Information support for knowledge-intensive business processes — combining workflows with document analysis and information retrieval

Andreas Abecker, Ansgar Bernardi, Heiko Maus, Claudia Wenzel

German Research Center for Artificial Intelligence (DFKI) GmbH

P.O.Box 2080, D-67608 Kaiserslautern, Germany
Phone:+49 631 205 3582, Fax:+49 631 205 3210

{aabecker,bernardi,maus,wenzel}@dfki.uni-kl.de

Abstract

Explicit modeling of business processes and their enactment in workflow systems have proved to be valuable in increasing the efficiency of work in organizations. We argue that enacted business processes - that is: workflow management systems - form a solid basis for adequate information support in complex and knowledge-intensive business processes.
To support this claim we demonstrate results from two different projects: The *VirtualOffice* approach employs workflow context information to support the high-precision document analysis and understanding in standard office settings; the combination of workflow context and document analysis techniques allow for the automatic handling of incoming paper mail with respect to the appropriate workflows. The *KnowMore* approach focuses on the support of people who work on knowledge-intensive tasks by automatic delivery of relevant and goal-specific information; the context of the workflow, an extended process model, and a detailed modeling of information sources are combined to this end. Both approaches show ways to proceed from workflow systems towards active knowledge management.

Introduction

Workflow Management is a widespread technology for automating structured business processes. Today it is mainly used to coordinate complex processes where many activities must be scheduled and dispatched among many possible agents. Further support comes from an integrated handling of application programs used in the process chain and a streamlined passing of application data and electronic documents flowing between different process steps.

Now consider today's complex and often document-driven business processes which are characterised by manifold interfaces and document and information exchange with the company's environment. Within such processes the respective employees use to deal and handle with a great amount of information and knowledge transferred by and embedded in explicit, often paper-based forms, letters, books, manuals, records, and all kinds of documents.

The knowledge and information applied to perform the several activities in a given process step typically are not represented in the WfMS. Nevertheless, it is usually available and often explicitly accessible somewhere in the information space of the company, since otherwise the employees would not be able to do their jobs.

For instance, imagine a workflow for buying the required hardware and software to equip a new project team. Here, the responsible clerks know which catalogues and websites to consult, or which colleagues to ask. Here, one would like the WfMS to automatically present relevant information sources to base the buying decisions on. Therefore, it should be possible to represent the service process supporting the actual business process by appropriate information and knowledge sources. To enact such a service process, the WfMS should possess interfaces for exchanging knowledge items with the surrounding system environment, which would be appropriate to enable powerful support services. Moreover, having such a semantically rich interface between the WfMS and its system environment would even allow to directly pump information items extracted from incoming documents to the appropriate places in the data models of the actual workflow instance. Such a mechanism would diminish the bottleneck at the company's interface to the "external world" which is often characterized by paper-based communication, or at least by electronic documents which obey other data models and conceptualizations than the company-internal information system. Such a semantically enriched interface between WfMSs and their environment could greatly improve their integration with the other parts of the enterprise information infrastructure, especially with document management systems.

Unfortunately, such considerations are not subject of current WfMS approaches ((Georgakopoulos *et al.* 1995), (Alonso *et al.* 1997)). Contemporary systems exhibit only a very thin interface for data exchange with other office applications such as text processing applications or corporate data bases.

In order to overcome this limitation, this paper will present two different solutions which describe how a

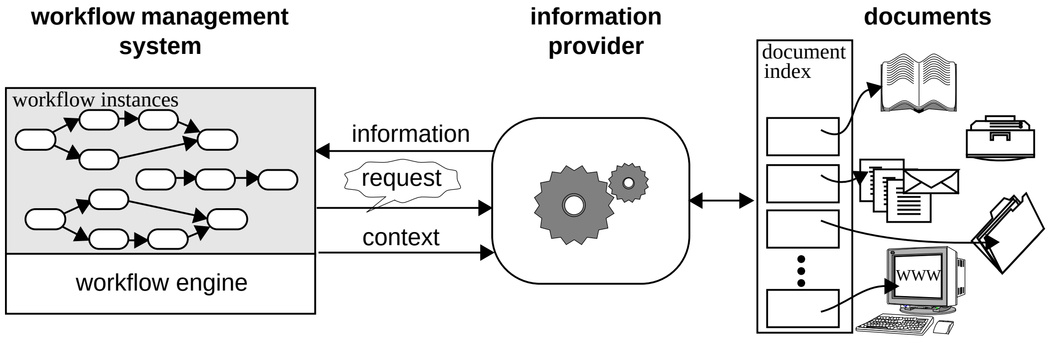

Figure 1: *Information Support for Workflows: A Common Description Frame*

true knowledge transfer between business processes and their surrounding information space can be established. Both approaches focus on process-embedded information delivery from documents.

