

Schedule-Aware Workflow Management Systems

R.S. Mans^{1,2}, N.C. Russell¹, W.M.P. van der Aalst¹, A.J. Moleman², P.J.M. Bakker²

¹ Department of Information Systems, Eindhoven University of Technology, P.O. Box 513, NL-5600 MB, Eindhoven, The Netherlands.

{r.s.mans,n.c.russell,w.m.p.v.d.aalst}@tue.nl

² Academic Medical Center, University of Amsterdam, Department of Quality Assurance and Process Innovation, Amsterdam, The Netherlands.

{a.j.moleman,p.j.bakker}@amc.uva.nl

Abstract. Contemporary workflow management systems offer work-items to users through specific work-lists. Users select the work-items they will perform without having a specific schedule in mind. However, in many environments work needs to be scheduled and performed at particular times. For example, in hospitals many work-items are linked to appointments, e.g., a doctor cannot perform surgery without reserving an operating theater and making sure that the patient is present. *One of the problems when applying workflow technology in such domains is the lack of calendar-based scheduling support.* In this paper, we present an approach that supports the seamless integration of unscheduled (flow) and scheduled (schedule) tasks. Using CPN Tools we have developed a specification and simulation model for schedule-aware workflow management systems. Based on this a system has been realized that uses YAWL, Microsoft Exchange Server 2007, Outlook, and a dedicated scheduling service. The approach is illustrated using a real-life case study at the AMC hospital in the Netherlands.

1 Introduction

Healthcare is a prime example of a domain where the effective execution of tasks is often tied to the availability of multiple scarce resources, e.g. doctors. In order to maximize the effectiveness of individual resources and minimize process throughput times, typically an appointment-based approach is utilized for scheduling the tasks performed by these resources. However, the scheduling of these appointments is often undertaken on a manual basis and its effectiveness is critically dependent on preceding tasks being performed on-time in order to prevent the need for rescheduling.

To illustrate the importance of the afore-mentioned issue, consider the small hospital process for diagnosing a patient shown in Figure 1. As a first step, the patient is registered. Next, a physical examination (task “physical examination”) of the patient takes place which is done by an assistant and a nurse. In parallel, a nurse prepares the documents for the patient (task “make documents”). When these tasks have been completed, a doctor evaluates the result of the test (task




Fig. 1. Running example showing schedule (S) and flow (F) tasks. The prefix “d:” indicates the average time needed for performing the task and prefix “r:” indicates which roles are necessary to perform the task. From each associated role, exactly one person needs to be assigned to the task. For both schedule tasks, indicated by the character “P” in the top-right corner of the task, the patient is also required to be present.

“consultation”) and decides about the information and brochures that need to be provided by the nurse (task “give brochures and information”). Figure 1 also shows the corresponding organizational model which specifies the roles being played by people in the organization.

From this example, it can be seen that a distinction can be made between two kinds of tasks. The tasks annotated with an “F” in the figure, can be performed at an arbitrary point in time when a resource becomes available and are called *flow tasks*. However, the tasks “physical examination” and “consultation”, annotated with an “S” in the figure, can only be performed when the required room is reserved, the patient is present, and the necessary medical staff are present for performing the specific task, i.e. these tasks need to be scheduled and performed at particular times. Therefore, we call these kinds of tasks *schedule tasks* as they are performed by one or more resources at a specified time.

For the consultation task in the figure, it is often the case that a doctor finds out at the actual appointment that some results from required diagnostic tests are missing. Consequently, this leads to wasted time for the doctor as a new appointment needs to be scheduled. Therefore, for the effective performance of schedule tasks it is vital that the whole workflow is taken into account in order to guarantee that preceding tasks are performed on-time thereby preventing the need for rescheduling and avoiding unproductive time for resources as a result of canceled appointments.

Workflow technology presents an interesting vehicle with which to support healthcare processes. Based on a corresponding process definition, Workflow Management Systems (WfMSs) support processes by managing the flow of work such that individual work-items are done at the right time by the proper person [2]. Contemporary WfMSs offer work-items through so-called work-lists. At an arbitrary point in time, a user can pick a work-item from this list and perform the associated task.

If we consider the implementation of this process in the context of a WfMS, we find that a significant dichotomy exists in that people are used to working

in a scheduled way, but this is not supported by current WfMSs. In contrast to administrative processes, healthcare processes invoke the coordination of expensive resources which have scarce availability. Therefore, it is of the utmost importance that the scheduling of appointments for these resources is done in an efficient way, that is suitable both for the medical staff and also for the patients being treated. To summarize, there is a *need to integrate workflow management systems with scheduling facilities*.

In this paper, we present the design and implementation of a WfMS supporting both schedule and flow tasks. In addition to the classical work-list functionality generally associated with workflow systems, the concept of a calendar is also introduced in order to present appointments for scheduled work-items to the people involved. Unlike traditional workflow implementations, our focus is on how WfMSs can be *integrated* with scheduling facilities rather than simply extending the functionality of a WfMS or a scheduling system (e.g. a scheduling algorithm). In other words, we investigate how scheduling facilities can be added to workflow systems in general.

An interesting problem in this context lies in the actual development approach taken to extending a WfMS with scheduling facilities. Our strategy for this is based on the use of CPN Tools, a widely used modeling and execution tool for Colored Petri Nets (CP Nets) [18], with which we developed a *comprehensive conceptual model capable of serving both as a specification and simulation model for the application domain*. Formalizing such a system using CP Nets offers several benefits. First of all, building such a net allows for *experimentation*. So, the model or parts of it can be executed, simulated, and analyzed which leads to important insights about the design and implementation of the system. Second, the hierarchical structuring mechanism of CP Nets allows for the modeling of large complex systems at different levels of abstraction. That is, CP Nets can be structured into a set of components which interact with each other through a set of well-defined interfaces, in a similar way to the components in a modular software architecture.

