupware systems "pick the low-hanging fruit", that is they automate that which is easy to automate - enactment of routine processes and providing information-sharing and communication services. The tasks performed within these processes are routine also, and are performed by relatively low skill workers. There is very little process-related support available for high skill (professional) knowledge workers or workers at intermediate skill levels . “Process aware” and “knowledge aware” support to enhance the effectiveness of intermediate and high-value knowledge workers is required, but much more difficult to achieve.

- Empowerment of users: current workflow management systems are suitable for routine processes and demand uniformity from users, effectively expecting them to behave like machines. This makes poor use of human abilities, even in the case of low-skill workers, and can have a demotivating effect. Informing people about the context of their work is a necessary short-term requirement, which should eventually lead to systems that encourage and support initiative, but would require workflow systems to be enriched with a semantic knowledge on the processes they enact. Future systems need to assist people in achieving their potential in their roles, which means encouraging initiative and adapting to human diversity rather than enforcing regimentation.

- Visual representations of the current status of a process instance, so that workers within a process can see how their activities fit into the ”big picture” [Zuboff, 1988]. This is actually intermediate between the “short term” requirements, which might be satisfied by incremental additions to current generation workflow, and the “medium term” ones which require a change in philosophy.

Long term

The following are factors driving process management development in the long term. Mostly they concern the need for organisations to function in a business environment that is increasingly uncertain and subject to change.

- Flexibility: one of the main drivers in process management is the need to be able to get new products and services to market quickly. This means that an organisation and its supporting infrastructure must be capable of enacting a wide variety of processes, with the actual set of processes active at a given time being easily changed.

- Evolvability: no matter how flexible an organisation is in the short term it will have to change in the longer term in response to changing markets, technology, etc. Change takes time, however, and the organisation must continue to operate. The organisation needs to be capable of gradual evolutionary change to avoid the current problems with legacy systems recurring in the future.

- Adaptiveness: currently, organisations and their business processes are seen as basically static, but subject to occasional discrete changes such as re-organisation or introduction of a new product and/or process. However, the frequency of change is increasing. In the future organisational models will be needed in which continuous change is the normal state of affairs. Such models most incorporate processes that "sense" the drivers for change (e.g. increasing demand for a product) and cause appropriate changes to the organisation model. Organisation software infrastructure will need to support such dynamic organisational models.

- Decentralised management: a management paradigm shift is currently under way motivated by the need for flexibility, evolvability and adaptiveness. This is variously described as a move from centralised to decentralised management, from management push to market pull, and from plan and build to sense and respond. This involves moving decision-making responsibility from central management to autonomous local units. The behaviour of the organisation as a whole is then the cumulative result of local decisions. The role of higher management is then one of defining performance metrics and other incentives by which local managers make decisions, and also of providing means by which the autonomous units can interact constructively. Workflow management systems are very much tied into a plan and build management style. A new approach to software infrastructure is required to support decentralised management. An agent-based approach seems well suited in this respect. Use of an agent-based approach does not in itself guarantee the benefits sought from decentralisation, however. A better understanding of how to apply agents and agent-based approaches to achieve the benefits is still required.

Dynamically changing organisations of the future will involve forming opportunistic organisational structures and dynamic supply chains. Theoretical work in Virtual Organisations and "switching" is related to planning approaches and can be used to build the workflow systems of the future [Mowshowitz, 2001].

1.4.3 Recommended actions

This section highlights a set of actions that will be able to attain some of the objectives (requirements) that were defined in the previous section.

- The potential for benefit from applying existing AI P&S techniques to short term requirements should be explored through case studies of large-scale real problems conducted jointly with domain experts and process management experts. It is apparent that process management problems can be posed as planning and scheduling problems. However, for the techniques to be adopted, it must be shown that they result in a significant benefit compared to current practices in the context of realistic business processes. The techniques must also integrate with other components of organisation software, and be usable by the typical software or business process engineer. PLANET cannot perform such case studies, but it should facilitate and encourage them. It should also publicise the results.

- Generally, links with process management and software engineering communities need to be strengthened by means of interdisciplinary events and other measures. Within the second phase of PLANET, PLANET II, a set of workshops have been organised in conjunction with people from WfMC and other relevant communities in order to further understand the gap between the two fields.

- To encourage case studies and increase awareness of process management challenges within the planning community, PLANET should collect examples of process management problem domains.

Chapter 2 Road Map themes

The rest of the document looks more closely at state of the art (including trends and current projects), research goals and open issues, and recommended actions within the themes shown in Figure 1.2. The boxes in the middle correspond to the main divisions shown in Figure 1.4. Some further explanation is needed where the two halves of Figure 1.4 do not align, however. Basically, activities in which people design or elicit models (possibly supported by software tools) are included in the modelling theme. Planning (as understood in AIP&S) is included with scheduling and resourcing.[1] Planning and scheduling can be performed as part of the enactment of processes as well as off-line. The optimisation and metrics theme recognises the need to measure attributes of the various models and provide feedback to improve desirable qualities. The feedback can be to the same box, e.g. measurements of the attributes of a process model can be used to optimise the process model. Larger-scale optimisation is also important, however. For example, measurements during execution of a process can be used both to adjust the theoretical production plan to better reflect reality, and to provide information to help improve it. These improvements can be fed back into the executing process, and so on.

The remaining two boxes represent "orthogonal issues": infrastructure and human factors. An important infrastructure issue is the establishment of a reference architecture for a highly-modular "AI-enabled" process management system, covering both design-time (off-line) and execution-time systems. The human factors box reflects the need to take into account (throughout process management) of the special characteristics of the people involved in performing the processes, and to form a symbiotic relationship between "man and machine". These two "orthogonal themes" are covered first.

[1] We will use indistinctly the terms "scheduling" and "scheduling and resourcing". Therefore, we will consider that scheduling techniques can also handle resources, and not only time information.

2.1 Human issues

In this section, we will focus on the state of the art problems, requirements, and future actions that relate to the fact that humans are in the loop with respect to workflow applications. This will condition how AI P&S techniques will/can be applied to this domain. Perhaps, the biggest difference between the type of domain requirements from the area of Workflow Management and those from the Intelligent Manufacturing area is precisely related to human issues (see report on the common workshop organised by the two related PLANET TCUs in [Aylett and Borrajo, 2002]).

2.1.1 Introduction

In relation to workflow management, human issues can be mainly divided into three main categories depending on the role of the user:

- Users involved in process design. This issue will be covered in more detail in Section 2.3. Also, for a deeper understanding of the knowledge engineering aspect of developing AI P&S tools, the reader can consult the Road map of the PLANET TCU on Knowledge Engineering.

- Users involved in process management tasks. Once the process has been modeled, it can be enacted. There is a human role (one can think of a finer division of this role into several) for monitoring the execution of activities, adjusting the workflow tool, or drawing conclusions from the observed behaviour of the enacted process. This issue strongly relates to monitoring tools, which are auxiliary to the main task of the workflow tools (routing documents among people). We will not discuss here the issues that this type of users enforce into the requirements of the tools to be developed, but they will certainly be important in the future to consider, once workflow tools incorporate many of the features that are being described in this document.

- Users of the workflow tool. Given that processes are enacted by people from the organisations, it is crucial to study the impact of fielding a workflow system within the organisation. As it usually happens with computer applications being used in organisations, there has to be a consensus of the utility of using such tools. In this section, we focus on this aspect.

2.1.2 Current state of the art

Most business processes require people to perform at least some of the tasks. Workflow management systems tend to view people and software resources in the same way: as means of carrying out process steps. At the present state of the art, they can be applied only to well defined, routine processes, introducing even more regimentation into dull, boring jobs. All too often, software is seen as a means of decreasing costs through automation and standardisation rather than as a means of enhancing value and quality of he product/service and hence customer satisfaction. It is also used to monitor productivity, thereby increasing pressure on workers further. Unfortunately the quantities measured tend to be those that are easy to measure (number of calls handled in a call centre, number of studied proposals, etc) rather than true measures of value contributed.

AI planning too has traditionally been concerned with automation of processes. The goal has been to build intelligent machines, i.e. to enable machines to perform activities that currently only people/animals can do. Little attention has been paid to amplifying human abilities, though there is a body of work in the area of mixed-initiative planning [Allen *et al.*, 1995, Tate *et al.*, 1998, Veloso *et al.*, 1997, desJardins *et al.*, 1999].

Computer-supported cooperative work [Beaudouin-Lafon, 1998] was a very active field in the early nineties, though activity seems to have died down recently, re-focusing on work within virtual teams. A stream of work within CSCW focuses on Tailorable Workflow systems [Kahler *et al.*, 2000], where workers can modify run-time functionality of workflow via preferences. Industrial psychologists work in the area of job design, which focuses on maximising the motivational characteristics that people experience in their jobs [Oldham, 1996], and have formulated factors that improve job satisfaction, for example autonomy, variety and responsibility. Technology such as workflow has the potential to both simplify and enrich the nature of work. Control over various aspects of work such as timing and method is thus considered crucial in workflow-type environments, with higher control leading to better productivity and work attitudes [Jackson *et al.*, 1993].