Although the approaches have been come from completely different motivations, they nevertheless fit into the same common description frame which is sketched in Figure 1.

The WfMS represents running business processes by workflow instances and executes them by means of a workflow engine. The second system is the mediator system which we call *information provider*.

The information provider gets an information request and some additional context information from the WfMS. It accesses the documents by some kind of document index. This index may consist of an inverted index file as required for information retrieval tasks, but it might be any more sophisticated document model as well. The retrieved information is handed over to the WfMS. So the interface between the two systems comprises three different kinds of information.

Central to all our considerations is the notion of context. Since the context required highly depends on the actual information request, the common description frame does not answer the questions where necessary context comes from and what it consists of. We argue that context delivery may require different system architectures. This statement will be clarified by introducing two different solutions instantiating the general description frame:

The next chapter describes the *VirtualOffice* project (see also (Wenzel 1998)) which integrates paper-based information into any kind of workflow activity. Afterwards, chapter 3 introduces the *KnowMore* project (see also (Abecker *et al.* 1998)) which aims at supporting so-called knowledge-intensive tasks by pro-active document delivery. After a review of some related work in chapter 4, we summarize in chapter 5 with some general insights on possible solutions for the questions how semantically enriched interfaces can be established and how context information can be dealt with practically. This will hopefully lead to ideas for the future work in integrating and further developing process-oriented information and knowledge management.

The VirtualOffice Scenario: Information Support by Paper Documents

Although the paperless office has been a buzzword for many years now, it still has not come into reach. On the contrary, the enactment of business processes by administrative workflows has even complicated the integration of paper documents since such workflows require electronic representations of all documents involved.

Looking a little bit closer, such workflows are characterized by heterogenous documents which belong to one common process and arrive in a chronological order. A typical example is an insurance process with initial applications for contracts, changes in the policies, annual invoices, and damage claims. Another example are business trips where the traveller has to fill out an application, the application must be confirmed, some invoices, e.g., for plane tickets, have to be payed in advance and, finally, several receipts must be accounted for. Certainly, there are a lot of similar examples but within this paper we will use a purchasing process within a company as a basis for our examples.

The *VirtualOffice* project which being conducted at DFKI Kaiserslautern tackles the above-mentioned problem of integrating paper-based information into workflows. In order to achieve a solution, a system for document analysis and understanding (DAU) is used as workflow application which even benefits from this integration: As stated in (Baumann *et al.* 1997a) context information from business processes improves the analysis of business letters and is therefore a crucial point of our research.

In order to get an impression how document analysis and understanding works, figure 2 reveals typical analysis steps and results. Usually, DAU systems are divided into two parts: The analysis front-end of the system works as provider of uniform information about the characters appearing in a scanned document along with additional logical attributes. This information is interpreted by the system's document understanding components which form the back-end.

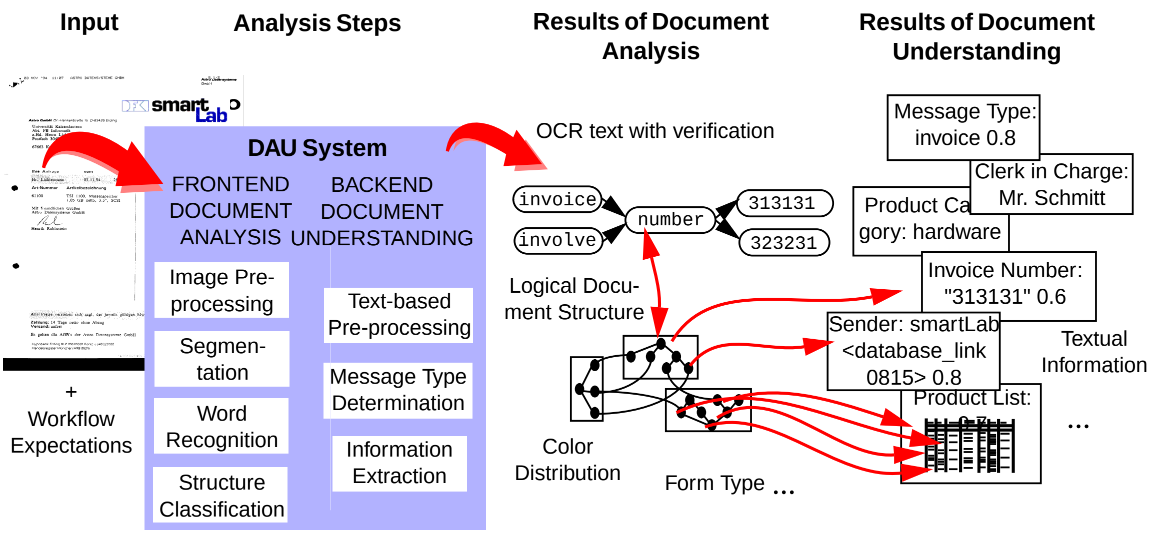

Figure 2: *Document Analysis and Understanding: Steps and Results.*

The *analysis* of a scanned document starts with low-level image preprocessing such as skew angle adjustment and upside-down detection. Afterwards, segmentation divides the document into geometrically connected components and identifies segments of characters, words, lines, and blocks. Then, text recognition explores the captured text segments, generates character hypotheses, and merges them into word hypotheses. Structure classification takes this given geometric structure to hypothesize the so-called logical objects of a document, e.g. title, author, chapter, etc.