In this way, we were able to use the conceptual model as a *specification* for the subsequent realization of the system. In order to realize the functionality contained in the conceptual model, we *incrementally* mapped it to an operational system based on widely available open-source and commercial-off-the-shelf (COTS) software. Although the conceptual model is detailed, it remains abstract enough, such that its components can be concretized in many different ways. We choose an approach based on the reuse of existing software. In total, the conceptual model consists of 30 nets, 238 transitions, 625 places, and in excess of 1000 lines of ML-code illustrating the overall complexity of the system. For the concrete realization of the system we used the open-source, service-oriented architecture of YAWL and Microsoft Exchange Server 2007 as the implementation platform.

The remainder of the paper is organized as follows. In Section 2 we explain how a workflow language can be augmented with information relevant for scheduling. In Section 3 we present the design of a WfMS integrated with schedul-

ing facilities, together with a concrete implementation. In Section 4 a concrete application of the realized system is presented. Section 5 discusses related work and finally Section 6 concludes the paper.

2 Flow and Schedule Tasks

In order to allow for the extension of a WfMS with scheduling functionality some concepts need to be introduced. It is assumed that the reader is familiar with basic workflow management concepts, like case, role, and so on [2]. Using the process shown in Figure 1, we will elaborate on how a workflow language can be integrated with scheduling functionality.

2.1 Concepts

We can distinguish between two distinct types of tasks. Flow tasks are performed at an arbitrary point in time when a resource becomes available. As only one resource is needed, it is sufficient to define only *one* role for each of them³. Consequently, these tasks can be presented in an ordinary *work-list*. For example, for the flow task “make documents” the work may either be performed by “Sue” or “Rose”.

Conversely, schedule tasks are performed by one or more resources at a specified time. As multiple resources can be involved, with different capabilities, it is necessary to specify which kinds of resources are allowed to participate in completing the task. To this end, multiple resources may be defined for a schedule task where for each role specified, only *one* resource may be involved in the actual performance of the task. For example, in Figure 1, the schedule task “physical examination” may be performed by “Jane” and “Rose”, but not by “Sue” and “Rose”. Note that a resource involved in the performance of a schedule task may also be a physical resource such as medical equipment or a room. Furthermore, for the schedule tasks the patient may also be involved which means that the patient is also a required resource for these tasks. Note that the patient is not involved in the actual execution of the task but is a passive resource who needs to be present whilst it is completed. For this reason, the patient is not added to any of the roles for the task, nor are they defined in terms of a separate role. Instead, it is necessary to identify for which schedule tasks the patient needs to be present.

For presenting the appointments made for schedule tasks to users, the concept of a *calendar* will be used. More specifically, each resource will have its own calendar in which appointments can be booked. Note that each patient also has his / her own calendar. An appointment either refers to a schedule task which needs to be performed for a specific case or to an activity which is not workflow related. So, an appointment appears in the calendars of all resources that are

³ There also exist approaches for which more roles may be defined, but this is not the focus of our work.

involved in the actual performance of the task. An appointment for a schedule task, for which a work-item does not yet exist, can be booked into the calendar of a resource. However, when the work-item becomes available it has already been determined when it will be performed and by whom. Note that sometimes work-items need to be rescheduled because of anticipated delays in preceding tasks.

In order to be able to determine at runtime the earliest time that a schedule task can be started, information about the duration of every task needs to be known. For example, in Figure 1, for each task the average duration is indicated by prefix “d:”. For example, one block represents one minute, which means that the task “physical examination” takes 60 minutes on average.

2.2 Formalization

Based on the informal discussion in the previous section, we now formalize the augmented workflow language. The definition of our language is based on WF-nets [2]. Note that our results are in no way limited to WF-nets and can be applied to more complex notations (BPM, EPCs, BPEL, etc). Note that WF-nets are the most widely used formal representation of workflows. A WF-net is a tuple $N = (P, T, F)$ defined in the following way:

- P is a non-empty finite set of *places*;
- T is a non-empty finite set of *tasks* ($P \cap T = \emptyset$);
- $F \subseteq (P \times T) \cup (T \times P)$ is a set of arcs (flow relation);
- There is one initial place $i \in P$ and one final place $o \in P$ such that every place or transition is on a directed path from i to o .

A WF-net can be extended in the following way, called a *scheduling WF-net* (sWF-net). A sWF-net is a tuple $N = (P, T_f, T_s, F, CR, Res, Role, R, Rtf, Rts, D)$, where:

- T_f is a finite set of *flow tasks*;
- T_s is a finite set of *schedule tasks*;
- $T_f \cup T_s = T$ and $T_f \cap T_s = \emptyset$, i.e., T_s and T_f partition T . So, a task is either a flow task or a schedule task, but not both;
- (P, T, F) is a WF-net;
- $CR \subseteq T_s$ is the set of schedule tasks for which the human resource for whom the case is being performed is also required to be present.
- Res is a non-empty finite set of *resources*;
- $Role$ is a non-empty finite set of *roles*;
- $R: Res \rightarrow \mathcal{P}(Role)$ is a function which maps resources on to sets of roles;
- $Rtf: T_f \rightarrow Role$ is a partial function which maps flow tasks on to roles;
- $Rts: T_s \rightarrow \mathcal{P}(Role) \setminus \{\emptyset\}$ is a function which maps schedule tasks on to at least one role;
- $D: T \rightarrow \mathbb{N}_0$ is a function which maps tasks onto the number of blocks that are needed for the execution of the task.

Note that Figure 1 fully defines a particular sWF-net.

