

Lars Ehrler · Martin Fleurke · Maryam Purvis  
Bastin Tony Roy Savarimuthu

## **Agent-based workflow management systems (WfMSs)**

**JBees: a distributed and adaptive WfMS with monitoring and controlling capabilities**

Published online: 8 November 2005  
© Springer-Verlag 2005

**Abstract** Workflow management systems (WfMSs) based on agent technology can cope with the rapidly evolving business environment better than most other systems as they are more flexible and open. In this paper we describe a possible architecture of such a system by means of our prototype WfMS called JBees. The combination of collaborating agents and the Coloured Petri Net (CPN)-formalism in JBees enables a flexible and adaptive system with the possibility of simulation, analysis, and monitoring of the process execution in order to identify potential inconsistencies and to provide appropriate information to the workflow administrator for the purpose of the process improvement.

**Keywords** Workflow management system · Adaptability · Software agents · Workflow monitoring · Coloured Petri nets

---

L. Ehrler · M. Fleurke · M. Purvis (✉) · B. T. R. Savarimuthu  
Department of Information Science, University of Otago,  
PO Box 56, Dunedin, New Zealand  
E-mail: tehrany@infoscience.otago.ac.nz

L. Ehrler  
Fakultät Informatik, Technische Universität, Dresden, Germany  
E-mail: lars@larsehrler.de

M. Fleurke  
Department of Computer Science, University of Twente,  
Enschede, The Netherlands  
E-mail: a.m.fleurke@student.utwente.nl

## 1 Introduction

Workflow management systems (WfMSs) (Meilin et al. 1998; Schael 1998; van der Aalst and van Hee 2002) are increasingly being used to manage business processes associated with distributed global enterprises. Some of the benefits of using a WfMS are:

- Ability to visualize the overall process and interdependencies between various tasks,
- Automation of the processes, and
- Automated coordination and collaboration between various business entities.

Existing, commercially available WfMS do not offer sufficient flexibility for distributed organizations that will be participating in the global market. These systems have rigid, centralised architectures that do not operate across multiple platforms (Shepherdson et al. 1998).

Improvements can be made by employing a distributed network of autonomous software agents that can adapt to changing circumstances.

In the past, WfMSs were used in more well-defined activities, such as manufacturing, where the processes tend to be more established and stable. But in the current climate WfMSs may be used for more fluid business processes, such as e-commerce, or in processes involving human interactions, such as the software development process. In such situations, at times, it is not always possible to predict in advance all the parameters that may be important for the overall processes. In particular some of the reasons for wanting an adaptive WfMS are as follows (Bandinelli et al. 1993):

- It may not be possible to specify all the process details associated with a complex process at the outset. The initial model may represent a high-level view of the process, which includes some of the sub-processes. Gradually some of these sub-processes may be refined as the stakeholders obtain more experience and knowledge of a particular process.
- Due to changes in the market, new requirements may be imposed which can impact the process definition. This change in the market may also include the availability of some new technologies, which may require the modification of the process as well.

In this paper, we describe an agent-based framework that provides a flexible infrastructure for incorporating dynamic changes to the currently executing process model. This architecture provides a mechanism for communication of distributed components in order to support inter-organizational WfMS.

In Sect. 2 we give a short overview of the work which has been done so far. Section 3 describes the architecture and the functionalities of our prototype WfMS, followed by an example in Sect. 4. We conclude with the results of our research and point out the work which should be done in the future in Sect. 5.

## 2 Background

### 2.1 Coloured Petri nets

Coloured Petri nets (CPNs) are used to model workflow systems, due in part to their sound mathematical foundations and to the fact that they have been used extensively for modelling distributed systems with concurrent activities (Jensen 1992). Coloured Petri nets consist of the following basic elements:

- Tokens, which are typed markers with values—in our implementation the type can be any Java class.
- Places (circles), which are typed locations that can contain zero or more tokens.
- Transitions (squares), which represent actions whose occurrence (firing), can change the number, locations, and value of tokens in one or more of the places connected to them. Transitions may have guards, which must evaluate to TRUE in order for a transition to fire. In a workflow model a transition may represent a task.
- Arcs (arrows) connecting places and transitions. An arc can have associated inscriptions, which in our implementation are Java expressions whose evaluation to multisets of token values affects the enablement and firing of transitions.

Some reasons for preferring Petri net modelling to other notations used for workflow modelling are given by (van der Aalst 1996):

- They have formal semantics, which make the execution and simulation of Petri net models unambiguous. It can be shown that Petri nets can be used to model workflow primitives identified by the Workflow Management Coalition (WfMC 1995).
- Unlike some event-based process modelling notations, such as dataflow diagrams, Petri nets can model both states and events.
- There are many analysis techniques associated with Petri nets, which make it possible to identify dangling tasks, deadlocks, and safety issues.

Currently, we are using JFern (Nowostawski 2003)—a CPN simulator and enactment engine to design and execute the models. We have modified the JFern editor to support hierarchical models, which enables the modeller to start with a coarse model of the process and gradually refine each of the sub-processes into a separate Petri net model.

