

Towards Collaborative Agent-Based Knowledge Support for Time-Critical and Business-Critical Processes

Andrea Freßmann¹, Rainer Maximini¹, and Thomas Sauer²

¹ University of Trier,
Department of Business Information Systems II,
54286 Trier, Germany

{rainer.maximini, andrea.fressmann}@wi2.uni-trier.de

² rjm Business Solutions GmbH, 68623 Lampertheim, Germany
t_sauer@rjm.de

Abstract. In this paper, we present the Collaborative Agent-based Knowledge Engine approach for supporting mobile workers performing time-critical and/or business-critical tasks within agile projects. By a combination of sophisticated knowledge management and a light-weight workflow model, this approach provides guidance and knowledge as required to perform the individual activities. Moreover, we discuss aspects for maintaining project history, as well as possibilities for integrating tools regarding computer-supported collaborative work already deployed in organizations.

1 Introduction

In recent years, agile development methodologies have been the subject of intense research [1]. Although agile methods have been put in opposition to more traditional approaches on the planning spectrum [2], agile projects still demand effective planning skills. As explained in [3], planning is required for arranging tasks effectively, i.e. in order to ensure that the agents involved are carrying out the most important tasks. Planning fulfills a documentation purpose as well, e.g. in order to make progress available to the project stakeholders.

However, planning and documentation is not sufficient to support projects including mission-critical or literally life-critical processes. While planning helps to identify *when* to perform these processes, agents working in knowledge-intense domains will require additional support for leveraging the tacit knowledge available in the organization in order to find out *how* to perform these processes effectively.

In addition, the lessons learned are valuable sources of information for future knowledge intense and creative tasks, which clearly includes management activities like designing an initial release plan. As pointed out in [4], this idea of iterative enhancement of the overall methodology is shared among agile approaches and more traditional planning methodologies found in domains like software engineering.

In this paper, we present the Collaborative Agent-based Knowledge Engine (CAKE) approach, which aims at supporting empirical processes [5], i.e. processes that are

mostly unpredictable and unrepeatable. The CAKE concept provides an infrastructure for constant measurement and control through intelligent and light-weight workflow modeling, leading to the idea of planning sketched above. Furthermore, knowledge-intensive tasks are supported by sophisticated knowledge management, which allows to present context-dependent information to agents carrying out unknown or unexpectedly difficult tasks.

In the following section, we will discuss background architecture for CAKE, and the CAKE approach is presented in brief. In Section 4, a use case demonstrates how to put this approach into practice. Finally, related and adjacent research will be discussed in Section 5.

2 Business and Time Critical Processes in the AMIRA Context

The presented approach of CAKE is developed domain independently but is motivated by the fire service domain within the AMIRA (Advanced Multi-modal Intelligence for Remote Assistance) project¹. This domain addresses both business critical and time-critical situations for mobile workers. While wearing operational kits or gloves, accessing information written on paper or stored on laptops is very cumbersome. Hence, the envisaged mobile workers wear head-sets to access diagnoses support by speech. The fire services demand highly flexible processes and collaborative working in the field. Different representative processes are worked out: First, operatives encounter rare problems and want to perform questions to a system or experts. Second, they need pro-active information support for optimizing their collaboratively working procedures. Third, they require support in report activities and review procedures.

2.1 Single Person Request While Collaboratively Working

A fire fighter extinguishes a fire in collaboration with colleagues and encounters a cylinder with unknown abbreviations of chemicals. While collaboratively working he sends a request to the system for getting information about which chemicals are in the cylinders. The response is only sent to the single mobile worker and the headquarters. The others could get this information from the headquarters if necessary.

2.2 Pro-active Context-Based Information Support

In time critical situations it is a demand for headquarters becoming aware of the activities of officers or fire fighters who work under the headquarters' control for making correct diagnoses. Hence, all interactions of the fire fighters with the system are monitored and logged by the system. Based on these logs context-based information is extracted, so that the headquarters are supported in getting corresponding guidelines, important information, and possible instructions for the mobile workers.