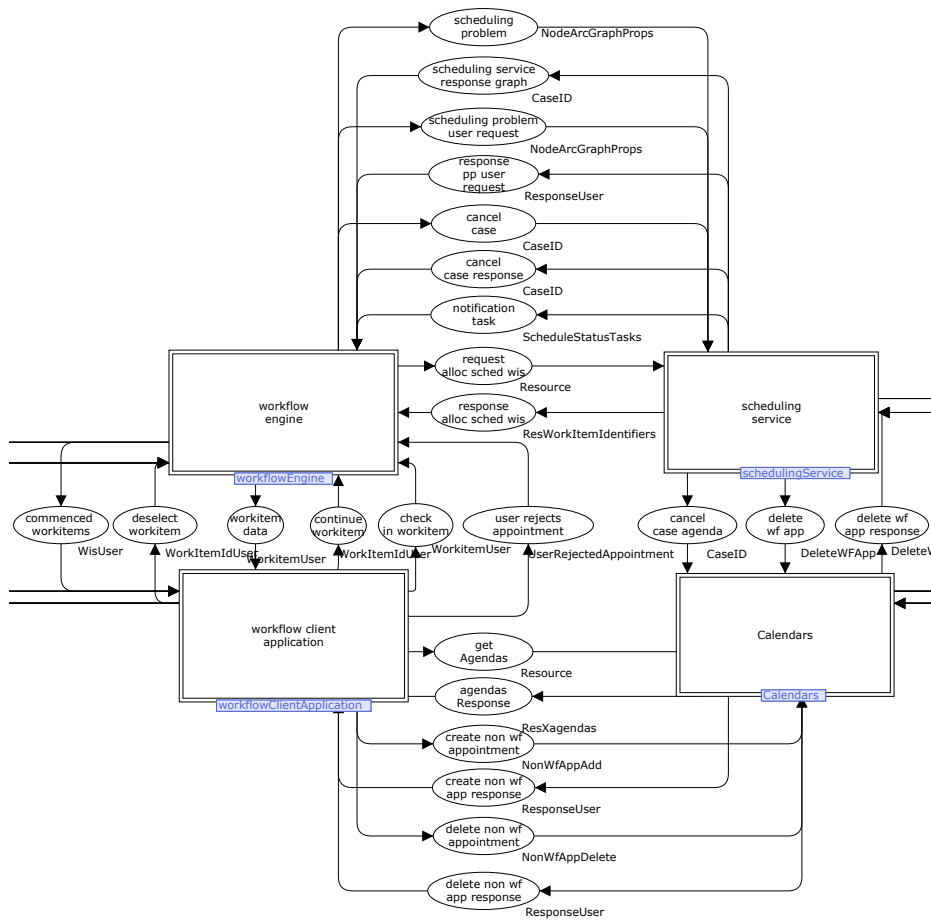


Fig. 3. The topmost model of the conceptual model realized in terms of CP Nets.

system allows for *testing* parts of the system that are implemented. Given that a CP Net consists of several components, we can “replace” one or more components in the CP Net by the concrete implementation of these components. This replacement is realized by making connections between the CP Net model and components in the actual system. As the CP Net is an executable model we can easily test numerous scenarios facilitating the discovery of potential flaws in both the architecture and the implementation.

Another important benefit of having a CP Net consisting of several components, is that it provides precise guidance in the configuration of software products, thereby allowing for the use of existing software. As will become clear below, whilst the specification model is detailed, it remains abstract enough, such that it allows components to be concretized in various ways.

3.2 Conceptual model

The complete conceptual model, defining the precise behavior of a workflow management system augmented with scheduling facilities, consists of a series of CP Nets in which several layers can be distinguished. Figure 2 shows the hierarchy of CP Nets in the CP Net model, together with the relationships between them. In total, there are 30 distinct CP Nets. An indication of the complexity of each net is expressed by the p and t value included for each them, showing the number of places and transitions they contain. In total, the whole CP Net consists of 238 transitions, 625 places and in excess of 1000 lines of ML code. The construction of the whole model required more than three months of work. This underscores the fact that it is a complex system. In order to illustrate the use of CP Nets in the development approach, the blocks in Figure 2 which are colored grey will be discussed in more detail.

Figure 3 shows the topmost net in the CP Net model and gives an idea of the main components in the system and the interfaces between them. As can be seen in the figure, there are four substitution transitions. They represent the major functional units in the system which are explained in detail in the next section. Each place which is connecting two substitution transitions forms part of the interface between the corresponding two components. Additionally, for each place, being part of such an interface, the direction of the arcs between the place and the two components specifies the sending (tail) and the receiving component (head) of a specific message.

3.3 Architecture

Now, we give a more detailed overview of the architecture of a WfMS integrated with scheduling facilities. As Figure 3 only provides a snippet of the whole conceptual model (it only shows 22 places of the 57 places that are part of an interface), Figure 4a shows the architecture of the conceptual model together with an illustration of the interfaces between the components (shown as clouds). Note the correspondence between Figure 4 and the top level CP Net model (Figure 3). The concrete implementation of the systems architecture is shown in Figure 4b. Both architectures illustrate the main components and the system is defined in a service oriented way. The components are loosely coupled and the interfaces are kept as compact and simple as possible. As the interfaces share the same numbering, it is easy to compare both sets of interfaces.

The architecture consists of four components. First of all, the *workflow engine* routes cases through the organization. Based on the business process definition for a case, tasks are carried out in the right order and by the right people. Once a task in a case becomes available for execution, the corresponding work-item is communicated to users via the *workflow client application* allowing it to be selected and performed by one of them. The scheduling service and the workflow client application communicate with the *Calendar* component in order to obtain a view on users' calendars and to manipulate their contents. Note that users can add /remove appointments that are possibly unrelated to the workflow.