Within *document understanding*, the generated word hypotheses are validated by dictionary look-up in the text-based preprocessing step. Now, the content-based part of analysis is invoked. First, the message type of a document has to be derived (e.g. confirmation of order by supplier xyz). This information is used to start a more in-depth analysis for the extraction of reference data, products, the date of the letter and so on. For more information on DAU, please refer to (Baumann *et al.* 1997b).

The basic scenario of our integration of DAU into WFMSs fits well into the general architecture presented in Figure 1: The DAU system represents an information provider and a company's mail box represents the collection of documents involved. The task which is requested by a workflow instance represents an information need of this particular instance. In order to support DAU in answering this request, the workflow instance transfers context information to the DAU system. Having finished the task, the DAU system hands over the data requested and consequently satisfies the information need of the workflow instance.

Within *VirtualOffice*, two different kinds of tasks are distinguished: The *process identification* task is stated whenever workflows inform DAU that they are waiting for particular documents. In this case, DAU analyses incoming documents in order to determine the correct workflow instance to which the document has to be assigned. This scenario is discussed in the following section and shown in Figure 3. The second task which we call *information extraction* is stated when workflows request information contained in particular documents. Then, DAU analyses the document in order to deliver the information requested. This scenario is subject to section 2.2 and displayed in Figure 4.

This chapter concludes with a summarization of all design and implementation extensions which have to be accomplished in order to integrate a DAU information provider.

Solving the Task of Process Identification

Process identification, i.e., assigning an incoming document to the corresponding workflow instance, is basically a match of information contained in documents with the corresponding data available in workflow instances. To achieve this, the relevant data in workflow instances have to be collected for two different reasons:

First of all, they build one kind of input for the instance match. Imagine there are a lot of open instances and a new document has to be assigned to them. In this case, the more data we have about the instances and the expected document (e.g, document type, sender, products mentioned and even possible references to other events within the same instance) the more accurate the match can be accomplished. On the other hand, these same data specify the analysis task for DAU since they define which items DAU has to search for.

The resulting scenario is an instantiation of our general scenario and shown in Figure 3. In this scenario, several activities within different workflow instances state the current context of these instances whenever new context is available. All these workflow context data are transferred to the context pool, a database

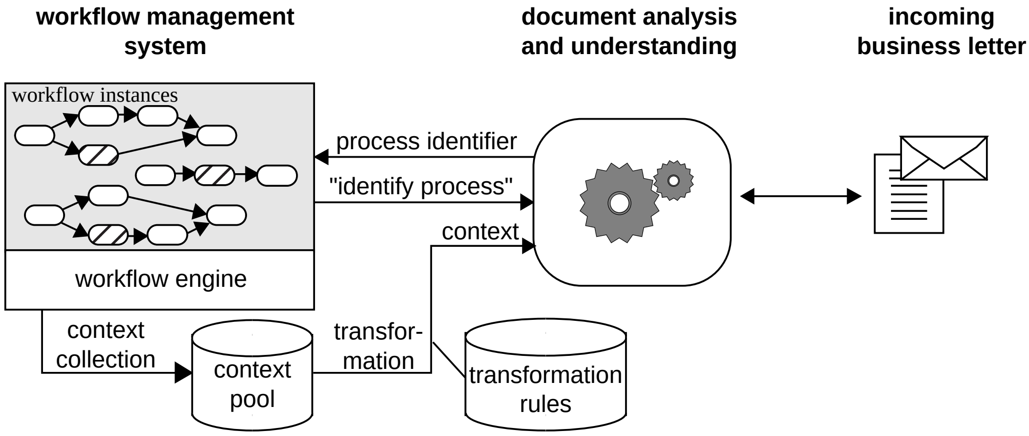

Figure 3: *Information Support by Process Identification for Paper Documents.*

available outside the WfMS.

There, the context is stored until an event occurs in a workflow instance which implies a response by a document (shown by hatched oval activities), i.e., the workflow has an information need which will be satisfied by an incoming document related to the event or workflow. This information need is stated by handing over the request of "process identification" to the DAU system and by formulating its context by means of a so-called *expectation.* It describes content and meaning of the expected document, additional information need such as a list of data which has to be extracted, and some administrative data in order to identify the workflow after a process determination was successful. In order to allow the DAU to interpret the expectation, it is connected to the DAU's document ontology which describes structure and content of the documents of the domain under consideration. Such an expectation is generated by using the context pool as input for a rule interpreter. The rule interpreter uses transformation rules which relate the context information of the context pool to possible contents of an expected document.