¹ Funded by the EU. Project partners are Kaidara Software, Fast Datasearch, DaimlerChrysler RIC, the University of Trier, and the Fire Service College.

2.3 Collaborative A-Posteriori Analysis

Collaborative a-posteriori analysis of the operations should be managed. This encompasses pro-actively asking the involved persons, headquarters and/or fire fighters, for information about their last actions concerning possible modifications to guidelines or other information used. Furthermore, it is possible to support methods for capturing information about the incident itself, e.g. in order to alleviate handover procedures. For achieving reliable information sources new or additional information is integrated into the databases.

3 The CAKE Approach

Motivated by the AMIRA context a concept of the collaborative working system CAKE is developed that acts as moderator between several services (e.g. search engines) and user interfaces for providing their communication and context-based information support. For example, information may be retrieved from other search services, and the user interface contains a speech service for converting speech-based requests into machine-processable requests.

For coping with knowledge intensive tasks required for context-based information support CAKE comprises a workflow engine manager. Furthermore, an agent framework enables arbitrary access and communication to different agent-based services mediated by CAKE. At last, for providing highly flexible tasks of mobile workers the collaborative working system integrates a planning component for modifying workflows at runtime.

3.1 Workflow Engine Manager

The CAKE approach describes collaboration using *workflow definitions*. Each workflow definition consists of a set of tasks, as well as a control flow relationship between them. The latter allows arranging the tasks in sequence, in parallel, or by using splits and branches, but does not cover data flow at all. This allows to reuse tasks in different application scenarios easily.

A *task* is either a descriptor for a complex activity (e.g. a real-world activity like “write report”) or a machine-executable program, which is denoted as *executor task*. The latter may be implemented as a Java class that can be incorporated into the workflow definition. For instance, an executor task for “send notification” may be defined for use within a workflow definition. Triggering a CAKE workflow definition is also covered by the task definition, so hierarchical decomposition of a complex activity description like “design component” is achieved by following the way a human agent would solve this instead of following a fixed process model.

In order to enact a workflow definition, an *agent* or an *agent role* has to be assigned. Agents may describe either *user agents* (i.e. human actors) or *information agents*, with the latter being connections to arbitrary information providers. Depending on the service provided by agents, information may be accessible read only or with write permissions: While an Internet search engine may be queried as an read-only source, a groupware calendar application deployed in an organization would be available for writing

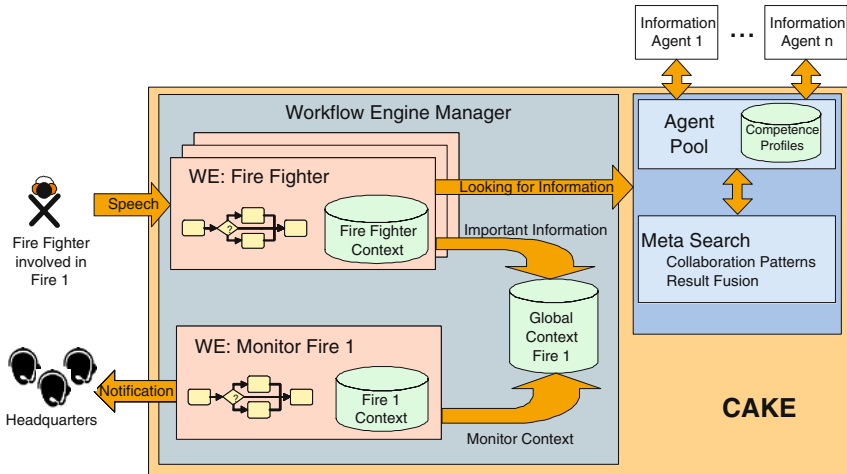


Fig. 1. The CAKE Approach

operations, too. Finally, agent roles describe the competences an agent has to possess in order to follow the workflow definition by providing appropriate agent characterizations.