As our focus is on how a WfMS can be *integrated* with scheduling facilities, we want to completely separate the scheduling facilities provided by the system from the engine. As a consequence, we have a separate *scheduling service* component which is responsible for providing scheduling facilities to the system (e.g. (re)scheduling of tasks). In order for the scheduling service to function correctly, all scheduling constraints imposed by the engine (which might be relevant to a scheduling decision) need to be sent to the scheduling service. To be more precise, the scheduling service receives a *scheduling problem*, which contains all relevant constraints for one case only. Based on these constraints, the scheduling service makes decisions with regard to the scheduling of schedule tasks for the case.

Informally, the scheduling problem is formulated as a *graph* which has *nodes* and *arcs* between nodes. Nodes, arcs and the graph itself may have properties represented as name-value attributes. The rationale for representing the scheduling problem using this data structure is that any information in the graph can be included which is deemed relevant. For a case, which is in a given state, we

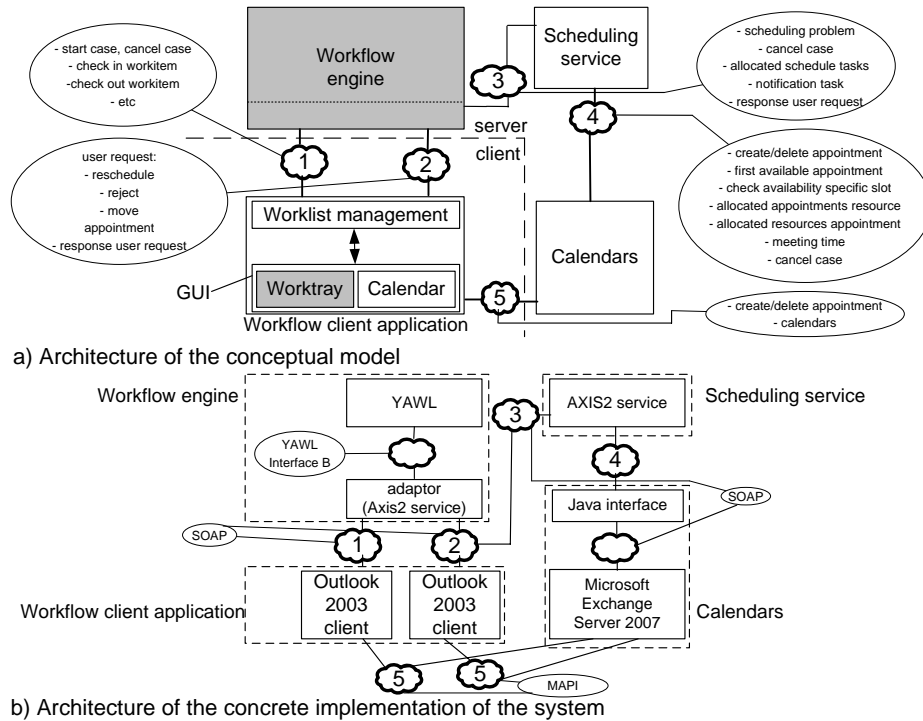


Fig. 4. Architectures of both the conceptual model and the concrete implementation of the system. There are four main components: (I) workflow engine, (II) scheduling service, (III) workflow client application, and (IV) calendars. The distinct interfaces are indicated by numbers.

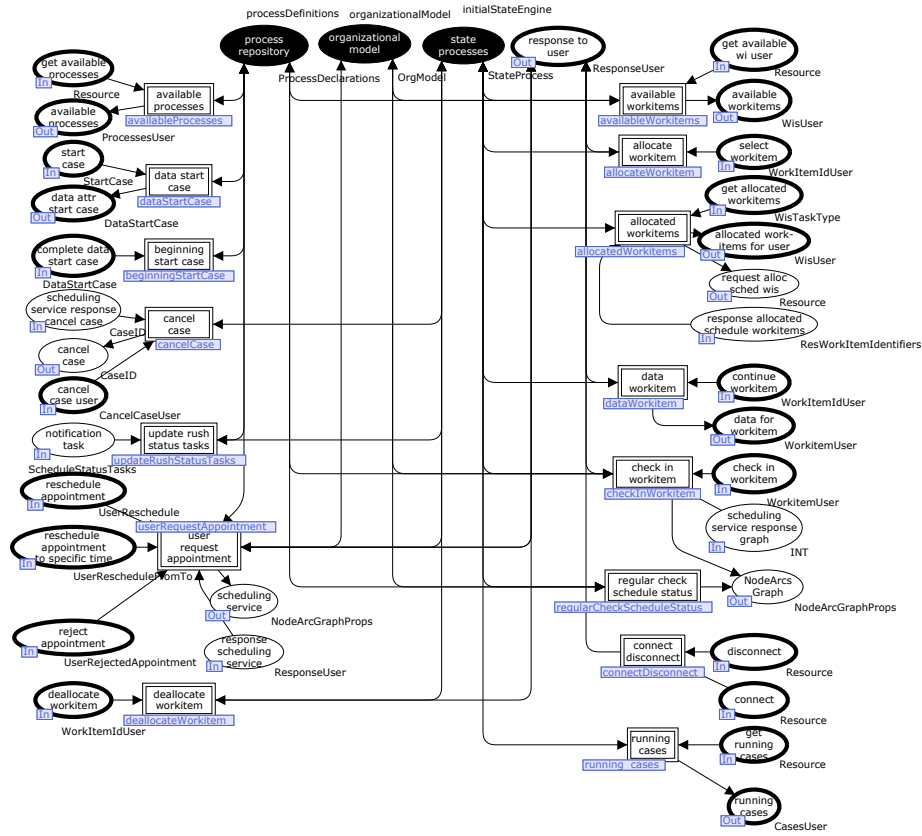


Fig. 6. CP Net component for the workflow engine component.

As a consequence of the execution of the scheduling service, the engine is informed about appointments for which limited time is left in which to finish work-items of preceding tasks. For these work-items, a warning is sent to the workflow client to indicate that limited time remains in which to finish the work-item.

Model.