### 2.2 Agent systems

Some commonly accepted characteristics of an agent are listed by Bradshaw (1997): reactivity, autonomy, collaborative behaviour, adaptivity, and mobility. The most important attribute is probably autonomy. Wooldridge (1999) (and similarly Shoham 1997) discusses the difference between agents and objects. They both mention three main distinctions, which are all related to the autonomy of agents: agents have control over their behaviour, agents

have flexible behaviour (they can behave reactively or proactively), and an agent often has its own thread(s) of control.

According to Sycara (1998), there are several benefits of using multiagent systems for building complex software. For example, multiagent systems can offer a high level of encapsulation and abstraction. Because agents are independent, every agent can decide by itself, what is the best strategy for solving its particular problem. The agents can be built by different developers; as long as they understand the communication, they can work together. A second important benefit is, that multiagent systems offer distributed and open platform architecture. Agents can support a dynamically changing system without the necessity of knowing each part in advance. This requires, however a matchmaking infrastructure.

Our system is based on the Java-based agent platform Opal (Purvis et al. 2002), developed at the University of Otago since 2000. It meets the standards of the Foundation for Intelligent Physical Agents (FIPA) for agent platforms and incorporates a modular approach to agent development (Nowostawski et al. 2001).

## 2.3 Former approaches to WfMSs

### 2.3.1 *WfMSs using software agents*

In the context of WfMSs, agent technology has been used in different ways (Joeris 2000). In some cases the agents fulfil particular roles that are required by different tasks in the workflow. In these cases the existing workflow is used to structure the coordination of these agents (Jennings et al. 2000; Nissen 2000). An example of this approach is the work by M. Nissen in designing a set of agents to perform activities associated with the supply chain process in the area of e-commerce (Nissen 2000).

In other cases, the agents have been used as part of the infrastructure associated with the WfMS itself in order to create an agent-enhanced WfMS (Stormer 2001; Wang and Wang 2002). These agents provide an open system with loosely coupled components, which provides more flexibility than the traditional systems.

Some researchers have combined both of these approaches (Chen et al. 2000), where an agent-based WfMS is used in conjunction with specialized agents that provide appropriate application-related services.

### 2.3.2 *Adaptive WfMS*

Adaptive workflows have been discussed for many years and many people have described what should be done (Chen et al. 2000; van der Aalst 2001). Only a few have proposed techniques to manage adaptability and only a small number of actual implementations have been made that tackle some aspects of adaptability (Purvis et al. 2001). Transferring running work cases to a new model is still a difficult issue. The work done in the paper (Purvis et al. 2001) describes a prototype, which provides some adaptability by

manual transfer of tokens in the new process model. This is indicated in a comparison of current WfMS that was done by van der Aalst et al. 2002.

The approach in much of the research is to define a limited set of possible transfers, such as inserting an extra task in sequence, skipping a task, or replacing a task with a new subnet. In this way the semantics of the workflow stay well defined. The drawback of this approach is that the set of possible transfers is very limited and/or the expressive power of the specification language is too restricted, limiting the possibilities for specifying a process. For example the approach of van der Aalst et al. (1999) is based on inheritance, but it requires that every workflow should be derived according to a limited set of transitions from some basic workflow definition.

### 2.3.3 *Monitoring and feedback*

Monitoring and feedback mechanism of workflow systems have been mentioned by researchers for many years. Few researchers have discussed the issues associated with monitoring and feedback (Cui et al. 1998; Muehlen and Rosemann 2000). When it comes to agent-based monitoring, there has been one proposed system (Wang and Wang 2002), but this lacks feedback of the process model using agents and also does not cater to distributed monitoring, which is central to any workflow system as described by van der Aalst et al. (2002).

## 3 **JBees**

### 3.1 The architecture

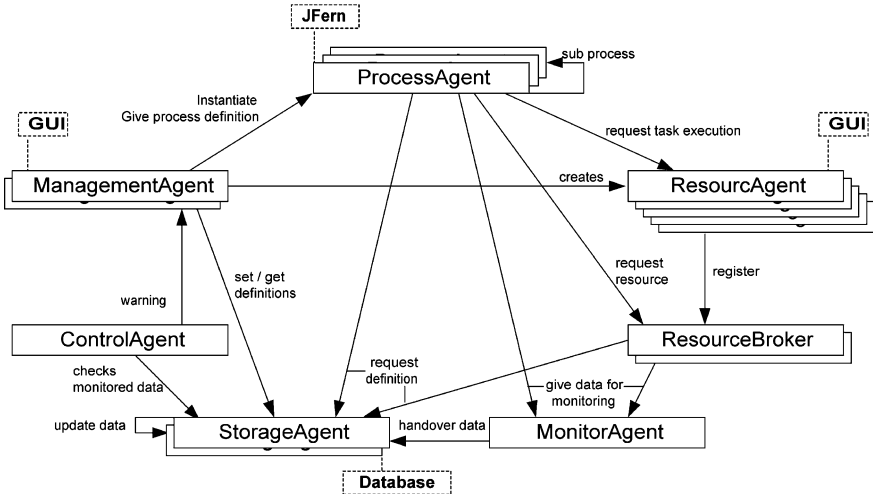
Our research is focussed on developing an agent-enhanced WfMS, where the work associated with running a WfMS has been partitioned among various collaborating agents that are interacting with each other by following standard agent communication protocols.