Thus, these rules transform the data scheme used in the context pool into the DAU domain ontology. For instance, a rule states that the receiver of the (already known) order is the sender of the expected corresponding invoice.

After generation, the expectation is stored at the information provider's (namely the DAU system's) side in a so-called *expectation set.* Since the expectation set stores all expectations of all processes, it contains the whole workflow context relevant from the DAU point of view. Access to the expectation set is triggered by any incoming document for which DAU has to perform a process identification. Note that the identification of the corresponding process takes all expectations of all workflow instances into account. Consequently, this task can be seen as inherent to the business letter domain. The result of the process identification task is a unique process identifier and, of course, the name of the document which is assigned to the process. The expectation of the correponding instance is deleted after the process identification has been verified within the instance.

This ends the process identification scenario. Data kept within the context pool are deleted when the corresponding process instance has come to an end.

Solving the Task of Information Extraction

Information support by incoming paper documents is not complete by only handing over the document to the corresponding workflow instance. Instead, the relevant information contained in the document should be transferred into an electronic representation. This is formulated by a second task called information extraction. The time for stating such an information extraction task is not restricted to process identification only. On the contrary, a workflow can also state a separate information extraction task after a document has already been assigned.

Figure 4 provides a conceptual view on the information extraction scenario: In this context, only one workflow instance is of further interest. Within this instance, the context pool is still updated after a process identification request whenever new context is available. If an activity states an information extraction request, a new expectation is generated as described above. The request now includes a list of identifiers according to the business letter ontology which have to be extracted from the attached business letter. In contrast to our process identification scenario, the DAU now has not only the raw paper image at its disposal but it can also access previous analysis results for this document image such as OCR and process identification results. This is accomplished by using an adminstrative document index.

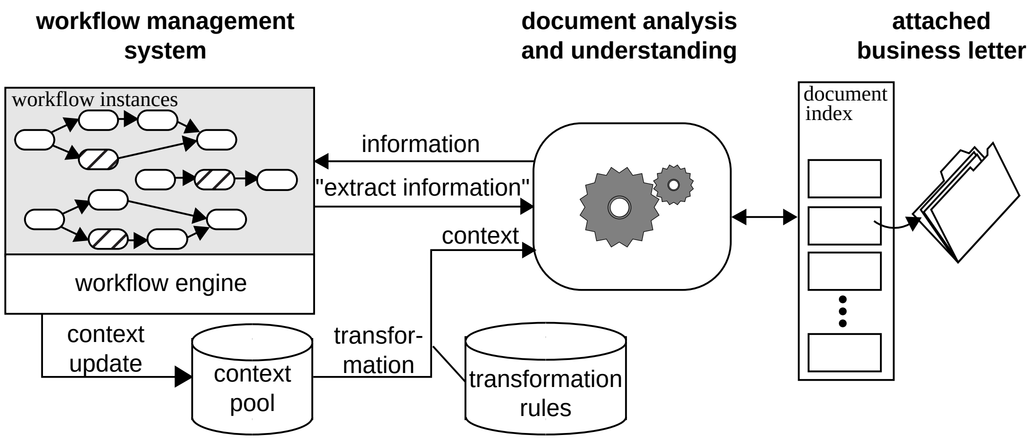

Figure 4: *Information Support by Information Extraction from Paper Documents.*

Implementation Issues

One requirement of the *VirtualOffice* project was to use commercial WfMSs. Therefore, we have to face some conceptual restrictions, e.g., it is impossible to extend the workflow model in order to provide the necessary context information. Thus, our solution uses the interfaces as provided by WfMSs and adds the concept of workflow context data. This allows us to incorporate any WfMS which is based on the WfMC standard for Interfaces 2/3 (Workflow Management Coalition 1998).

In order to extend a conventional workflow to be able to get an automatic information support from paper documents, the following working steps have to be accomplished: At the workflow side, the collection of context has to be modeled at build time of the workflow by including some actions which retrieve workflow context data and store them in the context pool. Furthermore, the workflow has to be extended with activities which state expectations and handle documents after their assignment. At the information provider's side, the ontology for the specification of DAU tasks and the domain ontology for documents have to be made accessible. Furthermore, DAU techniques need an interface to use information stored in the expectation set. Between the two systems, a set of transformation rules has to be specified.