The complete CP Net for the workflow engine component is depicted in Figure 6. The places tagged with an input or output port tag (shown by a rectangle with the text “In” or “Out” respectively) are part of the interface of the engine with other components. More specifically, the white colored places with a thick black line are part of the interface of the workflow client application. The white colored places with a thin black line are part of the interface of the scheduling service and the black colored places are places which keep internal engine information.

The substitution transitions have been defined such that each of them specifies a specific functionality provided by the engine. For example, the “cancel case” substitution transition models the actions taken by the engine in case

of cancellation of a case. A cancellation request is received via the “cancel case user” place. In that case, the corresponding case information is removed from the “state cases” place which contains the state information for all cases. Afterwards, the scheduling service is informed, via place “cancel case”, so that appointments made for the case are deleted. As an indication of the complexity of the engine it is worth mentioning that the flattened substitution transition comprises 63 transitions and 151 places.

Implementation.

The Engine component in the CP Net model can be replaced by a concrete implementation which allows it to be tested. The workflow component is realized (see Figure 4b) using the open-source WfMS YAWL [1] and a service which acts as an adaptor in between YAWL and the workflow client application. The adaptor service communicates with YAWL via “Interface B” [1]. The adaptor also communicates with the scheduling service using SOAP messages. However, the adaptor and the YAWL system are tightly coupled as large volumes of work-item and process related information are exchanged.

3.5 Workflow client application

Users working with the WfMS do so via the workflow client application which delivers the basic facilities that should be provided by this facility [2]. The component consists of a GUI and a work-list management component. The work-list management component serves as a layer between the engine and the GUI and takes care of the communication between them. The GUI component consists of a “worktray” and a “calendar” component where the “worktray” provides the same facilities as a classical worktray. The appointments that are created for schedule tasks are advertised via the calendar. Once a work-item becomes available for such an appointment, it can be performed via the calendar. In our approach, only one user can interact with the WfMS with respect to the completion of the work-item. This prevents concurrency issues where multiple users want to complete the same work-item.

With regard to the appointments that are made for schedule tasks, users can express their dissatisfaction with the nominated scheduling by requesting: (1) the rescheduling of the appointment, (2) the rescheduling of the appointment to a specified date and time, or (3) the reassignment of the appointment to another employee. Such a user request can be done as a single action and is the only supported means for the rescheduling of appointments by users. In addition, the workflow client also indicates whether limited time is left in which to undertake work-items in order to meet the schedule. Moreover, users are also allowed to add appointments to the calendar which are not workflow related (e.g. having dinner with friends).

As can be seen in Figure 4a, two interfaces are defined for the communication between the workflow client application and the engine. The interface with number “1” defines the standard communication that takes place between an engine and a workflow client application. The interface with number “2” defines methods added as a consequence of the scheduling facilities developed for the

system. For this interface, nothing is stored in the engine when these methods are called.

Model.

The corresponding CP Net model for the work-list management component is fairly complex (and is not shown here): the component's model contains 109 transitions and 239 places. However, in order to allow for experimentation, we also had to include some user behavior, like selecting a workitem and so on, which has nothing to do with the actual behavior of the system.

Implementation.

Similarly, the workflow client application component of the CP Net can be replaced by a concrete implementation. Once the Exchange Server was in place we could easily use the Microsoft Outlook 2003 client to obtain a view of a user's calendar. Furthermore, the Outlook client can be configured in such a way that it can act as a full workflow client application which can communicate with the WfMS via an adaptor service via the exchange of SOAP messages.

3.6 Scheduling Service

The scheduling service is responsible for providing scheduling facilities to the WfMS. Scheduling is done sequentially on a case-by-case base. Once a scheduling problem is received, the scheduling service needs to determine whether some of the schedule tasks need to be (re)scheduled. Moreover, several distinct issues need to be addressed of which we mention the most important ones.

First of all, the final scheduling of tasks needs to occur in the same order as the sequence of schedule tasks in the accompanying process definition for the case. Moreover, there should be sufficient time between two scheduled tasks. Also, when rescheduling appointments, the preceding constraints need to be satisfied. For example, in Figure 1, it needs to be guaranteed that first the "physical examination" is scheduled, followed by the "consultation" which needs to occur at a later time.

Second, for the actual scheduling of an appointment multiple roles can be specified for a schedule task. For each role specified a resource needs to be selected, i.e., the number of roles determines the number of resources involved in the actual performance of the task. If the patient for which the case is performed also needs to be present at an appointment, then this is also taken into account. The scheduling service only books an appointment in the calendars of these resources who need to be present during the performance of the task (i.e. the performers of the task and the patient (if needed)).

Third, the scheduling service is also responsible for determining whether limited time is left for performing preceding work-items for scheduled tasks. In such a situation, the engine needs to be informed. Moreover, the scheduling service is also informed about the cancelation of a case, so that all appointments related to the case can be removed. When too little time is left for performing preceding work-items for a scheduled schedule task, the corresponding appointment is automatically rescheduled which in this context can be seen as the most straight-

forward recovery action. However, different strategies can also be conceived for dealing with such situations. Potential solutions can be found in [4].

In this paper, we focus on integration aspects instead on devising new scheduling algorithms. Nevertheless, to demonstrate the approach that is used for the scheduling of appointments, we will briefly examine the implemented ‘naive’ scheduling algorithm. Of course it can be envisaged that more advanced scheduling strategies are possible.

The (re)scheduling of appointments is done automatically, which means that there is no user involvement. Starting with the tasks in the graph for which a work-item exists, it is determined which schedule tasks need to be (re)scheduled. Once we know that tasks are able to be scheduled, they are scheduled. Moreover, these tasks are scheduled on a sequential basis in order to avoid conflicts involving shared resources. However, we do not schedule any tasks which occur after a choice in the process as this can lead to unnecessary usage of available slots in the calendar. Moreover, we do not take loops into account.