JBees is based on Opal (Purvis et al. 2002) and uses the CPN execution tool JFern (Nowostawski 2003). A first description of JBees can be found in our previous paper (Fleurke 2003). Our enhanced system consists of seven Opal agents, which provide the functionality to control the workflow. Figure 1 shows these seven agents and their collaboration.

#### 3.1.1 *Management agent*

A management agent (on the left in Fig. 1) provides the user interface for the human workflow manager. It can:

- Create and delete role definitions and process definitions
- Instantiate a new process instance
- Create resource agents for new resources
- Simulate the execution of a process.



**Fig. 1** An overview over the JBees architecture

### 3.1.2 Storage agent

A storage agent (on the bottom in Fig. 1) manages the persistent data, for instance the definitions of tasks, roles and processes, and the monitored data. It also notifies all management agents if the data has changed (for example one management agent adds a definition, and the storage agent notifies all other management agents that there is a new definition).

### 3.1.3 Process agent

A process agent (on the top in Fig. 1) is responsible for the execution of one particular case. For each work case and for each sub work case a new process agent is created. The CPN model is provided by the management agent or by the “parent” process agent and the process agent uses JFerm to execute this model.

### 3.1.4 Resource agent

A resource agent (on the right in Fig. 1) is the user interface for the human resource or the interface for some tool which can do tasks automatically (such as printers and scanners). Every resource has its own resource agent. The agent represents the resource in the system and negotiates the resource allocation on behalf of the resource and exchanges the necessary information for the execution of a task.

### 3.1.5 Resource broker agent

A resource broker (on the right in Fig. 1) is responsible for the resource management. Every resource agent is registered with at least one of the

resource brokers and a process agent requests the resource broker to identify and allocate a suitable resource. Different strategies for the resource management can be incorporated just by replacing the agent with a new agent that employs a new strategy.

### 3.1.6 *Monitor agent*

A monitor agent (on the bottom in Fig. 1) gathers the data in the system that is necessary to analyse workflows, such as execution times and resource utilization. The monitor agent gathers this data associated with a particular case and (after the case is finished) sends this data to the storage agent for persistent storage.

### 3.1.7 *Control agent*

A control agent (on the left in Fig. 1) provides the feedback mechanism required for process re-engineering. The control agent continuously senses the anomalies or violations of the criteria specified by the manager and sends warning messages to the manager agent and also logs those messages. The manager agent or the human manager decides the appropriate corrective measures that have to be carried out.

## 3.2 Advantages of using an agent-enhanced WfMS

We outlined in Sect. 2.2 that the use of software agents facilitates the design of a distributed and open platform because agents are loosely coupled components forming an open system. New technologies and techniques can easily be incorporated by introducing new specialized agents into the system.

The use of such an open platform has advantages for building a workflow management system. Firstly it facilitates cooperation with other organizations. The essential pieces for interoperability are ontologies and interaction protocols. If the organizations agree on standard ontologies and protocols, they can easily cooperate even if they have completely different platforms and systems. Secondly, the use of software agents gives us the opportunity to have a flexible system. For example, the resource management strategy can change dynamically (by just starting an agent with a new strategy), resources can register or de-register, and process definitions can be changed more easily.

This flexibility provides adequate support in implementing an adaptive and distributed system. New process models (when they are required) can be dynamically incorporated into the system (see Sect. 3.5) and each agent can be located in a separate host on the Internet (see Sect. 3.7).

## 3.3 Execution of workflows

In order to process a new work case, the management agent creates a new process agent providing the definition of the process. The process agent starts the execution of this case by invoking JFern.

The process agent requests all the definitions of the tasks from the storage agent (the process definition contains only the names of the task—the definitions themselves are not stored in the CPN model so that they can be easily reused in different process definitions). Every task definition contains all roles that can execute this task. These roles are needed to request an appropriate resource from the resource broker (see Sect. 3.4). The process agent contacts the resource allocated by the resource broker directly to hand over the information necessary for executing the task.

### 3.4 Resource allocation

The workflow is described by a coloured Petri net. Normally, a transition representing a task is enabled if there is a job token that specifies the context of the current job (work case instance), and a resource token that indicates that there is a resource to do the task. The resources are taken from a CPN model place that holds all resources in the system. To provide more flexibility the resource required for each transition is not included in the model. Since the resources are decoupled from the model, they can be easily changed. In this approach resources have to be explicitly requested when required by a transition. This process (obtaining the required resources) is represented in another subnet not shown in the process model to avoid cluttering the net with the modelling elements that are not central to the model. The request for a resource contains the roles which can execute the task. The resource broker sends a request message to the appropriate resource agent which is capable of performing the task. The resource agent can either decline or accept the request. If the broker can successfully allocate a resource, it will return the name of the chosen resource to the process agent.

### 3.5 Adaptability

As indicated earlier, for each work case a new process agent is created and an appropriate CPN model is instantiated. The work case is represented by a token in the CPN model. While the workflow system is running there can be a need to modify the CPN model (due to a change in the business process). In this case there are several possible actions that can take place. The choice of these actions depends on the scope of the change requested and the extent to which it has to be applied to the existing work cases. In case the change has to be applied to new work cases, waiting in the queue to be processed, then we can easily instantiate the process agent with the new or modified model instead of the old model for these work cases. However, if the proposed change in the process should be applied to the running instances, it is necessary to make sure that the change does not violate the structural and semantic consistency of the model before we can transfer the state of the running instance to the new model. To accomplish this task, we use an algorithm that calculates the region in which transfers are unsafe (Fleurke 2004). The algorithm is an improved version of van der Aalst's algorithm (van der Aalst 2001) to calculate the minimal change region.