The KnowMore Scenario: Active Information Support by Workflow Integration

The *VirtualOffice* scenario emphasized the use of workflow-context information for improved document analysis and information extraction. Our second scenario has a different focus: The *KnowMore* approach targets on supporting a person working on some knowledge-intensive task by *actively*, i.e., without an explicit, detailed request by the employee, delivering context-sensitive and relevant information.

The KnowMore Scenario in an Example

To illustrate this approach we consider a snippet from a rather simple process: managing a contact to a potential customer at our research institute. After the initial contact (e.g., a telephone call), the relevant topics of interest are identified (for instance, specific technologies or tools potentially useful for the customer, or former projects dealing with similar issues as the customer's problems) and appropriate information material is selected (for example, a technology whitepaper, a brochure about a specific tool, or a project flyer). Having done this, a nice information package can be sent to the potential customer, whose reaction will then determine the further steps (e.g., arranging a meeting, or terminating the process due to lack of interest). Most of these activities can be considered knowledge-intensive. It is easy to imagine useful support for these activities: When selecting the information material to be sent active suggestions from the system would be helpful, supposed that the system takes into account the information from the activities done so far, e.g., the selected topics. An automatic, context-aware archiving of the results is useful when a similar process is started at a different time and / or location: The initial contact will then profit from information about earlier contacts to the same company or about similar cases.

Figure 5 shows a screenshot of our experimental system prototype. On the left, in the background, we see an editor window of the workflow application used to indicate relevant information material. To do this it is necessary to fill the text fields in the input mask. The *KnowMore* system provides support in the following way:

When the workflow engine starts this activity, the system takes the information needs associated to the activity and finds out whether some element of the OM is relevant to this task (i.e., whether there is some material which is relevant wrt. the topics identified). This suggested decision value is inserted in the user input mask offering a proposed solution (in the example, the suggestion comprises documents about the DFKI, a paper about corporate memories, and some material on the ESB system which is an example of an organizational memory). The suggested information elements are ordered according to their relevance computed by the retrieval function as well as to a predefined order based upon their information type. The information elements are then offered to the user as hyperlinks in the *KnowMore* information browser (we use standard web browsers with Java applets for the user interface).

The user is free to accept or dismiss the suggestions or to select different material according to personal knowledge. Whatever the choice, the system keeps track of the solutions and the workflow, and records the results automatically together with the relevant context information. If, sometime later, further material is to be selected, the system remembers the results of the earlier steps and modifies the suggestions accordingly. Thus the system actively offers supply to the user by providing context-specific relevant information.

Solving the Task of Active and Context-Aware Information Supply

In order to provide the services described, the *KnowMore* approach combines an extended workflow model with sophisticated information agents. As illustrated in Figure 6, the workflow actively asks to retrieve relevant documents. To this end, specific modeling activities are needed at process definition time, which facilitate the adequate enactment at runtime.

Process Definition Time.

- **Model business processes.** Model the overall business process with a conventional BPM or workflow tool.
- **Extended modeling of knowledge-intensive tasks**. The knowledge-intensive tasks (KIT) among the activities of a particular process require special attention. In order to facilitate the intended active information support, the KIT representation — illustrated in Figure 6 as a hatched square — extends the conventional description of a workflow activity by a support specification describing the respective information needs as generic queries or query schemes, together with the information agent responsible for their processing. At runtime, the agent's processing of the instantiated queries will deliver relevant information which help to execute the activity in question.
- **Attach extended information flow.** In order to instantiate the query schemata at runtime, thus exploiting situation-specific knowledge and context parameters, the retrieval process must have access to the workflow parameters. These parameters are represented in variables which are handled by the workflow environment. As they go typically beyond what is modeled in a conventional workflow specification we talk about KIT variables, although there is no conceptual difference between them and conventional workflow variables. They simply describe the information flow between tasks in the workflow, and are the communication channel between workflow and information retrieval agents. In Figure 6, hatched triangles indicate addional KIT variables.

  To enable the necessary reasoning for intelligent retrieval, the KIT variables must be embedded into a domain ontology (essentially, this means that their values must be of a type defined as an ontology concept).

  In summary, the KIT variables partially model the data flow in the process and represent the relevant context of the knowledge-intensive activities at runtime.
- **Model information sources.** The information which is actively offered to support particular tasks is taken from a variety of knowledge sources available in an organizational memory. These sources are of different nature, resulting in different structures, access methods, and contents. To enable the precise-content retrieval from the heterogeneous sources, a powerful representation scheme for uniform knowledge description is needed. To this end, structure and metadata, information content and information context are modeled on the basis of formal ontologies. The *document index* in Figure 6 is thus realized as a set of descriptions modeling the information sources and facilitating an ontology-based access and retrieval.

Process Enactment Time.

- **Enact knowledge service processes.** Whenever a knowledge-intensive activity is reached during workflow enactment, the workflow engine not only starts this activity, but in addition initiates an appropriate information agent (which is specified in the KIT description).