For the actual scheduling of an appointment, a search is started for the first opportunity where one of the resources of a role can be booked for the respective work-item. If found, an appointment is booked in the calendar of the resource. If the patient for which the case is performed also needs to be present at the appointment, then this is also taken into account. For example, for Figure 1, if a case is started, an appointment is created for task “physical examination” in the calendars of “Jane”, “Sue”, and the patient, or “Jane”, “Rose” and the patient.

Model.

The CP Net model which models the scheduling service consists of 49 transitions and 150 places. Moreover, modeling the scheduling behavior necessitated writing many lines of ML code.

Implementation.

The concrete implementation of this component of the CP Net is shown in Figure 4b. Here we see that the component is implemented in Java as a service which communicates with the WfMS via SOAP messages. However, in order to get a view of and to manipulate the calendar, the service also communicates via a Java interface with the Exchange Server which in turn exchanges information via SOAP messages.

3.7 Calendar

The Calendar component is responsible for providing a view on the calendars of users and for manipulating their contents. It is possible to create / delete appointments or to get information about the appointments that have been made. Moreover, the interface contains some convenience methods for deleting cases and finding the first available slot for a schedule task. Otherwise, large volumes of low-level information need to be exchanged whereas now only one call is necessary.

Model.

The CP Net model which models the scheduling service consists of 13 transitions and 28 places. Note that this model is relatively simple.

Implementation.

For the Calendar component we selected Microsoft Exchange Server 2007 as the system for storing the calendars of users. The big advantages of this system are its widespread use and the fact that it offers several interfaces for viewing and manipulating calendars.

3.8 Analysis

A serious drawback that we faced was that no meaningful verification of the CP Net was possible due to its size and complexity. Even more, as an unlimited number of business process models and users can be represented, state space analysis is impossible for the general case. Therefore, in [22], we present an approach in which one or more components in the CP Net are “replaced” by the concrete implementation of these components by making connections between the CP Net model and components in the actual system. As the CP Net is an executable model this made it easy to test numerous scenarios in a systematic manner. In this way, a multitude of flaws, ranging from small ones to serious design flaws, have been discovered and fixed in both the conceptual model and components in the actual system.

In the design and the implementation of the system, a naive algorithm has been used for the scheduling of appointments. This naive approach can lead to inefficient use of resources. In [21, 22], the conceptual model is used for simulation-based performance analysis in order to evaluate various scheduling approaches and to investigate the effects of our calendar-based approach on performance indicators such as the average waiting time for an appointment.

In total, these approaches lead to an increased confidence in the design, reliability, and correct operation of the resultant system.

4 Application

In this section, we demonstrate our approach and software in the context of a real-life healthcare scenario. To evaluate our approach, we have taken the diagnostic process of patients visiting the gynecological oncology outpatient clinic at the AMC hospital, a large academic hospital in the Netherlands. This healthcare process deals with the diagnostic process that is followed by a patient who is referred to the AMC hospital for treatment, up to the point where the patient is diagnosed, and consists of around 325 activities. However, for our scenario we will only focus on the initial stages of the process shown in Figure 7.

At the beginning of the process, a doctor in a referring hospital calls a nurse or doctor at the AMC hospital resulting in an appointment being made for the first visit of the patient. Several administrative tasks need to be requested before the first visit of the patient (e.g. task “first consultation doctor”). At the first consultation, the doctor decides which diagnostic tests are necessary (MRI, CT or pre-assessment) before the next visit of the patient (task “consultation doctor”). Note that for the MRI, CT and pre-assessment tasks we do not show

the preceding tasks at the respective departments that need to be performed in order to simplify the model presented.

For this scenario, we assume that the task “additional information and brochures” has been performed. Moreover, at the first consultation with the doctor it has been decided that an MRI and a pre-assessment are needed for the patient. So, by looking at the process model it becomes clear that the tasks “MRI”, “pre-assessment” and “consultation doctor” need to be scheduled. The result of the scheduling performed by the system for these tasks is shown in Figure 8a. Note that our case has “Oncology” as its process identifier and has “15” as its case identifier. Moreover, for the “consultation doctor”, “pre assessment”, and “MRI” examination, a doctor, an anesthetist, and MRI machine are needed respectively. Moreover, the patient is also required to be present.

In Figure 8a we can see that the “MRI” has been scheduled for 10:00 to 10:45 (see first column), the consultation with the doctor has been scheduled for 13:00 to 13:30 in the calendar of doctor “Nick” (see second column), and that the pre-assessment has been scheduled for 11:00 to 11:30 in the calendar of anesthetist “Jules” (see third column). At the far right, we can see the calendar of patient Anne who also needs to be present for the work-items mentioned, which explains why the previously mentioned appointments are also present in her calendar. For Anne we see that she is not available till 10 ’o clock which has influenced the actual scheduling. This is due to the fact that she can not manage to be at the hospital before 10 ’o clock by public transport. However, it is important that the “consultation doctor” task is scheduled after the “MRI” and “pre-assessment” task, which is also consistent with the corresponding process definition.

Now, let us assume that unexpectedly some maintenance for the MRI machine is necessary for that day, which will take until 13:30 hours to complete. Consequently, the MRI appointment needs to be rescheduled to 13:30 hours. The effect of this specific rescheduling request can be seen in Figure 8b. In this figure, the message box indicates that the MRI has been successfully rescheduled to the requested time. Moreover, in the calendar of Anne we can see that the MRI now takes place from 13:30 to 14:15. However, it was also necessary to reschedule the

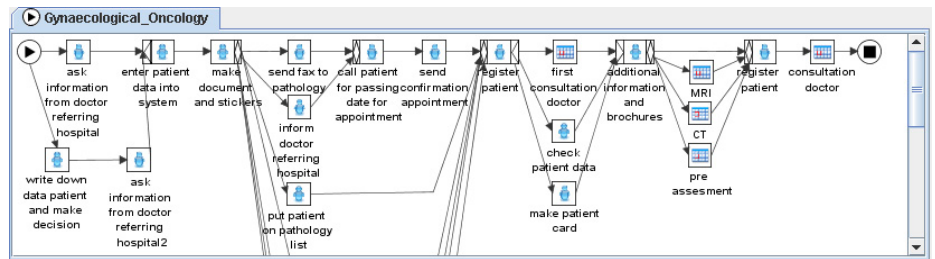


Fig. 7. Screenshot of the YAWL editor showing the initial stages of the gynaecological oncology healthcare process. The flow tasks are indicated by a person icon and the schedule tasks are indicated by a calendar icon.