2.1.3 Research issues

The main objective with respect to human issues would be the understanding of how to achieve a synergistic, symbiotic relationship between human workers, managers and software systems within business processes. In order to achieve this goal, some research issues arise, such as:

- How to involve users in controlling their coordination support and workflow planning systems? A decision to empower users so that they can control workflow systems requires appropriate interfaces and methods to make this user control possible. These interfaces and methods have to be carefully designed to take into account the expected variety of user backgrounds and programming skills. The academic fields of End User Development [Nardi, 1993] and Visual Programming [Burnett and McIntyre, 1995] have researched these issues in the general case of software programming and control. An interesting research direction would be to see how the general findings in these areas map onto the specific domain of workflow planning.

- How to take into account human issues when at the process modelling and definition stage of workflow planning systems? Two approaches that are used in mainstream software development are Participatory Design [Kuhn and Muller, 1993], which aims to involve user representatives in the design of new software, and input from industrial psychology such as job satisfaction factors.

- What is the optimal balance between users and software during the different stages of workflow planning and scheduling? Here again the work on mixed-initiative planning could help on finding automatic ways of balancing the control between the user and the software tool.

- If planning techniques are to be employed successfully, new means of visualisation and explanation need to be developed to reflect the combinatorial nature of AI planning. This is specially true when new very fast planners use search techniques that are difficult to explain to a human in case they have to collaborate at the search level.

2.1.4 Recommended actions

In order to carry on this research, a set of actions are recommended here:

- To run a set of trans-disciplinary workshops, which discuss the relationship between systems providing user-control of workflow and contributions from the areas of Participatory Design, Industrial Psychology, Visual Programming and End User Development. A report could summarise the workshop findings and possible new research directions can arise from it. A meeting as such would have the problem of first trying to use a common vocabulary, as we also suffer in the first meetings of this TCU.

- To develop a prototype demonstrator to test the feasibility of user control at different stages of the planning-driven workflow development and enactment. In order to do so, a prototype planning system would have to be built exhibiting different types of interaction, and then test with a set of people each setting.

2.2 Infrastructure

This section focuses on the issue of the languages, interfaces, standards, and software tools that have to be used in order to integrate the P&S systems within the information system of any organisation. The infrastructure can be divided into:

- an open, layered architecture and interface definitions that allow independently-developed modules to be combined in a flexible way; and
- the software modules that work within this architecture. PLANET's role should not be to favour one technology over another, or to enter the debate over free software, open-source software and/or proprietary software. Rather, it should aim to ensure that research, transfer and exploitation can be conducted effectively. Wherever possible, compliance with existing official and "de facto" standards and interoperability with existing solutions should be encouraged.

2.2.1 Current state of the art

The issue of the infrastructure related to a software tool to be used within an organisation has many different views, given that this type of software has to interact with most software of the organisation. Therefore, it would be very difficult to describe all possible software interfaces that would have to be studied for integration. Here, we only mention some of the standards that are available and would most probably be the closest ones to consider.

- The WfMC[2] has established a reference architecture based on five interfaces between workflow engines and other classes of associated software. Now is working on detailed definitions of these interfaces. Conformance to the WfMC standards by the major software developers is mixed, however. The interfaces they have defined are:

[2] http://www.wfmc.org

  - process model definition tools: defines the interface between modelling tools and workflow management tools. Basically, it uses WPDL (Workflow Process Definition Language), which is being re-designed towards XPDL (XML based). Since WPDL is stable, people willing to interface current planning and scheduling techniques or tools with workflow tools can already use this language for specifying the process models.
  - users: defines how to provide information to users of the workflow tool with respect to the work (tasks, activities) that they have to accomplish
  - automated systems: defines the interface with the other type of agent that can accomplish tasks, other software tools
  - other workflow systems: given that currently many organisations are trying to interface their processes with suppliers and clients, it is needed to define the interface among the workflow management tools of the other organisations
  - system administrator: any complex organisation should have a workflow management administrator that is in charge of monitoring, and controlling the overall performance of the system.

- The planning community on its side has defined a standard for domain specifications within planning tools, called PDDL [Fox and Long, 2002]. This is based on the work of a committee of planning experts. The definition of this standard has helped on carrying on a planning competition [McDermott, 2000, Bacchus, 2001]. A second standard in the planning community has been during some time ADL (Action Description Language) [Pednault, 1989].

- A three-four layer organisation architecture has now become standard practice in industry. The layers consist of user client software, back-end applications, and one or two layers of so-called middleware. More effort will have to be devoted to understanding what level(s) are affected by introducing P&S tools in the loop, and how they are interfaced with tools in other layers.

- There are various competing standards for middleware components including: CORBA (OMG), COM (Microsoft), java-based solutions (Sun Microsystems and others), etc. At some point it might be useful for AI P&S based tools to interface other systems using any of several of these components.

- There are also a number of emerging standards for interoperation of software agents from FIPA, DARPA (KQML/KIF), OMG, etc. Since the agent-based paradigm for programming of complex systems is gaining acceptance, it will be more important in the future to understand how to introduce P&S systems into agents, how several agents with planning and scheduling capabilities are able to share their (partial) plans/schedules, or collaborate to generate plans that are going to be executed by several agents. There is already some preliminary work done with this respect [Paolucci *et al.*, 1999, Camacho *et al.*, 2001, Dix *et al.*, 2001].

2.2.2 Requirements

A reference architecture and interface (language/api) covering modelling, build-time and execution is required to enable highly modular approach to research and application systems. A modular approach is necessary to facilitate:

- re-use of software and avoid wasteful duplication of effort;
- synergy between the work of research groups developing complementary technologies;
- exploitation of research results as add-on modules to "standard" software; and
- a steady flow of incremental enhancements from research into application.

The architecture must be compatible with the WfMC reference architecture and API and other standards.

2.2.3 Current research trends and active projects

It may be argued that software technology has developed by reducing the amount of information that a computational entity needs to know at runtime in order to be able to interoperate with other entities. Components have a more tightly defined interface than objects and provide interoperation primitives in the form of events that pass highly informative objects to receiving entities. Software agents that use communication languages based on speech acts are sometimes presented as being the logical next extension of this trend. A number of collaborative and individual research projects have been and continue to be conducted in the application of agents to workflow management. Some of these collaborative projects include: ADEPT,[3] ENTERPRISE, TBPM (UK collaborative projects), EURESCOM Project P815, or SWIM.

[3] http://www.informatik.uni-ulm.de/dbis/f&l/forschung/workflow/ftext-adept_e.html

2.2.4 Open issues

As it has already been said, it is important for PLANET to remain agnostic over middleware technology (CORBA vs. COM vs. Java Beans, .NET, Jini, etc.) until matters resolve themselves in the real world. However, to facilitate research collaboration, it would be better to pick one of these technologies on which to base a common research platform architecture. So, an open issue would be which one is closer to the demands of workflow tools and applications. Other related issues would be what type of structures have to be exchanged through different applications in relation to workflow and/or P&S tools.

2.2.5 Research goals

The issue of infrastructure is strongly connected to development of tools. From the point of view of a research in AI, one potentially important goal would be the development of planning and scheduling servers that can be accessed by software components in the same way as back-end application software. Given that the P&S community already have a standard language, such as PDDL, it can facilitate the effort to have such a server application.

2.2.6 Recommended actions

Actions to be carried out with respect to infrastructure, relate to the previously discussed issues:

- draw up a reference architecture that can help foster the development of different modules that can be integrated in many different ways;
- agree on interface standards for research collaborations within the network, using the reference architecture as the baseline; or
- setting up a working group to establish interface standards for research cooperation with this respect.

2.3 Domain and business modelling

This section is concerned with methods, tools, languages, etc. used to model businesses and other application domains. Therefore, it has a strong connection to the Road map on Knowledge Engineering being developed within the corresponding PLANET TCU.

2.3.1 Introduction

In business process management, the purpose is to design, improve or define more precisely the organisation and its processes. In AI P&S, the purpose is to find a way of modelling the domain (and problems) that enables planning and scheduling techniques to be performed effectively. These purposes are entirely compatible, and there are many similarities between the representation languages used. In business process management the objective consists normally on defining a set of processes (divided in activities). Instead, in AI planning, a space of plans/processes is defined (in terms of operators and/or task networks). That is, while in AI many different possible plans can be generated from the domain description, in BPM usually they only handle a small set of processes. With the growing requirement for flexibility and adaptability, modelling for process management is likely to move closer to the "planning" model in the future. There is also a difference, between business modelling and KE, in the stage of development at which they are applied. Business modelling is performed when establishing requirements, whereas KE is performed during early stages of software development.