  The information agent performs the actual retrieval of relevant information from the information sources. It relies on domain knowledge – available in a domain ontology – to realize an extended, ontology-based information retrieval, and utilizes the context information from the ongoing workflow — found in the instantiated KIT variables — in order to determine relevant information.

  Traversing the formal ontologies according to specified search heuristics (Liao *et al.* 1999), the information agent is able to extend and refine the given queries and to reason about the relevance of available information items.
- **Realize context-aware information storage.** Whenever a workflow activity results in the creation of an information item worth preserving, an appropriate information agent takes into account the current context of the process — available via KIT variables and access to the workflow control data. Thus the information is automatically linked to its creation

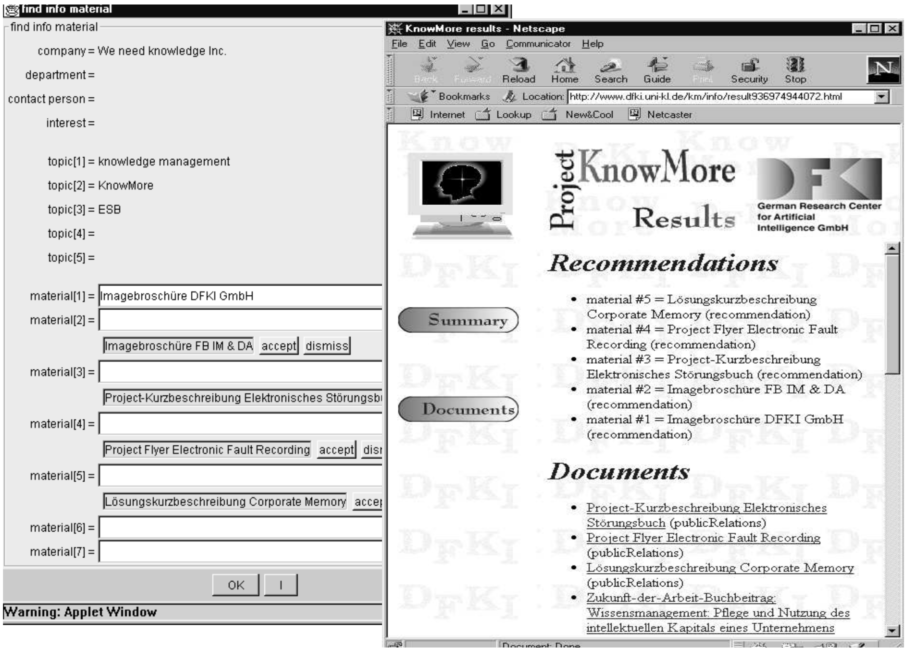

Figure 5: *Active Support by Suggesting Relevant Information Material.*

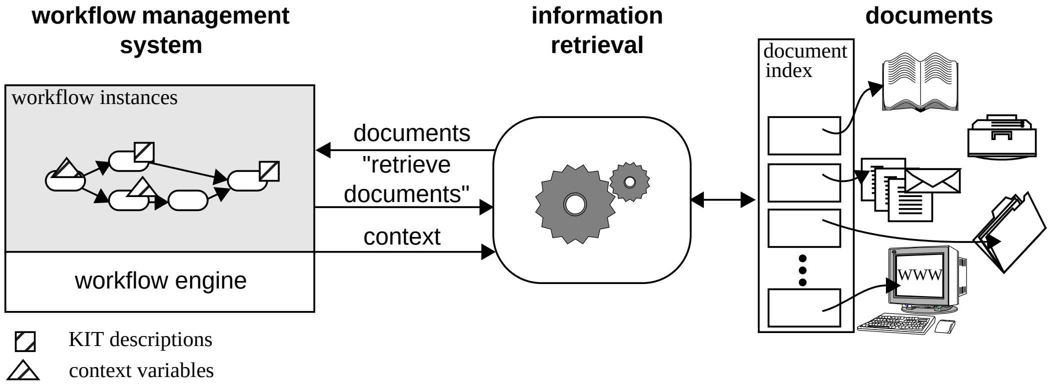

Figure 6: *Information Support by Information Retrieval for Knowledge Intensive Tasks.*

context and can be retrieved accordingly, should the need arise.

Implementation Issues

Figure 7 illustrates that the realization of the *KnowMore* system fits well into the standard architecture of a WfMS as promoted by the workflow management coalition (Workflow Management Coalition 1999): The KIT descriptions are an extension of the workflow relevant data handled by a WfMC-conform workflow engine. The worklist handlers start the relevant activities and trigger in addition the information agents specified in the KIT descriptions. These agents then access the context information in the extended workflow relevant data, evaluate parameters and search heuristics to find relevant information and offer the results to the user. In spite of the fact that the KnowMore prototype uses a self-developed workflow engine, the results are thus compatible with a vast amount of existing and commercially available workflow implementations.