In order to determine whether a transfer of the case is safe or not, the algorithm identifies the region of the process model that is impacted by the change. The impacted region is first identified by calculating the difference between the old and new net. Second, any other net elements that relate to the affected elements via parallelism are included. Together these elements form the critical region that is examined in order to determine if the transfer is safe or not.

If the job token of the running case is outside the corresponding region of the net, then the state is transferred to the proposed process model. Otherwise the transfer is not safe and requires a human manager to choose one of the following exception handling methods:

- Continue processing the instance according to the old definition
- Restart the instance with the new definition
- Delay the change until the job token gets during the workflow execution in a safe region and then change the process model
- Have the user manually transfer the token

A possible transfer is facilitated by our use of one agent for each work case. Therefore the work cases are executed independent of each other and a transfer from one case does not interfere with other work cases.

For more details concerning the algorithm used to identify this region, refer to the Masters degree thesis by Fleurke (2004).

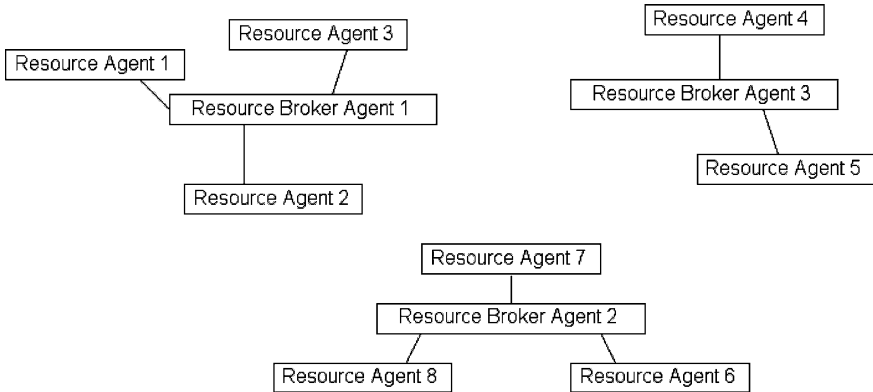
### 3.6 Reliability

One significant problem in the design of a distributed system is reliability. In particular, certain elements that are crucial to the execution of workflows might be prone to “central points of failure” (for example if they cannot be accessed or they fail themselves). A WfMS usually has two critical activities: the management of resources and the management of persistent data. If these elements are centralized and they fail, the whole system cannot function any more.

Our architecture based on agent technology provides an easy and flexible solution for this problem. A resource broker agent is responsible for the management of resources, but there can be several resource brokers. As you can see in Fig. 2, every resource agent is registered with at least one of the broker agents. If a resource broker fails or is not available, there are still other brokers which might be available and one of them can be contacted. Possibly a resource that is not ideal will be allocated in this case, however, the process will still be executed. The management of persistent data is managed in a similar way.

### 3.7 Distribution

An organization with offices/departments at different locations would require distributed workflow models. As shown in Sect. 2.2, the adoption of



**Fig. 2** Distributed resource management

agent technology facilitates the support for these distributed workflow modelling and execution. The mechanism of distribution in JBees has been carried out in the following ways:

1. The process model can be distributed. A model can be separated into several sub-processes which can be located on different hosts as they are executed by several process agents. The sub-processes could span various geographical locations. For example a shipping sub-process can be located in the shipping department while the billing sub-process can be located in the finance department of an organization—both departments may be in different cities or even countries.
2. The resource management can be distributed. Even if the process is executed in one location, resource brokers at several locations can be asked for a resource. Therefore the execution of a task can be delegated to a remote resource with possibly specialized skills, even though the process itself is located at a completely different location.
3. JBees incorporates an agent-based distributed database. A central database can become a bottleneck or, worse, a central point of failure in a system. Therefore we incorporated a distributed database, managed by agents. The system can have several storage agents—each agent with its own local database (a node of the distributed database) underneath. Data can be retrieved/stored by all agents through any of the storage agents. When any of the storage agents senses the change of data, it synchronizes the distributed database by informing all other storage agents to update the data.

### 3.8 Simulation

In the process of creating new process models these models need to be tested and evaluated for their efficiency. For this purpose JBees has a special simulation agent (a dedicated process agent). The human manager can

choose several parameters such as number of resources, number of cases or delay between actions, and the agent simulates the execution of the process. The data collected by the simulation agent can be analysed and also used for improving the process model.

### 3.9 Monitoring

Monitoring is an indispensable part of any WfMS. Every case that is executed in a process model has to be monitored for its various properties such as time taken to complete the process, the various resources employed, time taken by the resources to complete the tasks, and waiting times of the jobs in the queue to be served by a resource. Our initial work on monitoring can be found in the previously published paper (Savarimuthu et al. 2004).