Where workflow management systems are used, there are two distinct stages of modelling. The first is elicitation/documentation or design of the organisation model (including business processes, resources and organisational structure). This may be documented, analysed, and simulated using high level modelling tools such as ARIS[4] or the ones in `http://dmsweb.badm.sc.edu/bpr/aa-5.htm`. In this setting, representation formalisms and storage formats tend to be proprietary. The second step is to produce a lower level model suitable for execution by a workflow engine. At this level, there is work by the WfMC towards an official vendor-neutral process definition language, though still different workflow products differ considerably in the style and syntax of input required. It is generally oriented to the requirements of automated execution of the flow of documents and control between task-performing resources. It says little or

[4]http://www.ids-scheer.com

nothing about the nature and semantics of the tasks that are linked in this way. Although some high-level modelling tools do claim to generate workflow definitions for specific workflow engines, this capability is generally felt to be inadequate at present. Often engineers must write the workflow definition using a modelling tool associated with the workflow engine, using the high level model as a reference. It is hoped that tools from AI planning may help bridge this gap.

If planning is to be used for BPR problems, the first step would be to think at a high level of what inputs of a planner correspond to the knowledge that BPR tools use, as well as what output of the planner corresponds to what knowledge on BPR tools. At a high level, one could establish the following relation:

- Inputs of a planner:
  - Domain theory: usually composed of a set of operators in STRIPS-like (PDDL) language (described in terms of pre- and post-conditions). Each BPR domain (e.g. the accounting domain in an organisation) can be defined in terms of a set of activities (here, the terminology can vary and use other words as tasks, or, even, processes) that are performed by organisation agents (either human or software). Therefore, there is a strong relation between operators in planning and activities in BPR, but it is not clear yet how to go from an activity based representation (agents responsible of a task, resources to be used, time that it takes to perform it) to an operator based representation (pre- and post-conditions, and, in some cases, other issues such as time constraints).
  - Problem: in planning, problems are described in terms of an initial state and a set of goals. They represent particular instances of situations for which one would like to have a solution. For BPR, a problem might be described as a process that has to be designed (modelled) for a particular task to be performed within the organisation. For instance, modelling the purchasing of an organisation, or the process of installing a new telephone line at a given address.
  - Initial state: in planning, one has to specify the starting situation of the posed problem. In the case of the BPR domain one would have to represent all knowledge that the organisation has about itself and can be used for the modelling of a specific process within the organisation. For instance, the hierarchical and/or functional representation of the organisation, the resources that it can use in its processes, or the documents that are generated within the organisation and travel around, being filled in, or filed, etc.
  - Goals: they describe in planning what one would like to be true at the end of the solution of the problem; that is, a set of assertions that have be true in a final state. In the case of BPR, this might be represented by the business goal of the organisation with respect

to that process. For instance, a purchase has to be done, having in mind a set of time or cost constraints.

- Output of the planners: usually AI planners generate a plan or set of plans. A plan can be seen as a sequence of operator applications that can lead from the initial state to a state in which the goals are reached. In the case of BPR, most processes are sequences of activities, adding conditional branches. Therefore, one would have to work on the generation of conditional plans, if a "typical" BPR model wants to be built.

There has been some preliminary work on automatically generating a planning domain theory, plus a problem description, from a representation of an organisation set of activities, such as the one reported in [Rodríguez-Moreno *et al.*, 2001].

2.3.2 Current state of the art

The state of the art is characterised by the large number and diversity of representation languages and modelling tools and techniques available. This may be indicative of the importance of the topic and that a definitive solution is still a long way off. Only a subset of what is available is presented, and some other references are provided. The issues of knowledge modelling in this field have many different perspectives. Here, some of this points of view are presented.

Process management

Process management tools and languages focus on how to represent knowledge about how things are performed. There are many different modelling tools in the market already. They are often expensive and need skilled personnel, though there is considerable variation. A common criticism is that the modelling tools do not produce output in a format that is acceptable to process definition tools. Examples include: iThink (relatively simple to use, systems-oriented),[5] ARIS[6] and ARENA (both complex, include simulation tools), or ProSim/ProCap. For a more detailed list of tools, refer to `http://dmsweb.badm.sc.edu/bpr/aa-5.htm`. Also, there is a variety of languages for describing such knowledge. Examples of standard process description languages include: IDEFn,[7] PIF,[8] EPIF, PSL,[9]

[5]http://www.hps-inc.com/bus_solu/ithink/ithink.htm
[6]http://www.iwi.uni-sb.de/teaching/ARIS/aris-i/aris-e-i/index.htm
[7]http://www.idef.com/
[8]http://ccs.mit.edu/pif1.html
[9]http://www.mel.nist.gov/psl/

WPDL,[10] CPR (Core Plan Representation),[11] and SPAR.[12] Apart from that, there are many other proprietary languages associated with particular tools.

Knowledge management

Knowledge management deals with the task of explicitly representing what an organisation knows, not only about its environment, but also, and more importantly, about the organisation itself. What the organisation knows can be thought in terms of what the current people in the organisation know, or even knowledge that people that have worked in the past in the organisation and do not longer work there had. Perhaps, this has been one of AI's more important contributions to the field of computer science: what you know about the world should be *declarative and explicitly* represented in an inspectable format, so that you can later reason about that. Most current approaches to knowledge management rely on the creation and maintenance of an intranet with information about organisation projects (past or present). In case the organisation would have a well defined structure with such knowledge about the organisation and its people, then we could reason about this knowledge when modelling organisation processes. So, we could assign activities to the most appropriate people with respect to many different parameters.

AI Planning

As we have already said, plan description/modelling languages include: ADL (Action Description Language), PDDL (Planning Description Domain Language), TF (O-Plan language),[13] the domain description language of IxTeT, HSTS-DDL,[14] the ones in[15] and theory of action formalisms. The knowledge engineering TCU agreed that PDDL was seen to be deficient in that it is not equipped with a methodology or language structure that helped the planning domain modeller. STRIPS/PDDL was likened to a "low level" language - theoretically expressive but not pragmatically expressive enough. Also, the underlying STRIPS-assumptions were thought to restrict the usefulness of the language. Some effort is being devoted currently to developing modelling tools for use with AI planners. An example is the PLANFORM project that has generated the GIPO

[10] http://www.wfmc.org
[11] http://projects.teknowledge.com/CPR2/
[12] http://www.aiai.ed.ac.uk/~arpi/spar/
[13] http://www.aiai.ed.ac.uk/~oplan/
[14] http://www.cs.cmu.edu/afs/cs/project/ozone/www/PCP/hsts.html
[15] http://www.informatik.uni-ulm.de/ki/Biundo/publications/publications.html

tool [Simpson *et al.*, 2001], developed by people from Universities of Durham and Huddersfield at UK.[16]

Ontologies

Ontologies are a way of expressing what is known about a given domain by using a hierarchy of concepts, their relations, attributes, and a set of axioms [Heflin and Hendler, 2000, Davis *et al.*, 2002]. Perhaps, the most well known ontology effort has been CYC, developed by Douglas Lenat and his group [Lenat and Guha, 1990]. Work on organisation ontologies has been conducted by several groups. We could highlight in relation to organisations modelling:

- TOVE project at University of Toronto [Fox and Gruninger, 1998].[17]
- Enterprise [Stader, 1997] and TBPM [Stader *et al.*, 2000] projects at AIAI, University of Edinburgh.[18]
- MIT's Process Handbook is an evolving repository of business process knowledge. It is also available in "shell" form for organisations to populate with their own knowledge [Malone *et al.*, 1999][19]

A lot of work on ontologies has been sponsored by DARPA, e.g. the knowledge sharing effort, though it is not clear whether organisation modelling has been addressed specifically. See also the proceedings of the AAAI 94 workshop on AI and business process re-engineering for other approaches.

Software/knowledge engineering

Software engineering modelling languages and knowledge representation techniques also cover similar ground. The Unified Modelling Language (UML) has become dominant in object oriented software engineering and is increasingly being used in business modelling too. Recent extensions improve its usefulness in modelling processes. The CoRE (Controlled Requirements Expression) method originally developed by British Aerospace (Military Aircraft) and Systems Designers has been used by AIAI in conjunction with TF/O-Plan and is also the basis of the COGSYS EnCore tool for requirements engineering. The best-known knowledge engineering method is CommonKADS [Schreiber *et al.*, 1993].[20]

[16] http://scom.hud.ac.uk/planform/
[17] http://www.eil.utoronto.ca/eil.html
[18] http://www.aiai.ed.ac.uk/
[19] http://ccs.mit.edu/ph/
[20] http://www.commonkads.uva.nl/

From the connection with new ways of making business, through the Web, new standards on languages should be cited here. In the future, more and more organisations will describe, at some level of detail, their processes using formalisms developed for Web-based information exchange. The most well known example now is XML (eXtended Markup Language), though new ways of representing things are appearing in combination to ontologies standards, such as RDF, DAIML, OIL, or DAIML-OIL.