Related Work

Document analysis and understanding has reached a point where the resulting systems are successfully put on the market. But these DAU systems are mostly standalone applications with only one task, e.g., extracting all data from a form. Furthermore, the context used is static, e.g., given by predefined keywords. Up to now, the use of dynamic context as provided by business processes is still a unique feature of the *VirtualOffice* project (Wenzel 1998). Although commercial solutions provide an integration of document analysis and imaging into WfMSs (e.g., FormsRec AIDA solu-

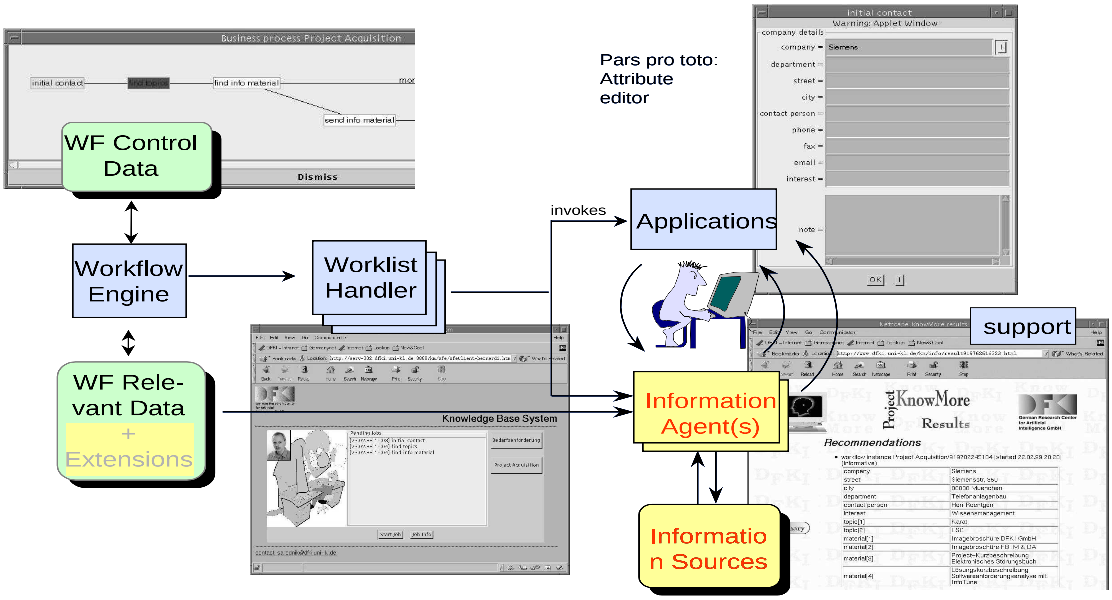

Figure 7: The KnowMore extensions fit well into the WfMC standard workflow architecture.

tion from ICR[1]), they only offer an isolated analysis of documents, i.e., they do not consider any open workflow instances waiting for a particular document. For example, the COI IntelliDoc solution[2] classifies documents according to a set of predefined keywords. Afterwards, the documents are routed to corresponding WfMS worklists to which the keywords are assigned.

Concerning Organizational Memory based support for running workflows, there are two interesting related projects. (Staab and Schnurr 1999) present a project conducted at the University of Karlsruhe which is very close in spirit to the *KnowMore* approach. Compared to our system, they explore in more depth the inferential power of ontology-based retrieval on top of the Ontobroker software (Fensel *et al.* 1998). They introduce the notion of context-based views for coupling workflow and retrieval which is the analogue to our information needs. The approach presented seems to be more developed in some respect than ours, for instance concerning XML-based document representations, or a deeper conceptual integration of workflow and service process. However, they make strong assumptions about the underlying ontologically annotated Intranet information sources, and it seems not completely clear how far their implementation is already. Another very interesting approach is the one by (Kaathoven *et al.* 1999) who describe the mechanism of knowledge flow chunks for synchronizing a conventional WfMS scheduling activities with a concurrently running OM-based task support system assisting in enacting specific activities. Their general considerations about conventional WfMS machinery and external knowledge services complement our interfacing ideas. But, they also seem to be still at the conceptual level of their work.

[1] `http://www.icr.de`

[2] `http://www.coi.de`

Conclusion

We demonstrated that the combination of business processes and their enactment in WfMSs with the information space of an organizational memory can lead to effective user support. Two different approaches have been presented to illustrate this claim:

The *VirtualOffice* project basically aims at improving the company's information logistics by a smoother transfer of information contained in paper documents into the electronic workflows. Thus its service can be applied virtually throughout the whole process and in each process which has interfaces to the external environment of the company where media breaks occur. It supports the business process object level where paper documents are an integral element.