At runtime, the *Workflow Engine* (WE) initializes an instance of a suitable workflow definition (in the following, *workflow instances* are shortly denoted as *workflow*). Separating these levels enables modifications on the workflow instances without changing the underlying workflow definitions. Beyond controlling tasks the WE contains the local *context* assigned to a single workflow instance that facilitates capturing, storing, and changing of context-based data. Consequently, the context is an information container for any kinds of objects used by the workflow. The context comprises *administrative data*, *workflow control data*, *workflow relevant data*, and *application data* that have been defined by Maus [6]. Due to the possibility of nested (sub)workflows, local contexts can be nested as well by following the concept of inheritance. Assigned to the higher-level workflow engine manager itself only one global context exists that is accessible by all workflow instances under control.

3.2 Agents Framework

From a more technical point of view, agents are represented as a combination of a *technology component*, a *competence profile*, and a *wrapper*. The technology component enables the service provided by the agent (e.g. speech recognition, search, information delivery), while the competence profile includes characterizations about the agent's competencies which are used for the agent role concept. The wrapper makes sure that communication to CAKE is based on a *unified data model*. Hence, wrappers act as interfaces between agents and CAKE for converting data. For being manageable, agents are able to register in the *agent pool* by publishing their competence profiles. Due to dynamical registrations the agent pool works highly flexible in allocating agents as pic-

tured in Figure 1. For instance, based on these competence profiles the most suitable information agent may be found for providing answers to requests performed by user agents.

Beyond simple mediations among agents, strategies based on best practice structures the schedule of agent communication. This information is stored in *collaboration patterns* which specify what to do when the agent firstly contacted is not able to support the user agent. By incorporating these collaboration patterns CAKE aims at providing representations of collaboration strategies among information agents and it aims at preserving universality with regard to domain applications. Collaboration patterns are represented using workflow definitions covering tasks like “sent request to CBR agent” or “sent request to both CBR agent and search agent”. These patterns base on both generic search strategies and domain-specific knowledge. Thus, they allow to leverage otherwise separate domain-specific knowledge within the CAKE approach. Sending parallel requests to several agents leads to a kind of meta search because of performing requests on different data sources. Consequently, further tasks are necessary within collaboration patterns that organize parallel requests to different information agents and the respective result fusions. For finding the most suitable collaboration pattern regarding the current request CAKE provides search facilities on workflow definition representations of the collaboration patterns. Hidden from the end users the search for suitable collaboration patterns is executed.

Furthermore, in order to support a dynamic agent pool the collaboration patterns only contain roles used for agent allocation. Hence, CAKE supports search facilities on agent competence profiles. According to the roles described in the collaboration patterns the most suitable agent is retrieved.

3.3 Pro-active and Context-Based Information Support

By combining the workflow engine manager and the agent framework, CAKE enables pro-active and context-based information support of user agents. The agent technology facilitates the integration of different data sources like external retrieval services or search engines, while the workflow technology enables coordination and collaboration among agents and allows to provide context-based information support.

For realizing the context-based information support, a workflow instance is assigned to one user agent, and the WE monitors all interactions (e.g. requests) of the user agent with this particular instance, whereby more than one user agents can be logged in association to one incident. All collected information is captured by the common parent workflow stored in attribute-value-representations, this information is stored in the global context database of the current incident. Based on the monitored interactions the context-based data already collected is enriched, which ultimately allows the WE to build the application data within the context that can be denoted as repository for all requests, responses, and inputs of the user agents.

The collected application data can be used for analyzing what is going on in the fire ground. CAKE retrieves additional information based on the application data by sending the collected data as request to the search engines, particularly to the search engines that are working on guidelines and working instructions. Consequently, CAKE gets results that match to the application data and contain guideline or working instruc-

tions that can be useful for users working in the fire ground. These guidelines can be presented to the headquarters.

The underlying model for building application data is a domain specific ontology. According to this ontology application data can be semantically interpreted and enables retrievals for context-based information, e.g. by using synonyms. A crucial issue is to develop a quality threshold when having achieved enough context-based data for performing the retrieval. Nevertheless, user agents can be pro-actively supported by context-based information. Otherwise, user agents can get notifications about context-based information of other user agents as shown in Figure 1.