2.3.3 Requirements

A long term requirement would be the integration of (most) languages and standards to improve interoperability; that is be able to generate tools that can handle several of these languages when using only one (by creating the appropriate interfaces, for instance). This would also allow to minimise requirements for re-training when changing to a new tool.

In relation to the usability aspect, a requirement would be to generate easy to use tools combining modelling and simulation capabilities. It seems like current tools are either very easy to use, but with a very poor representation mechanism underlying the tool, or have a powerful set of tools, but are too complex to be used without extended training.

Another requirement is the ability to integrate AI P&S domain models languages with the ones needed for process definition tools. This would allow to use their features, improving the usability ratio. Also, it would allow to re-use the already available processes that had been generated with those tools.

2.3.4 Current research trends and active projects

A lot of work is currently devoted to defining web-based languages, in the context of the Semantic Web. Perhaps, this is the more active point of view related to knowledge modelling, so future applications and tools will certainly used some form of language or standard arising from this field. These efforts have been described in the section on current state of the art. With respect to the workflow management world, as it has been mentioned before, the WfMC is working on a reference standard that would help on defining the interfaces of workflow tools with other systems, or among them.[21]

[21] http://www.wfmc.org

2.3.5 Open issues

Given that the generation of common interfaces, languages and standards is a very active field now due to the Web efforts, the open issues relate to how to integrate those efforts to workflow modelling and AI P&S modelling. So, it would be needed to answer to the following question: what is the best way to synthesise business process management, AIP&S and ontology modelling languages?

Also, domain experts have stated a requirement for new simulation tools. It is not clear why they believe this is needed. Can we clarify the requirement given that powerful simulation tools do exist? Why are existing tools not good enough? It is the price or the ease of use?

2.3.6 Research goals

Related to what has been said, a research goal would be to define a modelling language for organisations and their activities that has, at least, the following set of features:

- domain experts are comfortable with: they can understand it and write in it;
- has a rigorous semantics;
- has textual and graphical representations, so that there is consistency; between this two types of presenting information to the users of workflow tools
- is executable, for simulation purposes;
- can be easily (syntactically) translated into other standard formalisms;
- is mathematically formal so that the (static and dynamic) properties of models can be analysed; and
- planning techniques can be applied to.

Another goal would be defining means of creating and managing a library of processes in the spirit of the field of Case Based Reasoning. This should include, among other features, means of verification and analysis of redundancy, fast and appropriate retrieval of related processes, or easy update and version control of stored processes.

2.3.7 Recommended actions

In order to generate the above mentioned modelling language and Case Based approach, some potential actions to be taken could be:

- define a taxonomy of description languages, which includes features from current "de facto" standard languages, plus features from related fields, such as Semantic Web languages, or XPDL from WfMC.
- write a report documenting an agreed comparison and classification of Process Management, AI P&S and Ontology modelling languages
- write a comparative survey of existing software tools
- organisation of a "hands-on" workshop on languages for modelling in relation to the fields of AI P&S, the Web, e-Commerce, or workflow management. A preliminary workshop has been organised held in conjunction with ECAI'02. See TCU Web page for a report on that workshop.

2.4 Planning and scheduling

In this section, we will only describe from a high level what interactions may have AI P&S with planning and scheduling from the point of view of workflow management.

2.4.1 Introduction

The field of AI planning and scheduling is a very active one currently, mainly due to the arrival of new techniques that have enormously increased the solvability horizon of planners. We will not overview here all that has been done in the field. A good overview of the field until its publication can be found in [Allen *et al.*, 1990]. For an updated account of research that is going on since then, one should consult recent conference proceedings in the field, such as ECP, AIPS, AAAI, or IJCAI. Also, relevant papers are published in journals such as *Artificial Intelligence Journal* or *Journal of Artificial Intelligence Research*.

Finally, some authors are currently preparing text books on the subject, such as Ghallab, Nau and Traverso on one side, and Kambhampati on the other.

There are two aspects that are relevant in the relation of AI P&S to workflow management applications.

- In many workflow applications, combinatorics are significant. It is usual to see, in a medium-big organisation, many people involved (or potentially involved) in different processes. Generating all possible combinations of assignments of people to tasks and computing the optimal combination is an NP-hard problem. AI P&S due to its heuristic nature can help on reducing the complexity of this process, even if an optimal solution is not always found.

- As explained in the section on Human issues, in many cases it is very important the support to users decisions, rather than just the automation of their tasks. In order not to help the users, it is sometimes better leaving them with options and showing them alternatives, than providing them with a fixed set of tasks in a fixed order. Given that one of the main claims of AI is the explicit representation of knowledge, this can help on providing explanations to users, or computing different alternative solutions with different quality criteria (as it will be discussed in Section 2.6).

2.4.2 Current state of the art

Even if, from a research point of view, the field of AI P&S is a very active one, from an application point of view, its transfer to industry is rather slow. Some characteristics of the current state of the art are:

- There are many planning techniques already and software freely available.[22] Every two years a planning competition is held with better and better techniques, whose implementations are made publicly available.

- However, these techniques and systems need to be applied by an expert on planning technology due to the lack of interfaces, and the strong knowledge on the technique that is needed to understand which technique to apply for each problem, and how to make better use of it (parameterisation). That is, the application of current AI P&S technology has to be done on a case-by-case basis.

- There are, though, some commercial scheduling tools (e.g. ILOG Schedule). These are not entirely general purpose, so they need to be customised for new applications. Also, they need to be used by experts in the tool.

[22]See PLANET repository at http://scom.hud.ac.uk/planet/repository

- With respect to applications, there have been already some successful industrial applications of planning and scheduling, which have shown its profitability. From elevators, satellites, or robots control to the design of ships, there is a wide range of systems that are being used (or have been used) in the industry.

2.4.3 Requirements

From all that has already been said, it is obvious that the next step consists on the integration of scheduling into process management (and execution) tools. So, we have to understand what steps to follow in order to incorporate this AI techniques into currently used workflow tools. One of the roles of PLANET consists on attracting people from industry, so that they understand the usefulness of AI techniques, so industry and academia are able to generate such integration. On one side, workflow tools have to provide a high level description that AI P&S techniques can use to generate plans+schedules. On the other side, AI P&S techniques have to be able to handle in an integrated way plans, resources and time, which is crucial for most workflow applications.

2.4.4 Current research trends and active projects

There is a large list of current research trends in planning that would take long to detail. Some of the most well known techniques and approaches currently are:

- Integrated planning and scheduling. Industrial applications require that planning and scheduling not be two separated processes, but they are integrated one way or another. Examples of known integrated systems can be found in [Ghallab and Laruelle, 1994, Muscettola, 1994, AIP, 1998].

- Fully instantiated models of planning. They rely on some form of search for solutions in the space of instantiated operators and states. Examples are the SATplan [Kautz and Selman, 1996], Graphplan [Blum and Furst, 1995], or HSP [Bonet and Geffner, 2001] approaches.

- Generating schedules that guarantee certain behaviours at run-time, i.e. that are robust against limited changes in the environment. This is very important for workflow, given that there is a lot of interaction with the environment through the execution of tasks that people carry out.

- Constraint based approach to scheduling. This is related to the PLANET TCU on Dynamic Scheduling, so we refer the reader for more information to this document.

- Mixed initiative planning. The goal of these approaches is to introduce the human in the modelling loop. This has been described on Section 2.1.

- Decision-theoretic planning. In many domains, the world is uncertain and/or non deterministic (execution of the same action does not always arrive to the same state). In those domains, it is needed to explicitly reason about probabilities of fluents to be true, as well as actions to cause certain effects. One way to approach this type of planning task is by representing the problem as a Markov Decision Problem (MDP) and solving it by using either Dynamic Programming [Bellman, 1957] or using Reinforcement learning [Kaelbling *et al.*, 1996]. Some approaches to solve this type of problems can also be found in [Blythe, 1994, Boutilier *et al.*, 1999].

- Model checking. It is based on the idea that planning can be formulated as verifying some formal properties of a formula [Giunchiglia and Traverso, 1999].

- Applications. We are currently witnessing many interesting real applications of AI P&S technology. Among them, we can cite the projects at Nasa-JPL (such as Mars missions, imaging sequencing of telescopes, or the Deep Space One experiment) [Muscettola *et al.*, 1998, Muscettola and Smith, 1997], DARPA projects (such as MOABS, rescue projects, or scheduling of military operations), virtual agents projects (such as the travel planners, or the Electronic Elves [Chalupsky *et al.*, 2002]), elevator control, or satellites control.