Our architecture incorporates the modules to examine, analyse, and display the properties of the workflow system. This data is stored after simulation (as described in Sect. 3.8) and also after enactment of every case. Due to the distributed nature of the system we need new strategies for monitoring the system in enactment mode. The information is distributed and work cases can have modified process definitions, several sub-work cases, and several agents associated with the case, which are independently working. Our approach is to have a dedicated monitor agent collecting all information. Information which belongs to the same case will be collected and integrated by this agent. This data can be used to improve the workflow while executing (such as employing more resources of a role that is in demand).

In both cases (simulation and enactment), the data is sent to the storage agent to store this data persistently. The human manager of the WfMS can choose any of the properties to analyse the state of that property at any particular point of time. The data obtained from simulation and enactment can be used for drawing graphs for the purpose of monitoring. We have integrated the JFreeChart, an open source Java API (Object Refinery Ltd JFreeChart. <http://www.jfree.org>) with our framework to draw and display graphs using the data collected by our system.

### 3.10 Feedback mechanism

During the enactment and also while simulating various scenarios, the user/manager of the workflow system specifies certain criteria that have to be looked into continuously and should be constantly compared with the data obtained from the simulation or enactment. The user could specify whether the criteria holds for specific processes or for all processes. The criteria used in the feedback mechanism are implemented using database queries. The manager constructs these queries depending on which criteria he/she needs to monitor and also the frequency at which the control agent should report violations of these criteria.

When the controlling agent senses anomalies or violations of these criteria it sends a warning message to the management agent. An example of these

criteria could be the overall completion time for any given case not exceeding more than a given value or the resource utilization not decreasing below a certain value. The controlling agent also logs the messages to the warning log so that the human manager can handle that particular problem if he/she is not available at that time. The aim of this design is to capture all possible critical conditions for continuous optimization of the processes.

## 4 Example

### 4.1 Simple example

To illustrate the functionality of the system, we chose a simple process of ordering a book. The process model of this is shown in Fig. 3.

After the customer orders a book, the inventory has to be checked whether a copy of this book is there and the credit rating of the customer has to be checked. These checks are done in parallel to speed up the process and afterwards the results are evaluated. Based on the evaluation it is decided whether the order can be processed or should be rejected. Assuming that the processing of the request has been approved, the shipping of the book and the sending of the bill are done in parallel. Finally the results of shipping and billing activities are archived so that one can handle possible customer complaints.

Suppose a new order arrives. The first step of creating a new process agent and getting all task definitions is shown in Fig. 4.

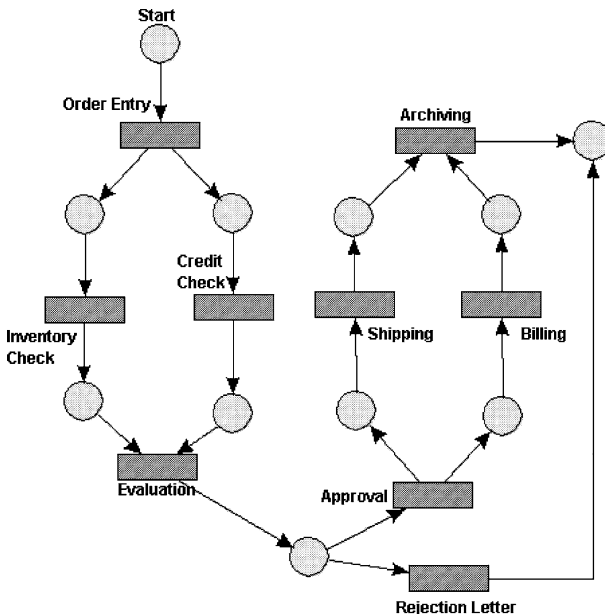


Fig. 3 The CPN of the example workflow




Fig. 4 The sequence diagram for starting a new work case

The process agent “process1” starts executing the work case by putting a job token in the place called Start (shown in Fig. 3). This activates the transition Order Entry.

According to the task definition, a resource of the role processor is needed. The process agent asks the resource broker for a resource of this role. The broker will allocate a resource of this type and return the name of this resource to the process agent. In case it cannot allocate a resource it will send a message to the process agent that it has failed to allocate a resource. Then, the process agent has to decide how to handle this scenario (for instance by asking another resource broker, or waiting a certain time and asking the same resource broker again).

After the resource executes the task and returns the result of the task, the JFern engine gets notified that the task Order Entry has been successfully executed and it continues executing the CPN model. The process of allocating a resource and executing a task is shown in Fig. 5 and this process is repeated for every task.

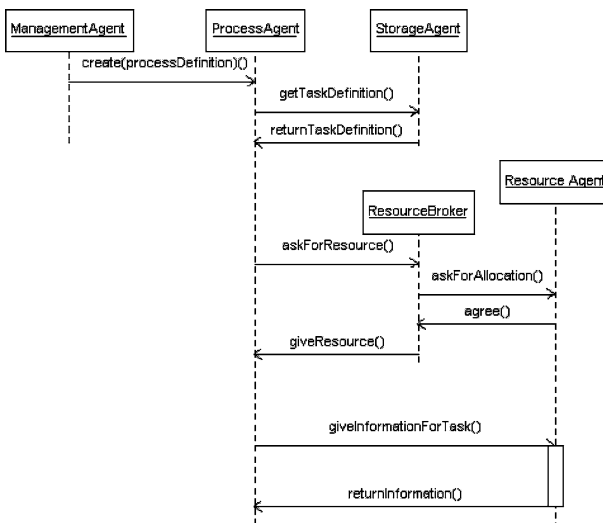


Fig. 5 The sequence diagram for allocating a resource and executing a task

## 4.2 Demonstration of adaptability

Suppose the company changes the way an order is dealt with. The process model in the right of Fig. 6 is the new definition of this workflow. The tasks, shipping and billing, are now executed in series, which changes the degree of parallelism.