3.4 CBR-Driven Planning Support

The CAKE workflow definitions can be used for planning activities which are expected to be performed during project enactment. Each workflow definition describes an individual activity, which may either be refined by introducing tasks (including sub-workflows) and a control flow among them, or which may be defined abstractly in terms of a “black box”. This allows coarse-grained planning as required by many application domains in order to represent capricious situations. These situations occur in many application domains incorporating creative or knowledge-intense processes which make it impossible to lay out detailed procedures beforehand (e.g. because selecting a concrete procedure depends on context parameters: a fire fighter has to know about the type of fire before an applicable procedure to extinguish it can be chosen.)

As explained above, by separating workflow definitions from workflow instances, CAKE supports late planning by allowing to apply changes to workflows even during their execution. In addition, abstract tasks allow to specify workflow definitions may at any level of detail, which in combination with late planning leads to support of weakly structured workflows [7]. These concepts overcome limitations of “classic” process models and workflow enactment control known from business process modeling, which are unsuitable in agile environments.

During workflow execution, a user agent may do late planning in order to further refine the situation within the current context when approaching an abstractly defined workflow definition. However, in time and business-critical situations this is insufficient: For instance, while extinguishing a fire, a fire fighter cannot wait until late planning has been completed by the headquarters. In order to overcome this, late planning may be backed by previously recorded planning activities, leading to adaptive workflows supported by Case Based Reasoning (CBR) [8].

CBR technology enables a similarity-based retrieval by incorporating further experience: When proceeding to an abstractly planned task the WE allows the corresponding user agent to retrieve a suitable workflow definition in a special workflow database. In that scope, ad-hoc planning is facilitated during runtime. Procedures how to retrieve suitable workflows are described by collaboration patterns, so when looking for potential replacements for the abstract tasks, domain or organization-specific constraints are respected.

In order to reuse existing knowledge within the organization, the CAKE data model is used to characterize semantics of workflow definitions and agents. For workflows definitions, goals and metrics may be defined with respect to the underlying domain

ontology (e.g. “workflow goal is to produce a report”). Goals and metrics are represented using attribute-value pairs based on the unified CAKE data model, which is also used to specify agent characterizations or context data as explained above.

Further modifications on the subworkflow instance can be done by the user agent for adapting. The CBR-driven retrieval mechanism bases on similarity measures derived from the context-based data that annotates the workflow models. In particular, the domain ontology is fundamental for the similarity measure. Though it will be a challenge to develop the similarity measures in detail.

By logging changes to workflow definitions and contexts, the CAKE workflow engine manager allows user agents to conduct further analysis on deviations from a previously laid out workflow definition or additions which have been necessary during workflow enactment. This leads to a Plan-Do-Check-Act cycle, and allows to transform tacit knowledge of the project participants (e.g. experience of a senior fire fighter) into an explicit workflow definition. Notably, domain-specific knowledge is kept aside from the workflow definitions and accessed through the information agents, hence the workflow definitions may be shared across the organization (e.g. between different fire departments). This enables to capture knowledge in the sense of an organizational memory [9].

For business and time critical situations as discussed above, a demand for documentation afterwards is obvious. For example, project stakeholders may request additional reports after having inspected artifacts that have resulted from workflow enactment, or they may request further information on the workflows themselves that led to the results. Thus, the data logged by the workflow engine manager may also be used to automatically generate documentation to handle a-posteriori requests from user agents.

4 Example: A CAKE Use Case

In this section, an example is given how to put the CAKE approach presented above into practice. Especially to fire service domains, the presented approach has to be tailored to the specific domain requirements. For example, the UK fire brigades are strictly organized in a hierarchical order. By following this organizational structure only the incident commander (IC) is equipped by the hands-free information support based on CAKE because he or she is in charge of the incident and of several fire fighters whose number depends on the incident’s size. The responsibilities of the IC comprise information gathering, decision making, and keeping contact to the headquarters, the control center to which the IC has a connection via radio. The fire fighters report to the IC and carry out his or her commands.