2.4.5 Open issues

In relation to planning and scheduling for workflow management, one of the main issue relates to how to best combine human capabilities with planning and scheduling. This is related to mixed-initiative approaches in the sense of deciding what is the type of work that can be automated, and which is the type of work that should not be automated. There might be different reasons for not implementing fully automated systems, such as humans like doing some planning and scheduling tasks, or they like to be in control of some issues. In those cases, AI people should get software to work out the combinatorics and present results for person to use or modify.

2.4.6 Research goals

The goals for AI P&S in this context are related to these open issues. Perhaps a crucial goal is creating the appropriate interfaces, so that planners are easy to use for domain experts. This could imply limited capability, though.

2.4.7 Recommended actions

One action that remains to be performed consists on the definition of graduated reference problems in the field of workflow planning and scheduling. This is one of the goals of this TCU for the next future, that we hope to achieve at least partially by the end of the NoE period.

2.5 Enactment / execution

This section describes the issues related to the application of a pre-defined process model (or plan) in the real environment.

2.5.1 Introduction

If we are dealing with workflow tools, enactment deals with the application of a process model in the target organisation. In case we are dealing with planning, execution deals with the application of plans to achieve the proposed goals. There are many common aspects that emerge from the comparison between the enactment of workflow processes and the execution of plans, such as tasks like monitoring, control, exception handling, adaptation, or interaction with the environment.

2.5.2 Current state of the art

Most of the current work on execution of plans belongs to the field of robot planning or integrated manufacturing, given that most work on planning has been devoted to developing faster or more powerful planners, and there have been very little applications of planning techniques within organisations. Within AI P&S, there has been work done in the following issues related to execution:

- conditional planning: in many domains it is very hard to think "a priori" about all possible outcomes of the actions in the environment, or the agents need to gather information (sense the environment) while they are executing a plan in order to know what should be performed next. Conditional planning generates plans with branches for solving this problem [Draper *et al.*, 1994, Kushmerick *et al.*, 1995, Peot and Smith, 1992]. When executing a plan, every time alternatives are found, the current state of execution is consulted and one branch is selected.

- decision-theoretic planning: this has been described in the Section on Planning and Scheduling.

- reactive planning: most work in robotic tasks deals with two types of planning: deliberative and reactive [Brooks, 1986]. Deliberative planning is used to generate high level descriptions of sequences (or sets of sequences) of actions to be applied, without consideration of the actual details of the plans. When execution begins, control is assigned to a reactive component that decides what to do in the real environment. It selects the next action to be performed according to the current state of the system and the desired goals. In many cases, the reasoning is as simple as a pre-defined algorithm, and in other cases, it performs a very narrow local search to decide what to do next. Usually, the reactive behaviour has been learned by using many different techniques, although in some systems it is an "ad-hoc" procedure built from scratch. The PLANET Roadmap of the TCU of Robot Planning provides an extensive description of this type of planning approach. Examples of pure reactive, or hybrid deliberative-reactive planning systems are [Arkin and Balch, 1997, Firby, 1996, Georgeff and Lansky, 1987, Kaelbling, 1987].

2.5.3 Requirements

In order to apply planning-execution techniques to enactment of workflow models, there is a need to consider the following set of requirements:

- techniques for monitoring execution: given the complexity of many organisation processes, it is very important to be able to continuously look at the enactment of the processes. Monitoring should report on tasks that are being delayed, aborted, or resumed. Sometimes the identification of any of these issues is very difficult. For instance, it is very common that people forget to notify the computer (monitoring software) the completion of tasks, or changes in tasks development. Also, in most situations, people tend to delay the execution of tasks as much as possible, allowing very few time for reaction.
- techniques for exception handling: related to the previous requirement, once a problem is encountered, there is a need of defining procedures for recovery of the flaw. There have been some work in the field of planning with respect to re-planning, and recovery from failures that could be of some help to workflow failures. From the point of view of workflow systems, exception handling is usually performed by “ad-hoc” procedures that applied given that a problem is detected. In contrast, by allowing a declarative representation of operators, the system might be able to reason about possible failures and how to solve them.

2.5.4 Open issues, research goals, and recommended actions

Analysing the current research trends and the state of the art on the execution of plans, the following is a set of subjects that are pending to be solved:

- how to combine user preferences: usually a process that is being enacted is composed of many tasks to be performed by people with different roles and different qualifications. Assigning a task to a human is usually performed having in mind the set of roles that s/he is able to perform. However, there are other aspects that are worth considering such as user competence (even if a person is able to perform a given task, in what type of tasks is s/he really good at?), or user preferences (what type of tasks does s/he really enjoy performing?). Also, given that some tasks have to be performed by a group, this arises issues such as how to arrange the most productive group.

Some of this issues are strongly related to the currently very intensive field of knowledge management. This generates the goal of developing theories that define and reason about user models with respect to the assignment of tasks

- flexible working with overall plan: in most cases, detailing all possible aspects of the tasks can cause users to loose interest on their work. It is a better policy to provide some level of freedom for the execution of tasks, allowing decisions to be made by the human when executing a task. This is analogous to the integration of deliberative and reactive planning in robot tasks. A given degree of reactivity, allows the robot to be better prepared to cope with uncertainty and non determinism. A given degree of freedom, also allows the user to be prepared for uncertainty and non determinism, but also influences his/her way of looking at the work. The issue is how to combine the overall plan with the specific interests of the humans that have to carry out the plan steps. The associated goal would be the definition of models that allow to generate processes at various levels of detail, and interleave the execution of a high level plan with a somewhat reactive component.

- can plan repair or re-planning techniques be helpful in exception handling / jeopardy management?: as mentioned before, there has been some work done from the planning perspective with respect to handling plan failures during execution. It is not clear how this work should help and/or influence the workflow jeopardy or failures. It would be needed to study the sets of possible failures that can occur within the enactment of a process, and the set of repair procedures for those failures.

- how to provide personalised view of process (visualisation of big picture): another of the features that users find very important when performing a task of a process is knowing issues such as: why am I doing this?, where does this document come from?, or who should read this document afterwards? All of them deal with the problem of giving the users the ability to inspect at a certain level of detail the connections between the activity they are performing and the overall picture of the whole set of processes of the organisation. A research goal in this respect would be the description of variable visualisation techniques for parts of processes and the relationships among the processes of an organisation, having in mind security issues.

- how to combine and interleave plans for multiple humans (agents): if a distributed plan has been generated, the execution of that plan should monitor the interactions of the plans for each agent and combine the executions in the most effective way. Also, it should solve problems arising from the failure in an agent plan that has connection with other agents plans. The definition of a protocol of communication and negotiation between agents plans, execution of plans, failures, and repair methods would be needed.

2.6 Adaptation, optimisation and metrics

In this chapter, we will discuss an increasingly important aspect of workflow enhancement: how processes can be optimised/adapted according to design or enhancement problems. We will also discuss about the metrics considered for changing the processes.

2.6.1 Introduction

In general, there are two places in the application of workflow technology to organisation processes in which changes to the processes are involved:

- Design phase: when designing a given set of processes, the user might want to obtain an optimal process model according to a set of metrics and constraints. Usually, time and cost have been the only metrics considered for optimisation. Also, optimisation has been mainly a manual process, helped by the use of (sophisticated) simulation and analysis tools.
- Enhancement phase: when a process is being enhanced, many mismatches (might) occur between the designed process and its actual implementation. The role of adaptive workflow would be to feed the design and/or enhancement with those mismatches in order to optimise/adapt the process to the real situations.

Following the analogy between the process of applying planning technology and workflow technology that appears in chapter 0, there are several aspects that workflow and planning have (or not) in common with respect to optimisation:

- Design phase: the goal of both tasks (planning and workflow enhancement) is to obtain a process (plan) to be enhanced (executed) in the "real" world. However, while workflow has always considered optimisation (of time and cost) as a part of its design phase, it has not always been the case for planning. In the case of planning, the main emphasis has traditionally been on satisfying a goal, rather than on finding an optimal plan. This is mainly so, due to the already inherent complexity of finding "a" plan in many complex problems. When plan quality is considered, it has been mainly computed as "plan length", instead of using any user defined metric.

  However, there is a growing interest in the planning community for solving problems searching for optimality, or at least for better solutions [Drabble *et al.*, 2002, Nareyek, 2001]. In some cases, planners try to find an optimal plan according to a predefined criteria, such as makespan (total time to execute a partially ordered plan) or number of steps in the solution (in case of a totally ordered plan) [Haslum and Geffner, 2000, Williamson and Hanks, 1994]. In other cases, they take a plan as input and try to improve it [Ambite and Knoblock, 2001]. Others learn control knowledge to guide the planner towards "good" solutions [Aler *et al.*, 2002b, Borrajo and Veloso, 1997, Estlin and Mooney, 1996, Iwamoto, 1994, Pérez and Carbonell, 1994, Ruby and Kibler, 1992].