Our architecture, which assigns one process agent to every work case, gives the flexibility to handle this change:

- New work cases start with the new definition.
- Running work cases are handled individually. This means every single process agent decides whether the work case is in a safe region or not. If the token is in the safe region, the agent changes the definition to the new process model as shown on the right of Fig. 6. In our example a safe work case would be for a situation in which the Approval transition has not been fired (in other words, there should be no token in the dashed region shown in the diagram on the left of Fig. 6). If Approval has fired, there appears to be no automated algorithm that can decide how to transfer the case. In this case, the human manager can decide what to do with unsafe work cases: either to go on with the old definition or allow manual transfer of the token by the human manager.

The use of one agent for every work case allows us to easily change a process only for a specific number of particular work cases. The manager only has to select the agents that are running these work cases and give them a new definition.

## 4.3 Demonstration of distribution

Let us now assume that the company has several branches, for example, suppose the shipping department is located in Dunedin and the customer

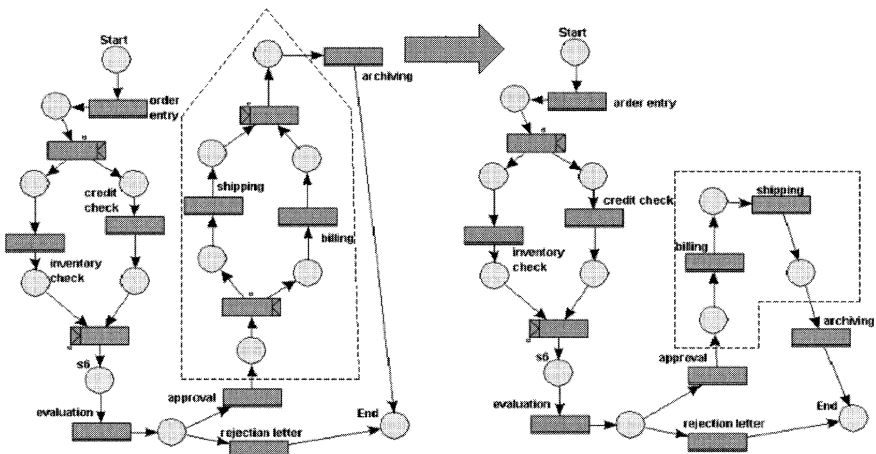


Fig. 6 The old and the new nets showing the unsafe region

service is located in Auckland. Because of this distributed nature of the company, it is desirable to have a hierarchical workflow and run the sub-processes at the local hosts.

The top-level CPN model for this example is shown in Fig. 7. The subnets are well defined. Now the process agent gets this top-level CPN model and executes it. Every transition has an attribute that indicates if it is a sub-process, a task, or a control transition. Depending on the value of this attribute, the process agent executes the action code of the control transition, requests a resource for the task or creates a new process agent for the sub-process. In case of a new sub-process it starts a new process agent (which can be located at another host if necessary). This process of the distributed workflow hierarchy can be arbitrarily extended. Note that there is no single Petri net that controls the entire process. There is one agent that controls the top-level process, but the agents for the sub-processes work independently. The results of the sub-processes are integrated accordingly with the top-level process. Sub-processes are only created when they have to be executed, so they automatically have the latest versions of the sub-process definition.

#### 4.4 Monitoring

In simulation enactment mode, the various tasks of the process model were associated with a set of values for the process parameters. These parameters

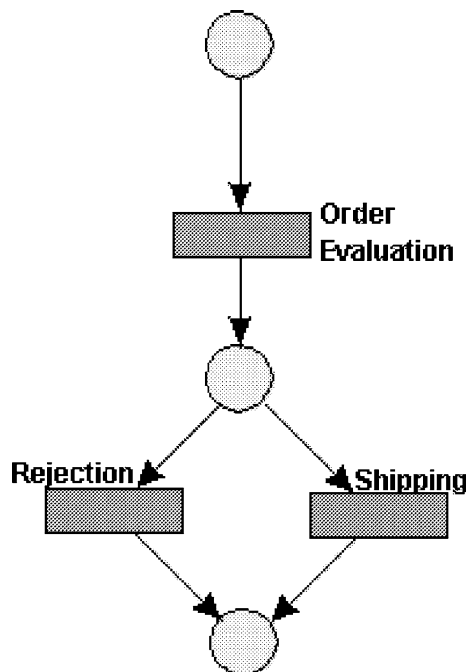


Fig. 7 The overall process of the distributed workflow

include the number of resources, their availability, and their ability (time taken) to do these tasks to simulate various “real life” scenarios. By varying these process parameters, various test scenarios can be examined. The storage agent stores the results of the simulated cases.