A use case is sketched to support the IC and the headquarters in performing a routine operation of extinguishing a fire and keeping safe the incident environments. The incident has happened on an industrial production plant. A global workflow is assigned to the incident in CAKE containing a simple workflow definitions “Fire 1”. For the IC a standard procedure has been prepared beforehand that is now carried out by the incident commander. This procedure is also represented by a simple workflow definition “extinguishing a fire”. Before the incident begins the IC logs into the CAKE system that starts his or her new workflow instance assigned.

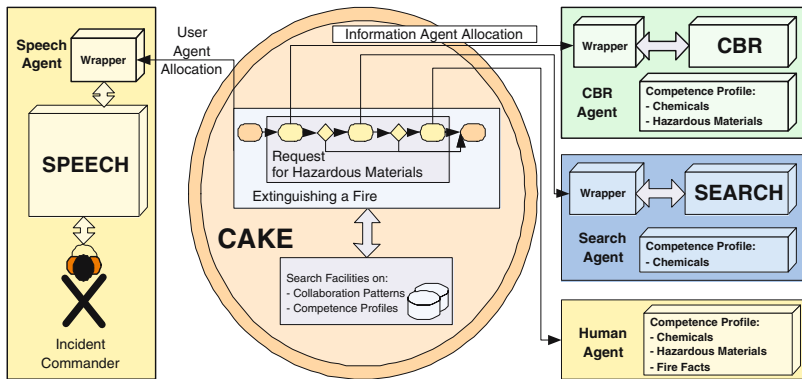


Fig. 2. Illustrations of the Agent Framework tailored to the described Example

In practice, the IC delegates the fire fighters to park all appliances safety, to restrict the zone around the incident, and to begin with extinguishing the fire. Fire fighters collect all available information about the incident and report it to the IC. While working the IC is connected with CAKE using a hands-free voice recognition device that allows the IC to utilize the CAKE information support. When a fire fighter finds a cylinder labeled with an unknown production code, he tells it to the IC who is able to perform a request for detailed information utilizing CAKE: Does this cylinder contain a hazardous material and does it need a special and careful handling?

In order to process the request as illustrated in Figure 2, CAKE enhances the initial workflow instance with the task “request for hazardous materials” and makes use of the respective collaboration patterns. The collaboration pattern says first to ask the CBR-based information agent for seeking on structured chemical data as the agent is an expert for both chemicals and hazardous materials. When no suitable data is retrieved the information agent is asked that works on unstructured chemical data. When no suitable results are retrieved the IC is informed how to connect to a human expert for chemical information. Here, because the CBR-based agent is an expert for chemicals and hazardous materials it delivers the requested information and provides three options in form of short descriptions. The options are sent to the user agent in order to be transformed into natural language, then it is read to the IC. The IC can choose among these results for getting the whole text that contains the information that the found chemical material is Acetylene that has to be dealt in a special manner. Furthermore, the IC can ask for instructions how to deal with the Acetylene cylinder, e.g. first to apply cooling spray from lashed Akron branches on it, to continue cooling for 24 hours, to check if water does not evaporate from cylinder surface when the spray is stopped and to check with a thermal imaging camera does not reveal the presence of any heat. If this is the case, then the cylinders have to be removed into safe area.

Further on the technical level, CAKE has monitored the interactions between IC and CAKE, therefore, the initial workflow is enhanced with this request for chemical materials. Both the request and the response are captured in the context application data of the IC’s workflow instance. CAKE monitors this context for extracting critical

and important information like Acetylene cylinder involved. In this case, CAKE stores this information in the global context that captures all information about the incident. Additional information about the Acetylene cylinders can be retrieved by utilizing the information agent for pro-actively supporting the headquarters with information. Although it is often the case that the headquarters do not know what is going on in the incident ground exactly because of their busy work. Here, the headquarters are informed that Acetylene cylinder have been found and the headquarters can be forewarned of potential hazardous situations, e.g. that in most cases when Acetylene is involved other hazardous materials are involved as well. The headquarters can decide whether they instruct the incident commander for searching for other materials.