- Enhancement phase: the second main goal of both tasks is to enhance (execute) the designed process (plan). Here, we also find some differences between workflow and planning. Workflow enhancement is currently very widely done, so most organisations that have been (re-)designing their processes are following them. However, very few applications of planning systems have been built and used. Therefore, from the optimisation/adaptation point of view, there are many more lessons to be learned from workflow applications than from planning applications. Since optimisation/adaptation coming from the enhancement (execution) needs to know what types of failures can occur within the execution of a process (plan), we might have more information coming from workflow.

Listing all possible metrics is an infinite task. However, there are some that have been considered in many applications:

- Cost: measured by whatever means. Currently, ABC analysis is commonly carried out within business processes.
- Time: usually measured as time steps of the process, or the makespan.
- Quality: e.g. defect rate in a product, delays and dropped packets in a network.
- Value of the end-product: e.g adding an extra processing stage may increase the value of the end product more than it increases the cost.
- Flexibility: the ability to change processes quickly is important.
- Processes that are highly optimised with respect to cost or time may well be inflexible.
- Robustness: the probability of success of the processes.

A related issue is the use metrics to motivate and assess the performance of people. Inappropriate metrics can have the opposite effect to that intended. For example if targets are perceived as impossible, then people will ignore them. Thus if a target is made more demanding it may in fact decrease performance. Similarly, taking a call centre as an example, an "obvious" performance metric is number of calls handled per day. However, this encourages staff to keep calls short, which may mean that poor answers are given leading to more calls. This improves apparent productivity, but customer satisfaction goes down. The "correct" productivity metric must take into account whether the caller was satisfied, but this is more difficult to measure.

In the next sections, we discuss issues related to optimisation with respect to: open questions; research results; barriers to technology transfer; and software and application requirements.

2.6.2 Current state of the art

The following is a set of results that might be used to approach the open questions of previous section:

- There are all types of mature optimisation techniques coming from AI and operations research such as: heuristic search; genetic algorithms; or linear/dynamic programming.
- There have been some approaches on planning for better solutions and learning to plan for better solutions that have been mentioned in the introduction of this section.
- Also, recently there is an interest towards using multiple criteria and considering them for planning or scheduling [Drabble *et al.*, 2002].

2.6.3 Requirements

Here, we discuss what the workflow tools and applications should have in order to allow optimisation:

- Integration with process design and enhancement tools: optimisation and adaptation procedures should be integrated on one hand with process modelling techniques (for obtaining good models), and, on the other, with process enactment tools (for adapting the models according to actual enactment of the processes)
- Interaction with the user: an important aspect of the tools consists on allowing the user to interact with the optimisation and adaptation procedures so that s/he is able to direct towards process models that comply with user expectations

- User-definable metrics and optimisation parameters: the user should be able to provide in a given language descriptions on how metrics should be computed, as well as parameters for controlling how optimisation and adaptation should be performed

2.6.4 Open issues, research goals, and recommended actions

The set of open questions with respect to optimisation/adaptation and metrics are:

- Do workflow applications really need metrics different than time and cost? If we are going to define tools for performing adaptation/optimisation according to user defined metrics, we should first make sure that users will need different types of metrics. A possible recommended action would be to survey in some organisations about this aspect.
- What language should we use to provide those metrics to the system? We should study what are good languages for describing those metrics, so that potential users of the tools are able to easily define metrics by themselves. PDDL2.1 has advanced on defining such language, by allowing the user to specify them in the language. Other planner-specific approaches allow also defining quality-based criteria [Borrajo *et al.*, 2001].
- If multiple agents are used, how should their respective metrics be combined/negotiated? Should it be left to execution time or should it be worked out before execution starts?
- What is the set of possible failures of a process (plan)? Although this question also appears in the section on execution, within this section, it refers to the generation of plans that are optimised according to, for instance, less probability of failure
- How should workflow enhancement influence optimisation/adaptation? This issue is related to the plan repair techniques in the execution section
- Where should design/enhancement optimisation knowledge come from? There might be three different types of sources: experts on a given domain (they usually know what models are wrong and why, what resources should be assigned to what task, etc.); experts on BPR or workflow enactment (usually they work on consultancy firms and provide advice on how different organisations implement their processes); learning from past executions of the workflow or from the history of the processes execution in the organisation.

- Are there experts on resolving failures of execution, or anticipating problems? This issue is related to the previous one. Usually, in big organisations there is people in charge of this task that could be of great help
- Can the systems recognise a "good" solution? Or how do we define procedures for computing how good a model is?
- How should the interaction with the user be integrated when optimising? Optimising a process might result in a less intelligible process, so an analysis on what is preferred.

When trying to apply optimisation to process design/enhancement, the following is a list of possible and actual problems:

- The user might not know/distinguish when s/he needs optimisation.
- How does the user describe optimisation and metrics knowledge?

Chapter 3 Summary and conclusions

This document has presented the second version of the PLANET R&D Road Map for AI Planning and Scheduling applied to Workflow Management. In an applied discipline such as this, a Road Map must not only identify research challenges, but also match them to current and projected end-user requirements. It must also consider the process by which the results are incorporated into the tools of the trade of the end-users and application developers. Furthermore, necessary preconditions for successful application of the results must be taken into account. This version is an important step towards a coherent strategy, but is not itself the definitive answer. The Road Map needs to be a living document that is developed and updated and regular intervals.

3.1 Main achievements

One of the main achievements to date has been to develop an understanding of how the “world view”, vocabulary, challenges, etc. of Business Process / Workflow Management relate to AI Planning and Scheduling. This has been possible because of the active participation of a small number of workflow and process management experts from end-user organisations and consultancy companies. The site visit to BT to gather information on existing (non-AI) software applications was also extremely valuable in this regard. The TCU must make every effort to involve more end-user representatives (not just researchers, but problem owners) from a spectrum of industries. A number of commercial software vendors are registered on the TCU mailing list but have not as yet participated actively. It is important for such organisations to become actively involved.

For planning techniques to be of practical use they must be integrated with, or must interface to, commercial workflow management systems (WfMS) and other related software.

Requirements have been classified as short, medium and long term as follows:

- short term: address short-comings in current-generation process management software. The most important items in this category are: integration of scheduling and resource allocation/management algorithms into workflow management software; and incorporation of a planning capability to enable a WfMS to modify the process instance automatically during execution, to cope with failure, changed objectives, and other exceptions.
- medium term: Current generation workflow software handles high volume routine processes, typically involving low-skill workers. The medium term requirements concern extending this support to high-skill knowledge workers. This may involve, for example, building process awareness into software tools.
- long term: More radical (e.g. adaptive self-organising) approaches addressing the need for organisations to function in a business environment that is increasingly uncertain and subject to change.

This document has also made a start on identifying planning techniques and research goals that address these requirements. In addition to the application of planning and scheduling algorithms we discussed: advantages to be gained from using AI plan representations for processes, ideas from plan execution (especially in uncertain environments), and work on adaptation optimisation and metrics. Further work remains to be done, however, to identify specific research goals and projects. Two further topics are also discussed: human issues and infrastructure. It is important to remember that much of the work in a business process is performed by people. Often technology is seen primarily as a means of cutting costs through automation rather than enhancing value by enabling people to work more effectively. The result of treating people like machines is often demotivation, high staff turnover, loss of productivity, etc. In addition, human qualities are under-utilised. There is a danger that must be guarded against that planning and scheduling techniques may make this situation worse rather than better. The discussion of infrastructure mainly focuses on the need for a reference architecture and interface standards to AI-based software tools to be integrated with each other, with conventional process management software, and with the general organisation infrastructure.

3.2 Main recommendations

The TCU main objective was to play a useful role in closing the gap between industry and academic research. However further work is needed:

- to make researchers aware of the real challenges and constraints of the workflow domain;
- to make application and tool developers aware of what AI planning and scheduling research has to offer;
- to address practical issues of integrating planning and scheduling technology into suites of application software, and of making the techniques usable by typical software engineers, analysts, etc.
- to form a consensus on medium and long term research goals. The Road Map should be seen as a living document and be extended and updated regularly.

Bibliography

[AIP, 1998] *Proceedings of the AIPS-98 Workshop Integrating Planning, Scheduling and Execution in Dynamic and Uncertain Environments*, Pittsburgh, PA, June 1998. AAAI Technical Report WS-98-02.

[Aler *et al.*, 2002a] Ricardo Aler, Daniel Borrajo, David Camacho, and Almudena Sierra-Alonso. A knowledge-based approach for business process reengineering, SHAMASH. *Knowledge Based Systems*, 15(8):473–483, 2002.