Fig. 8. Screenshot of the calendars for the MRI, consultation with the doctor, and the pre-assessment before and after rescheduling.

appointment with doctor “Nick” which will now take place from 14:30 to 15:00. As can be seen in Figure 7, this rescheduling step is necessary as the task “consultation doctor” occurs after the “MRI” task and the task “register patient” falls in between these two tasks and takes 15 minutes.

5 Related Work

Analysis of the healthcare research shows that significant work has been done on the problem of appointment scheduling. Examples of such research efforts are appointment scheduling for outpatient services [10, 17] and operating room scheduling [9]. However, most of these studies focus on a single unit instead of situations in which a patient may pass through multiple facilities. Another approach is described in [27] in which the online problem of scheduling multiple appointments on a single day is considered for outpatients thereby only considering the appointments which need to be scheduled. In our research, we take the scheduling of work-items for the whole workflow into account together with the current state of a case.

Our work is also related to time management in workflows. For example, in [23, 15] the authors focus on the satisfiability of time constraints and the enforcement of these at run-time. In addition, there is also research on the problem of

the scheduling of tasks by WfMSs. For example, [7, 12, 25, 26] present algorithms for the scheduling of tasks. In contrast, we focus on the augmentation of a WfMS with scheduling facilities instead of just presenting new scheduling algorithms.

The work presented in [11] is somewhat similar to ours as it presents different architectures for a WfMS in which temporal aspects are explicitly considered. However, the temporal reasoning facilities are added as core functionality to the engine. In this paper, we propose a different approach where this kind of functionality is realized through a separate service in the system. In this way, loose coupling is guaranteed which means that our approach can be generalized to any WfMS (or even to multiple engines at the same time).

Multiple people can be involved in the actual performance of a schedule task. However, in our approach, only one user can interact with the WfMS with respect to the completion of a work-item. In [3, 8, 14] reference models to extend the organizational meta model with a team concept allowing for the distribution of work to teams are proposed. By doing so, advanced mechanisms are offered for the performance of work by such a team. Additionally, in [3, 8] a language is discussed for defining work allocation requirements to people.

For various systems, CP Nets have been used to formalize and validate functional requirements. A good overview can be found in [13]. For example: the formalization of the design of the so-called worklet service, which adds flexibility and exception handling capabilities to the YAWL workflow system [5]; formalizing the implementation of a healthcare process in a workflow management system [20]; presenting a model-based approach to requirements engineering for reactive systems [16]; expressing architectural requirements and assessing middleware support in terms of decoupling [6]; and formalizing the implementation of a patient record system by modeling task descriptions [19]. Related to this is [24], in which CP Nets are used for specifying the operational semantics of *newYAWL*, a business process modeling language founded on the well-known workflow patterns⁴. To the best of our knowledge we are not aware of any work in which a comprehensive conceptual model of similar size and complexity is used for the specification of a newly developed system. Moreover, the specific component-based construction of the conceptual model allowed for the incremental mapping of the functionality contained in the conceptual model to an operational system based on both open-source and COTS software (YAWL, Microsoft Exchange Server, and Microsoft Outlook).

6 Conclusions and Future Work

In this paper, we have presented the design and implementation of a WfMS augmented with calendar-based scheduling facilities. Instead of just offering work-items via a work-list, as is the case in most existing WfMSs, they can also be offered as a concrete appointment in a calendar taking into account which preceding tasks are necessary and whether they have been performed.

⁴ For more information about workflow patterns, see <http://www.workflowpatterns.com>

Our approach demonstrates that the use of CP Nets, for constructing a conceptual model of the system to be realized, provides valuable insights in terms of understanding the problem domain and identifying the behavior of the system. Moreover, the same conceptual model provides a comprehensive specification on which to base the ultimate realization of the required functionality. We have incrementally mapped it to an operational system using widely available open-source and commercial-off-the-shelf (COTS) software. For the concrete realization of the system we used YAWL (as a workflow engine) and Microsoft Exchange Server 2007 / Outlook (as the platform providing calendar and user interaction). This demonstrates that although the specification model is detailed, it remains at a sufficient level of abstraction to allow its constituent components to be concretized in various ways. Moreover, it also shows that our ideas can, for example, be applied to a variety of WfMSs and scheduling systems.

Finally, to test the feasibility of our approach, we plan to evaluate the operation of our resultant system in a real-life scenario at the AMC hospital.