Finally, the IC and the fire fighters complete their operation, after that they create the operation report in order to fulfill the end task of the workflow definition. They decide that for future inspections, querying the company profile should be mandatory in order to prepare for giving safety instructions more carefully.

In the next section, we will discuss related work to the CAKE approach. Finally, we will summarize the core ideas presented above, and discuss possible future work.

5 Related Work

Providing workflow management support for agile methodologies has been discussed before [10]. This approach follows the idea of “heavy agile” [11] methods by suggesting to augment the specific methodology of Extreme Programming with additional documentation, formality, and tools, in order to support larger undertakings like distributed-team projects. However, because of limitation to a specific methodology, scope of application is limited.

In order to model workflows, various concepts have been proposed in the past decades, however none gained broad acceptance. Most rely on process description languages focusing on task dependencies. For instance, formal languages like MVP-L [12] have been designed specifically for expressing relationships between the various aspects of a software project, and to provide a formal execution model. Other efforts propose state and activity charts as means of workflow specification and execution [13]. While these efforts have their advantages to detect infeasible or suboptimal configurations, they require complex tool support, because of their rather non-intuitive model representation.

Workflow management systems have been discussed before as a valuable source of information for supporting knowledge-intensive tasks [6]. A promising approach is the concept of weakly-structured workflows that provides knowledge-intensive tasks [7]. Here, benefits of explicit process models, workflow-type control, and information support are investigated for process-oriented Knowledge Management support. Building an organizational memory based on process models and workflow information has been used in various systems, as described in [14]. In [15], the author extends this idea by suggesting a framework for explicitly expressing information needs and sources, and how to support team members by proactively providing them. Modeling knowledge creation and management within a workflow environment focused on weakly-structured workflows has been discussed in [16]. However, distributing knowledge to the user

agents remains a challenge, and none of the approaches discussed above has been designed to cover the special demands of time and business critical-situations as described in section 2.

An approach for adaptive workflow enactment using multiple agent systems [17] makes use of Business Process Execution Language for Web Services (BPEL4WS or short BPEL) [18]. Here, the XML-based language for expressing the composition of Web services is applied for coordinating agents. Finally, using information agents for providing project history information is known from the field of software metrics [19], and deriving documentation by tracing workflow execution has been suggested in [20]. The idea is to look for patterns within a project history in order to create agile documentation [21] as required, and without obliging users to enter additional information.

6 Conclusion

Motivated by mission-critical or literally life-critical domains from the AMIRA project, in this paper we presented the CAKE approach, a concept for coping with business and time-critical processes. Based on requirements derived from the fire service domain, CAKE rejoins several approaches for supporting application-driven scenarios. Besides providing workflow knowledge and agent mediation the CAKE approach also integrates planning skills and agile documentations as well.

CAKE is currently getting implemented for the application domains of roadside assistance and fire services. In future work, we will research how CAKE can be utilized for providing support in other domains, where knowledge-intense and creative processes are present. This includes, but is not limited to, software engineering and medical processes.

Further applications, e.g. for supporting geographically dispersed collaboration between experts, are taken into account as a generalization of the ideas behind the CAKE approach.

Acknowledgement. The authors acknowledge the European Commission for funding AMIRA under grant number FP6, project IST-2003-511740.

References

1. Abrahamson, P., Salo, O., Ronkainen, J., Warsta, J.: Agile software development methods: Review and analysis. Number 478. VTT Publications (2002)
2. Boehm, B.: Get ready for the agile methods, with care. *IEEE Computer* **35** (2002) 64–69
3. Beck, K., Fowler, M.: *Planning Extreme Programming*. Addison-Wesley (2000)
4. Glazer, H.: Dispelling the process myth: Having a process does not mean sacrificing agility or creativity. *Crosstalk: The Journal on Defense Software Engineering* **14** (2001) 27–30
5. Advanced Development Methods, Inc.: *Control chaos: Living on the edge. the origins of scrum*, <http://www.controlchaos.com> (1996)
6. Maus, H.: Workflow context as a means for intelligent information support. *Lecture Notes in Computer Science* **2116** (2001) 261–274