[Aler *et al.*, 2002b] Ricardo Aler, Daniel Borrajo, and Pedro Isasi. Using genetic programming to learn and improve control knowledge. *Artificial Intelligence*, 141(1-2):29–56, October 2002.

[Allen *et al.*, 1990] James F. Allen, James Hendler, and Austin Tate (eds.). *Readings in Planning*. Morgan Kaufmann, 1990.

[Allen *et al.*, 1995] James F. Allen, Lenhart K. Schubert, George Ferguson, Peter Heeman, Chung Hee Hwang, Tsuneaki Kato, Marc Light, Nathaniel G. Martin, Bradford W. Miller, Massimo Poesio, and David R. Traum. The TRAINS project: A case study in building a conversational planning agent. *Journal of Experimental and Theoretical AI (JETAI)*, 7:7–48, 1995.

[Alonso *et al.*, 1997] Gustavo Alonso, Divyakant Agrawal, Amr El Abbadi, and C. Mohan. Functionality and limitations of current workflow management systems. *IEEE Expert*, 12(5), 1997.

[Ambite and Knoblock, 2001] José Luis Ambite and Craig A. Knoblock. Planning by rewriting. *Journal of Artificial Intelligence Research*, 15:207–261, 2001.

[Arkin and Balch, 1997] Ronald C. Arkin and Tucker R. Balch. AuRA: principles and practice in review. *Journal of Experimental and Theoretical Artificial Intelligence*, 9(2), 1997.

[Aylett and Borrajo, 2002] Ruth Aylett and Daniel Borrajo. Workflow and intelligent manufacturing: Common problems, common solutions? *Planet News*, (3):25–26, 2002.

[Bacchus, 2001] Fahiem Bacchus. AIPS'00 planning competition. the fifth international conference on artificial intelligence planning and scheduling systems. *AI Magazine*, 22(3):47–56, Fall 2001.

[Beaudouin-Lafon, 1998] Michel Beaudouin-Lafon. *Computer Supported Cooperative Work (CSCW)*. John Wiley and Sons Ltd, 1998.

[Bellman, 1957] R. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ., 1957.

[Blum and Furst, 1995] Avrim L. Blum and Merrick L. Furst. Fast planning through planning graph analysis. In Chris S. Mellish, editor, *Proceedings of the 14th International Joint Conference on Artificial Intelligence, IJCAI-95*, volume 2, pages 1636–1642, Montréal, Canada, August 1995. Morgan Kaufmann.

[Blythe, 1994] Jim Blythe. Planning with external events. In *Proceedings of the International Conference on Uncertainty in AI*, 1994.

[Bonet and Geffner, 2001] Blai Bonet and Hector Geffner. Planning as heuristic search. *Artificial Intelligence*, 129(1-2):5–33, 2001.

[Borrajo and Veloso, 1997] Daniel Borrajo and Manuela Veloso. Lazy incremental learning of control knowledge for efficiently obtaining quality plans. *AI Review Journal. Special Issue on Lazy Learning*, 11(1-5):371–405, February 1997. Also in the book "Lazy Learning", David Aha (ed.), Kluwer Academic Publishers, May 1997, ISBN 0-7923-4584-3.

[Borrajo *et al.*, 2001] Daniel Borrajo, Sira Vegas, and Manuela Veloso. Quality-based learning for planning. In *Working notes of the IJCAI'01 Workshop on Planning with Resources*, pages 9–17, Seattle, WA (USA), August 2001. IJCAI Press.

[Boutilier *et al.*, 1999] Criag Boutilier, Thomas Dean, and Steve Hanks. Decision-theoretic planning: Structural assumptions and computational leverage. *Journal of Artificial Intelligence Research*, 11:1–94, 1999.

[Brooks, 1986] Rodney A. Brooks. A roboust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2(1):14–23, March 1986.

[Burnett and McIntyre, 1995] M. Burnett and D. W. McIntyre. Guest editors' introduction: Visual programming. *IEEE Compute*, 28(3):14, March 1995.

[Camacho *et al.*, 2001] David Camacho, Daniel Borrajo, José Manuel Molina, and Ricardo Aler. Flexible integration of planning and information gathering. In Amedeo Cesta, editor, *Sixth European Conference on Planning (ECP'01)*, Lecture Notes in Artificial Intelligence, Toledo (Spain), September 2001. Springer Verlag.

[Cesta, 2001] Amedeo Cesta, editor. *Proceedings of the Sixth European Conference on Planning (ECP'01)*, Lecture Notes in Artificial Intelligence, Toledo (Spain), September 2001. Springer Verlag.

[Chalupsky *et al.*, 2002] Hans Chalupsky, Yolanda Gil, Craig A. Knoblock, Kristina Lerman, Jean Oh, David V. Pynadath, Thomas A. Russ, , and Milind Tambe. Electric elves: Agent technology for supporting human organizations. *AI Magazine*, 23(2):11–24, Summer 2002.

[Davis *et al.*, 2002] J. Davis, D. Fensel, and F. van Harmelen. *Towards the Semantic Web: Ontology-Driven Knowledge Management*. John Wiley, 2002.

[desJardins *et al.*, 1999] M. desJardins, E. Durfee, C. Ortiz, and M. Wolverton. A survey of research in distributed, continual planning. *AI Magazine*, 4:13–22, 1999.

[Diimitrios Georgakopoulos, 1995] Amit Sheth Diimitrios Georgakopoulos, Mark Hornick. An overview of workflow management: From process modeling to workflow automation infrastructure. *Distributed and Parallel Databases*, 3:119–153, 1995.

[Dix *et al.*, 2001] J. Dix, Héctor Muñoz-Avila, and Dana S. Nau. Impacting shop: Putting an AI planner into a multi-agent environment. *Annals of Mathematics and AI*, 2001.

[Drabble *et al.*, 2002] Brian Drabble, Jana Koehler, and Ioannis Refanidis, editors. *Proceedings of the AIPS-02 Workshop on Planning and Scheduling with Multiple Criteria*, Toulouse (France), April 2002.

[Drabble, 1999] Brian Drabble. Task decomposition support to reactive scheduling. In Susanne Biundo and Maria Fox, editors, *Proceedings of the 5th European Conference on Planning (ECP-99)*, pages 200–212, Durham (UK), 1999.

[Draper *et al.*, 1994] Denise Draper, Steve Hanks, and Daniel Weld. Probabilistic planning with information gathering and contingent execution. In *Proceedings of the Second International Conference on AI Planning Systems, AIPS-94*, pages 31–36, Chicago, IL, June 1994. AAAI Press, CA.

[Estlin and Mooney, 1996] Tara A. Estlin and Raymond J. Mooney. Multistrategy learning of search control for partial-order planning. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, volume I, pages 843–848, Portland, Oregon, August 1996. AAAI Press/MIT Press.

[Fikes and Nilsson, 1971] Richard E. Fikes and Nils J. Nilsson. Strips: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2:189–208, 1971.

[Firby, 1996] James R. Firby. Modularity issues in reactive planning. In *Proceedings of the Third International Conference on AI Planning Systems*, pages 78–85, Edinburgh (UK), May 1996.

[Fox and Gruninger, 1998] M. Fox and M. Gruninger. Enterprise modelling. *AI Magazine*, pages 109–121, Fall 1998.

[Fox and Long, 2002] Maria Fox and Derek Long. *PDDL2.1: An Extension to PDDL for Expressing Temporal Planning Domains*. University of Durham, Durham (UK), February 2002.

[Georgakopoulos *et al.*, 1995] Dimitrios Georgakopoulos, Mark F. Hornick, and Amit P. Sheth. An overview of workflow management: From process modeling to workflow automation infrastructure. *Distributed and Parallel Databases*, 3(2):119–153, 1995.

[Georgeff and Lansky, 1987] Michael P. Georgeff and Amy L. Lansky. Reactive reasoning and planning. In *Proceedings of AAAI-87 Sixth National Conference on Artificial Intelligence*, pages 677–68, Seattle, WA (USA), July 1987.

[Ghallab and Laruelle, 1994] M. Ghallab and H. Laruelle. Representation and control in IxTeT, a temporal planner. In *Proceedings of the 2nd International Conference on AI Planning Systems*, 1994.

[Ghallab *et al.*, 2002] Malik Ghallab, Joachim Hertzberg, and Paolo Traverso, editors. *Proceedings of the Sixth International Conference on Artificial Intelligence Planning Systems (AIPS-02)*, Toulouse (France), April 2002. AAAI Press.

[Giunchiglia and Traverso, 1999] Fausto Giunchiglia and Paolo Traverso. Planning as model checking. In *Proceedings of the Fifth European Conference on Planning (ECP'99)*, pages 1–20, 1999.