In execution mode, every time an action is started or finished, the process agent sends a message with the exact time of these events to the monitor agent. This is also done by the resource broker in the event of resource registration or deregistration. The monitor agent collects these data. After it gets the notification message from the process agent that the work case is finished, it combines the data and sends it to the storage agent for persistent storage.

Figure 8 shows the overall utilization of the resources and Fig. 9 shows the waiting time for various tasks in similar cases.

#### 4.5 The feedback mechanism

The human manager specifies the various criteria that are to be monitored continuously for violations. For example, the criteria could be “the completion time for a case of process X should not exceed twice the average of completion times of all cases of the same process (process X)”. The controlling agent looks for anomalies to these criteria and reports them to the manager agent. The manager agent might then decide to initiate the appropriate feedback action or display the warning so that the human manager can take appropriate action. This feedback mechanism helps in continuous process improvement.

## 5 Conclusion

In this paper we have described the architecture of a prototype agent-based workflow management system JBees. The use of coloured Petri nets as process formalism and of collaborating software agents as the basis of the systems provides more flexibility than existing WfMSs. JBees provides a high

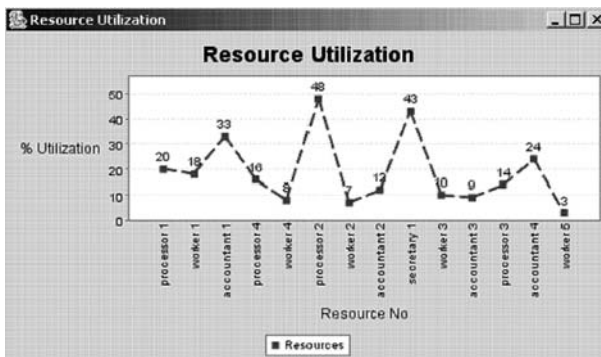


Fig. 8 The percentage utilization of the resources



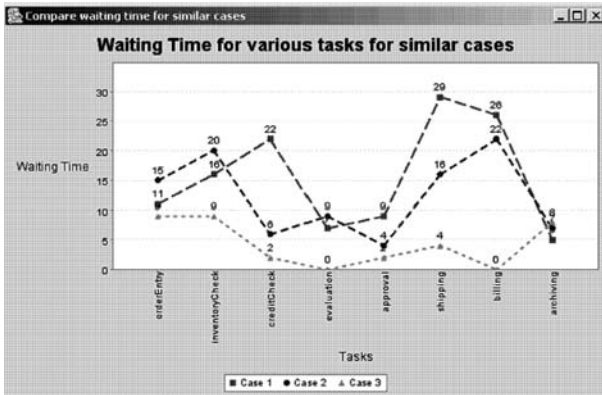


Fig. 9 The waiting time for various tasks in similar cases

degree of distribution and can deal with various levels of adaptability. We implemented the described architecture and this paper showed solutions for the problems arising while implementing this system. Our existing framework has been endowed with monitoring and feedback mechanisms, so that the various processes and cases can be studied, analysed, and the required feedback can be given to the workflow manager.

As part of our future work we intend to integrate into the currently implemented framework one of the existing CPN tools that supports formal analysis so that we can examine the model for certain properties such as reachability. This information as well as the monitoring information can be used to improve the processes in order to optimize the effectiveness of the system.

Another part of our future work will be the integration of sophisticated resource management. One step would be to store all tasks that have been performed by a particular resource (a task history) to allow more efficient resource scheduling. There are lots of resource management strategies already developed which we could as well apply to our system. The prototype described is available under the GNU Lesser General Public License (Free Software Foundation 2000) on the internet (Department of Information Science 2004).

**Acknowledgements** The authors wish to thank Mariusz Nowostawski for his help in implementing the system and improving JFern, and Prof. Martin Purvis for supporting this work. This work has been supported by the Otago Research Grant ORG-0103-0304. Lars Ehrlar thanks the Evangelisches Studienwerk Villigst e.V., Schwerte, Germany, for their financial support.

## References

- van der Aalst WMP (1996) Three good reasons for using a Petri-net-based Workflow Management System. In: Navathe S, Wakayama T (eds) Proceedings of the international working conference on information and process integration in enterprises (IPIC 1996), Cambridge, pp 179–201