7. van Elst, L., Aschoff, F.R., Bernardi, A., Maus, H., Schwarz, S.: Weakly-structured workflows for knowledge-intensive tasks: An experimental evaluation. In: Proceedings of the 18th International Workshops on Enabling Technologies: Infrastructures for collaborative enterprises, Linz, Austria, IEEE Computer Society Press (2003) 340–345
8. Weber, B., Wild, W., Breu, R.: Cbrflow: Enabling adaptive workflow management through conversational case-based reasoning. In Funk, P., Calero, P.A.G., eds.: Advances in Case-Based Reasoning, Proceedings of 7th European Conference, ECCBR 2004. LNAI3155, Madrid, Spain, Springer Verlag, Berlin-Heidelberg (2004) 434–448
9. Wargitsch, C., Wewers, T., Theisinger, F.: An organizational-memory-based approach for an evolutionary workflow management system - concepts and implementation. In: HICSS '98: Proceedings of the Thirty-First Annual Hawaii International Conference on System Sciences-Volume 1, Washington, DC, USA, IEEE Computer Society (1998) 174
10. Maurer, F., Mertel, S.: Process support for distributed extreme programming teams. In: Proceedings of the International Workshop on Global Software Development (ICSE'02 GSD), Orlando, FL (2002)
11. Highsmith, J.: Agile Software Development Ecosystems. Addison-Wesley (2002)
12. Bröckers, A., Lott, C.M., Rombach, H.D., Verlage, M.: MVP-L language report version 2 (1997)
13. Wodtke, D., Weißenfels, J., Weikum, G., Dittrich, A.K., Muth, P.: The mentor workbench for enterprise-wide workflow management. In Peckham, J., ed.: Proceedings of the ACM SIGMOD International Conference on Management of Data, Tucson, AZ, ACM Press (1997) 576–579
14. Abecker, A., Bernardi, A., Ntioudis, S., Mentzas, G., Herterich, R., Houy, C., Muller, S., Legal, M.: The DECOR toolbox for workflow-embedded organizational memory access. In: Proceedings of the 3rd International Conference on Enterprise Information Systems, ICEIS 2001. (2001) 225–232
15. Holz, H.: An incremental approach to task-specific information delivery in se processes. In: Proceedings of the 18th IEEE International Conference on Automated Software Engineering. (2003) 295–298
16. Papavassiliou, G., Mentzas, G.: Knowledge modelling in weakly-structured business processes. *Journal of Knowledge Management* 7 (2003) 18–33
17. Buhler, P.A., Vidal, J.M.: Towards adaptive workflow enactment using multiagent systems. Volume 6., Springer Science + Business Media, Inc. Manufactured in The Netherlands (2005) 61–87
18. Andrews, T., Curbera, F., Dholakia, H., Goland, Y., Klein, J., Leymann, F., Liu, K., Roller, D., Smith, D., Thatte, S., Trickovic, I., Weerawarana, S.: Business process execution language for web services, version 1.1. specification, BEA Systems, IBM Corp., Microsoft Corp., SAP AG, Siebel Systems (2003)
19. Johnson, P.M., Kou, H., Agustin, J., Chan, C., Moore, C., Miglani, J., Zhen, S., Douane, W.E.: Beyond the personal software process: Metrics collection and analysis for the differently disciplined. In: Proceedings of the 2003 International Conference on Software Engineering, Portland, OR (2003)
20. Sauer, T.: Using design rationales as agile documentation. In: Proceedings of the 18th International Workshops on Enabling Technologies: Infrastructures for collaborative enterprises, Linz, Austria, IEEE Computer Society Press (2003) 326–331
21. Ambler, S.: Agile documentation. Essay, <http://www.agilemodeling.com/essays/agileDocumentation.htm> (2001)