[Haslum and Geffner, 2000] Patrick Haslum and Hector Geffner. Admissible heuristics for optimal planning. In *Proceedings of the Fifth International Conference on AI Planning Systems (AIPS-2000)*, pages 70–82, 2000.

[Heflin and Hendler, 2000] Jeff Heflin and James A. Hendler. Dynamic ontologies on the web. In *Proceedings of the AAAI-2000*, pages 443–449, 2000.

[Iwamoto, 1994] Masahiko Iwamoto. A planner with quality goal and its speed-up learning for optimization problem. In K. Hammond, editor, *Proceedings of the Second International Conference on Artificial Intelligence Planning Systems (AIPS94)*, pages 281–286, Chicago, Illinois, June 1994.

[Jackson *et al.*, 1993] P.R. Jackson, T.D. Wall, R. Martin, and K.P. Davids. New measures of job control, cognitive demand and production responsibility. *Journal of Applied Psychology*, 78:753–762, 1993.

[Jacobson *et al.*, 1995] Ivar Jacobson, Maria Ericsson, and Agnet Jacobson. *Object Advantage: Business Process Reengineering With Object Technology.* Addison-Wesley, 1995.

[Kaelbling *et al.*, 1996] Leslie P. Kaelbling, Michael L. Littman, and Andrew W. Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, pages 237–285, 1996.

[Kaelbling, 1987] Leslie P. Kaelbling. An architecture for intelligent reactive systems. In M. Georgeff and A. Lansky, editors, *Reasoning about plans and actions*, pages 395–410. Morgan Kaufmann, San Mateo, California, 1987.

[Kahler *et al.*, 2000] Helge Kahler, Oliver Stiemerling, Volker Wulf, and Anders Mørch. Computer supported cooperative work. *Journal of Collaborative Computing*, 9(1), 2000.

[Kautz and Selman, 1996] Henry Kautz and Bart Selman. Pushing the envelope: Planning, propositional logic, and stochastic search. In *Proceedings of AAAI-96*, pages 1194–1201, 1996.

[Kuhn and Muller, 1993] S. Kuhn and M. Muller. Participatory design - introduction. *Communications of the ACM*, 36(4):24–28, June 1993.

[Kushmerick *et al.*, 1995] Nicholas Kushmerick, Steve Hanks, and Daniel S. Weld. An algorithm for probabilistic planning. *Artificial Intelligence*, 76(1-2):239–286, 1995.

[Lenat and Guha, 1990] D. B. Lenat and R. V. Guha. *Building large knowledge-based systems. Representation and inference in the Cyc project.* Addison-Wesley, Reading, MA (USA), 1990.

[Malone *et al.*, 1999] Thomas W. Malone, Kevin Crowston, Jintae Lee, Brian Pentland, Chrysanthos Dellarocas, George Wyner, John Quimby, Charley Osborne, Avi Bernstein, George Herman, Mark Klein, and Elissa O'Donnell. Tools for inventing organizations: Toward a handbook of organizational processes. *Management Science*, 45(3):425–443, 1999.

[McDermott, 2000] John McDermott. The AIPS'98 planning competition. *AI Magazine*, 21(2), 2000.

[Mowshowitz, 1994] Abbe Mowshowitz. Virtual organization: A vision of management in the information age. *The Information Society*, 10(4):267–288, 1994.

[Mowshowitz, 2001] Abbe Mowshowitz. Virtual organization: The new feudalism. *IEEE Computer*, 34(4):110–112, 2001.

[Muscettola and Smith, 1997] Nicola Muscettola and David Smith. On-board planning for new millenium deep space one autonomy. In *Proceedings of IEEE Aerospace Conference*, Snowmass, CO, 1997.

[Muscettola *et al.*, 1998] Nicola Muscettola, P. Pandurang Nayak, Barney Pell, and Brian C. Williams. Remote agent: To boldly go where no AI system has gone before. *Artificial Intelligence*, 103(1-2):5–47, 1998.

[Muscettola, 1994] Nicola Muscettola. *Intelligent Scheduling*, chapter HSTS: Integrating Planning and Scheduling, pages 169–212. Morgan Kaufmann, 1994.

[Myers and Berry, 1999] Karen Myers and Pauline Berry. Workflow management systems: An AI perspective. Technical report, Artificial Intelligence Center, SRI International, Menlo Park, CA (USA), 1999.

[Nardi, 1993] Bonnie A. Nardi. *A Small Matter of Programming*. MIT Pres, Cambridge, MA, 1993.

[Nareyek, 2001] Alexander Nareyek, editor. *Working notes of the IJCAI'01 Workshop on Planning with Resources*, Seattle, WA (USA), August 2001. IJCAI Press.

[Oldham, 1996] G. R. Oldham. *International Review of Industrial and Organisational Psychology*, chapter Job design. John Wiley, 1996.

[Paolucci *et al.*, 1999] Massimo Paolucci, Onn Shehory, Katia P. Sycara, Dirk Kalp, and Anandeep Pannu. A planning component for RETSINA agents. In *Proceedings of the Agent Theories, Architectures, and Languages Conference*, pages 147–161, 1999.

[Pednault, 1989] Edwin Pednault. ADL: Exploring the middle ground between Strips and the situation calculus. In *Proceedings of the Conference on Knowledge Representation and Reasoning*, pages 324–332, 1989.

[Peot and Smith, 1992] Mark A. Peot and David E. Smith. Conditional nonlinear planning. In *Proceedings of the First International Conference on AI Planning Systems*, pages 189–197, College Park, Maryland (USA), 1992.

[Pérez and Carbonell, 1994] M. Alicia Pérez and Jaime G. Carbonell. Control knowledge to improve plan quality. In *Proceedings of the Second International Conference on AI Planning Systems*, pages 323–328, Chicago, IL, 1994. AAAI Press, CA.

[Rodríguez-Moreno *et al.*, 2001] María Dolores Rodríguez-Moreno, Daniel Borrajo, and Daniel Meziat. Planning-based generation of process models. In Lee McCluskey and Alfredo Milani, editors, *Proceedings of the ECP-01/PLANET Workshop on Automated Planning and Scheduling Technologies in New Methods of Electronic, Mobile and Collaborative Work*, pages 22–34, Toledo (Spain), September 2001.

[Ruby and Kibler, 1992] David Ruby and Dennis Kibler. Learning episodes for optimization. In *Proceedings of the Ninth International Conference on Machine Learning*, pages 379–384, Aberdeen, Scotland, 1992. Morgan Kaufmann.

[Schreiber *et al.*, 1993] A.T. Schreiber, B.J. Wielinga, and J.A. Breuker. *KADS: A Principled Approach to Knowledge Engineering*. Academic Press, London (UK), 1993.

[Simpson *et al.*, 2001] R. M. Simpson, T. L. McCluskey, W. Zhao, R. S. Aylett, and C. Doniat. An integrated graphical tool to support knowledge engineering in ai planning. In Amedeo Cesta and Daniel Borrajo, editors, *Preprints of the Sixth European Conference on Planning (ECP'01)*, Toledo (Spain), September 2001.

[Stader *et al.*, 2000] Jussi Stader, Jonathan Moore, Paul Chung, I. McBriar, M. Ravinranathan, and Ann Macintosh. *The Workflow Handbook 2001*, chapter Applying Intelligent Workflow Management in the Chemicals Industries, pages 161–181. Workflow Management Coalition (WfMC), October 2000.

[Stader, 1997] Jussi Stader. An intelligent system for bid management. *The International Journal of Project & Business Risk Management*, 1/3:299–314, 1997.

[Tate *et al.*, 1998] Austin Tate, Jeff Dalton, and John Levine. Generation of multiple qualitatively different plan options. In *Artificial Intelligence Planning Systems*, pages 27–35, 1998.

[Tate *et al.*, 2000] Austin Tate, John Levine, Peter Jarvis, and Jeff Dalton. Using AI planning technology for army small unit operations. In *Proceedings of the 5th International Conference on Artificial Intelligence Planning Systems (AIPS-00)*, pages 379–386, 2000.

[Veloso *et al.*, 1997] Manuela M. Veloso, Alice M. Mulvehill, and Michael T. Cox. Rationale-supported mixed-initiative case-based planning. In *Proceedings of the AAAI/IAAI-97*, pages 1072–1077, Providence, RI, 1997.

[Williamson and Hanks, 1994] Mike Williamson and Steve Hanks. Optimal planning with a goal-directed utility model. In K. Hammond, editor, *Proceedings of the Second International Conference on Artificial Intelligence Planning Systems (AIPS94)*, pages 176–181, Chicago, Illinois, June 1994.

[Zuboff, 1988] Shoshana Zuboff. *In the Age of the Smart Machine: The Future of Work and Power*. Basic Books, New York, NY, 1988.