The *KnowMore* project is not so much focused on the object level optimization of whole, and arbitrary workflows. Instead, it concentrates on knowledge-intensive core processes of a business, and knowledge-intensive tasks at the heart of those processes. Not the processes themselves are improved, but a meta-level is added, an additional knowledge-service process which helps the user working on her tasks.

Both examples emphasize the use of the business process models and their enactment in the WfMS as indispensable source of context information. Looking in some detail at this common goal, we can extract the following observations:

Successful coupling of workflow and knowledge-intensive application requires some formal semantics as the basis of communication. That is, we need a mapping from the context information modeled in the workflow to the ontology of the information provider. The examples illustrate two principal approaches to this end:

- When the workflow modeling is extended especially with respect to the support of knowledge-intensive activities, it might be possible to consider the available ontology already at process modeling time. The resulting KIT representation is then well-grounded in the ontology of the information provider.
- On the other hand, if the already-defined workflow is to be left untouched, it is possible to add the necessary semantics by defining transformation rules which map the terms already used in the workflow model to the ontology of the information provider.

To effectively access the workflow context information it seems useful to treat the workflow instances as first order citizens in the information world. Two approaches have been presented to realize this goal:

- The explicit modeling of context information in KIT variables extends the data flow model in the workflow. This is suitable if the workflow engine controls when and what is transferred outside.
- The handling of a separate context pool is suitable to deal with situations where a large amount of workflow instances provide volatile context information which is required by activities outside of the workflow's influence. The continuous collection of context information and its organization according to the needs of the supporting application (e.g. the DAU component) even allows to handle context information which the WfMS only creates temporarily and which otherwise might not be available when needed.

In summary, the extension of the workflow scenario is a suitable way to provide active and extended services and to transform available information into process-oriented actionable knowledge. Some of the sketched solutions seem still a little bit clumsy. However, our observations might indicate some further work worth investigating which might result in suitable adaptations of future WfMS in order to ease the utilization of workflow context information.

References

A. Abecker, A. Bernardi, K. Hinkelmann, O. Kühn, and M. Sintek. Toward a technology for organizational memories. *IEEE Intelligent Systems*, May/June 1998.

G. Alonso, D. Agrawal, A. El Abbadi, and C. Mohan. Functionality and limitations of current workflow management systems. *IEEE Expert*, 1997.

S. Baumann, M. Malburg, H. Meyer auf'm Hofe, and C. Wenzel. From paper to a corporate memory : A first step. In *KI-97 Workshop on Knowledge-Based Systems for Knowledge Management in Enterprises, Freiburg, Germany*, 1997.

Stephan Baumann, Majdi Ben Hadj Ali, Andreas Dengel, Thorsten Jäger, Michael Malburg, Achim Weigel, and Claudia Wenzel. Message extraction from printed documents - a complete solution. In *Fourth International Conference on Document Analysis and Recognition (ICDAR 97), Ulm, Germany, August 18-20*, 1997.

D. Fensel, St. Decker, M. Erdmann, and R. Studer. Ontobroker: The very high idea. In *Proc. 11th Int. Florida AI Research Symposium (FLAIRS-98)*, May 1998.

D. Georgakopoulos, M. Hornick, and A. Sheth. An overview of workflow management: From process modeling to workflow automation infrastructure. In *Distributed Databases 3*. Kluwer Academic Publishers, Boston, 1995.

Rob van Kaathoven, Manfred A. Jeusfeld, Martin Staudt, and Ulrich Reimer. Organizational memory supported workflow management. In August-Wilhelm Scheer and Markus Nüttgens, editors, *Electronic Business Engineering – 4. Internationale Tagung Wirtschaftsinformatik, Saarbrücken, Germany.* Physica Verlag, Heidelberg, 1999.

Minghong Liao, Knut Hinkelmann, Andreas Abecker, and Michael Sintek. A competence knowledge base system for the organizational memory. In *XPS-99 - 5. Deutsche Tagung Wissensbasierte Systeme, Würzburg, Germany*, number LNAI 1570 in Lecture Notes in Artificial Intelligence, Berlin, Heidelberg, New York, March 1999. Springer-Verlag.

Steffen Staab and Hans-Peter Schnurr. Knowledge and business processes: Approaching an integration. In *OM '99 - Proceedings of the international Workshop on Knowledge Management and Organizational Memory (IJCAI-99). Stockholm, Sweden*, 1999.

C. Wenzel. Integrating information extraction into workflow management systems. In *Natural Language and Information Systems Workshop (NLIS/DEXA 98), Vienna, Austria*, 1998.

Workflow Management Coalition. Workflow Client Application Application Programming Interface (Interface 2 & 3) Specification. WFMC-TC-1009, Version 2.0, 1998.

Workflow Management Coalition. Terminology and Glossary. WFMC-TC-1011, Version 3.0, 1999.