- van der Aalst WMP (2001) Exterminating the dynamic change bug: a concrete approach to support workflow change. *Inf Syst Front* 3(3):297–317
- van der Aalst WMP, van Hee K (2002) *Workflow management: models, methods, and systems*. MIT Press, Cambridge
- van der Aalst WMP, Basten T, Verbeek HMW, Verkoulen PAC, Voorhoeve M (1999) Adaptive workflow: an approach based on inheritance. In: Ibrahim LM, Drabble B (eds) *IJCAI 1999 Workshop on intelligent workflow and process management: the new frontier for AI in business*, pp 36–45
- van der Aalst WMP, Hofstede AHM, Kiepuszewski B, Barros AP (2002) *Workflow patterns*. Technical report, Queensland University of Technology <http://tmitwww.tm.tue.nl/research/patterns/>
- Bandinelli S, Fuggetta A, Ghezzi C (1993) Process model evolution in the spade environment, Technical Report No. 14, ESPRIT-III Project GOODSTEP (6115). *IEEE transactions on software engineering*, 19(12):1128–1144
- Bradshaw J (1997) An introduction to software agents. In: Bradshaw J (ed) *Software agents*. MIT Press, Cambridge, pp 3–46
- Chen Q, Hsu M, Dayal U, Griss ML (2000) Multiagent cooperation, dynamic workflow and XML for ecommerce automation. In: 4th international conference on autonomous agents, Barcelona, Spain
- Cui B, Odgers Z, Schroeder M (1998) An in-service agent monitoring and analysis system. In: 11th IEEE international conference on tools with artificial intelligence, Chicago, USA, pp 237–244
- Department of Information Science, University of Otago. JBees. <http://jbees.sourceforge.net> (2004)
- Fleurke M (2004) JBees, an adaptive workflow management system—an approach based on petri nets and 10 Lars Ehrler et al. agents. Master's Thesis, Department of Computer Science, University of Twente
- Fleurke M, Ehrler L, Purvis M (2003) JBees—an adaptive and distributed framework for workflow systems. In: Ali Ghorbani, Stephen Marsh (eds) *Workshop on collaboration agents: autonomous agents for collaborative environments (COLA)*. Halifax, Canada, pp 69–76 <http://www.cs.unb.ca/~ghorbani/cola/proceedings/NRC-46519.pdf>, 2003. National Research Council Canada, Institute for Information Technology
- Free Software Foundation (2000) GNU Lesser General Public License
- Jensen K (1992) *Coloured Petri nets—basic concepts, analysis methods and practical use, vol 1. Basic concepts*. EATCS monographs on theoretical computer science. Springer, Berlin Heidelberg New York
- Jennings NR, Faratin P, Norman TJ, O'Brien P, Odgers B (2000) Autonomous agents for business process management. *Int J Appl Artif Intell* 14(2):145–189
- Joeris G (2000) Decentralized and flexible workflow enactment based on task coordination agents. In: 2nd international bi-conference workshop on agent-oriented information systems (AOIS 2000 @ CAISE\*00), Stockholm, Sweden, iCue Publishing, Berlin, pp 41–62
- Meilin S, Guangxin Y, Yong, Shanguang W (1998) *Workflow management systems: a survey in proceedings of IEEE international conference on communication technology*, [scw.cs.tsinghua.edu.cn/cscwpapers/ygxin/WfMSSurvey.pdf](http://scw.cs.tsinghua.edu.cn/cscwpapers/ygxin/WfMSSurvey.pdf)
- Muehlen MZ, Rosemann M (2000) Workflow-based process monitoring and controlling—technical and organizational issues. In: 33rd Hawaii international conference on system sciences, Maui, HI, USA
- Nissen ME (2000) Supply chain process and agent design for E-commerce. In: 33rd Hawaii international conference on system sciences
- Nowostawski M, Bush G, Purvis MK, Cranefield S (2001) A multilevel approach and infrastructure for agent-oriented software development. In: *International workshop on infrastructure for agents, MAS and Scalable MAS*, [http://www.umcs.maine.edu/~wagner/workshop/01\\_nowostawski\\_bush\\_purvis\\_et\\_al.pdf](http://www.umcs.maine.edu/~wagner/workshop/01_nowostawski_bush_purvis_et_al.pdf), 2001
- Nowostawski M (2003) JFern—Java-based Petri net framework
- Purvis M, Purvis M, Lemalu S (2001) A framework for distributed workflow systems. In: 34th Annual Hawaii international conference on system sciences (HICSS-34)
- Purvis MK, Cranefield S, Nowostawski M, Carter D (2002) Opal: a multi-level infrastructure for agent-oriented software development. *The information science*

- 
- discussion paper series no 2002/01, Department of Information Science, University of Otago, Dunedin
- Savarimuthu BTR, Purvis M, Fleurke M (2004) Monitoring and controlling of a workflow management system. In: Proceedings Australasian Workshop on Data Mining and Web Intelligence (DMWI2004)
- Schael T (1998) Workflow management systems for process organisations. Lecture notes in computer science, vol 1096. Springer, Berlin Heidelberg New York
- Shepherdson JW, Thompson SG, Odgers B (1998) Cross organisational workflow coordinated by software agents. In: CEUR workshop proceedings No 17 cross-organisational workflow management and coordination, San Francisco, USA
- Shoham Y (1997) An overview of agent-oriented programming. In: Bradshaw J (ed) Software agents. MIT Press, Cambridge, pp 271–290
- Stormer H (2001) AWA—a flexible agent-workflow system. In: Workshop on agent-based approaches to B2B at the 5th international conference on autonomous agents (AGENTS 2001), Montreal, Canada
- Sycara KP (1998) Multiagent systems. *AI magazine* 19(2):79–92
- The workflow management coalition (1995) The workflow reference model
- Wang M, Wang H (2002) Intelligent agent supported flexible workflow monitoring system. In: Advanced information systems engineering: 14th international conference, CAiSE 2002, Toronto, Canada
- Wooldridge MJ (1999) Intelligent agents. In: Weiss G (ed) Multiagent systems. MIT Press, Cambridge, pp 27–77