References

1. W.M.P. van der Aalst, L. Aldred, M. Dumas, and A.H.M. ter Hofstede. Design and Implementation of the YAWL System. In *Proceedings of CAiSE'04*, 2004.
2. W.M.P. van der Aalst and K.M. van Hee. *Workflow Management: Models, Methods, and Systems*. MIT Press, Cambridge, MA, 2002.
3. W.M.P. van der Aalst and A. Kumar. A Reference Model for Team-Enabled Workflow Management Systems. *Data and Knowledge Engineering*, 38(3):335–363, 2001.
4. W.M.P. van der Aalst, M. Rosemann, and M. Dumas. Deadline-based Escalation in Process-Aware Information Systems. *Decision Support Systems*, 43(2):492–511, 2007.
5. M.J. Adams. *Facilitating Dynamic Flexibility and Exception Handling for Workflows*. PhD thesis, Faculty of Information Technology, Queensland University of Technology, 2007.
6. L. Aldred, W. M.P. van der Aalst, M. Dumas, and A.H.M. ter Hofstede. On the notion of coupling in communication middleware. In R. Meersman and Z. Tari, editors, *Proceedings of the International Symposium on Distributed Objects and Applications (DOA 05)*, volume 3761 of *LNCS*, pages 1015–1033, 2005.
7. C. Bettini, X.S. Wang, and S. Jajodia. Temporal Reasoning in Workflow Systems. *Distributed and Parallel Databases*, 11(3):269–306, 2002.
8. J. Cao, S. Zhang, and X. Zhang. Team Work Oriented Flexible Workflow Management System. In X. Meng, J. Su, and Y. Wang, editors, *Advances in Web-Age Information Management*, volume 2419 of *LNCS*, pages 189–200, 2002.
9. B. Cardoen, E. Demeulemeester, and J. Beliën. Operating room planning and scheduling: A literature review. *To appear in European Journal of Operational Research*, 2009.
10. T. Cayirli and E. Veral. Outpatient Scheduling in Health Care: A Review of Literature. *Product Operations Management*, 12(4):519–549, 2003.
11. C. Combi and G. Pozzi. Architectures for a Temporal Workflow Management System. In H. Haddad, A. Omicini, R.L. Wainwright, and L.M. Liebrock, editors, *Proc. of the 2004 ACM symposium on applied computing*, pages 659–666, 2004.

12. C. Combi and G. Pozzi. Task Scheduling for a Temporal Workflow Management System. In *Thirteenth International Symposium on Temporal Representation and Reasoning (TIME'06)*, pages 61–68, 2006.
13. CPN Group, Aarhus. Overview of projects documenting the large-scale practical use of CP-nets and their tools. <http://www.cs.au.dk/CPnets/intro/industrial.shtml>.
14. L. Cui and H. Wang. Research on Cooperative Workflow Management Systems. In W. Shen, Z. Lin, and J.-P.A. Barthès, editors, *Computer Supported Cooperative Work in Design I*, volume 3168 of *LNCS*, pages 359–367, 2005.
15. J. Eder, E. Panagos, and M. Rabinovich. Time Constraints in Workflow Systems. In M. Jarke and A. Oberweis, editors, *Proceedings of CAiSE '99*, volume 1626 of *Lecture Notes in Computer Science*, pages 286–300. Springer-Verlag, Berlin, 1999.
16. J.M. Fernandes, J.B. Jørgensen, and S. Tjell. Requirements Engineering for Reactive Systems: Coloured Petri Nets for an Elevator Controller. In *Proceedings of the 13th Asia-Pacific Software Engineering Conference (APSEC 2007)*, pages 294–301. IEEE Computer Society Press, 2007.
17. D. Gupta and B. Denton. Appointment scheduling in health care: Challenges and opportunities. *IIE Transactions*, 40(9):800–819, 2008.
18. K. Jensen, L.M. Kristensen, and L. Wells. Coloured Petri Nets and CPN Tools for Modelling and Validation of Concurrent Systems. *International Journal on Software Tools for Technology Transfer*, 9(3-4):213–254, 2007.
19. J.B. Jørgensen, K.B. Lassen, and W.M.P. van der Aalst. From task descriptions via colored Petri nets towards an implementation of a new electronic patient record workflow system. *International Journal on Software Tools for Technology Transfer (STTT)*, 10(1):15–28, 2007.
20. R.S. Mans, W.M.P. van der Aalst, N.C. Russell, P.J. Bakker, A.J. Moleman K.B. Lassen, and J.B. Jørgensen. From requirements via colored workflow nets to an implementation in several workflow systems. *Transactions on Petri Nets and Other Models of Concurrency (ToPNoC) III*, 5800:25–49, 2009.
21. R.S. Mans, W.M.P. van der Aalst, N.C. Russell, P.J. Bakker, and A.J. Moleman. Process-Aware Information System Development for the Healthcare Domain - Consistency, Reliability, and Effectiveness. In F. Leymann, S. Sadiq, and S. Rinderle-Ma, editors, *Business Process Management Workshops 2009*, LNBIP, 2009.
22. R.S. Mans, W.M.P. van der Aalst, N.C. Russell, P.J.M. Bakker, and A.J. Moleman. Model-based Development and Testing of Process-Aware Information Systems. In A. Alimohammad, A. Meixner, and M. Popescu, editors, *Proceedings of The First International Conference on Advances in System Testing and Validation Lifecycle (VALID 2009)*, pages 129–134. IEEE Computer Society Press, 2009.
23. O. Marjanovic and M. Orłowska. On Modeling and Verification of Temporal Constraints in Production Workflows. *Knowledge and Information Systems*, 1(2):157–192, 1999.
24. N. Russell, W.M.P. van der Aalst, and A.H.M. ter Hofstede. Designing a Workflow System Using Coloured Petri Nets. *Transactions on Petri Nets and Other Models of Concurrency (ToPNoC) III*, 5800:1–24, 2009.
25. P. Senkul and I.H. Toroslu. An Architecture for Workflow Scheduling under Resource Allocation constraints. *Information Systems*, 30(5):399–422, 2005.
26. G.B. Tramontina, J. Wainer, and C.A. Ellis. Applying Scheduling Techniques to Minimize the Number of Late Jobs in Workflow Systems. In *Proc. of the 2004 ACM Symposium on Applied Computing*, pages 1396–1403, 2004.

27. I. Vermeulen, H. La Poutré, S.M. Bohte, S.G. Elkhuisen, and P.J. Bakker. Decentralized Online Scheduling of Combination-Appointments in Hospitals. In *Proceedings of ICAPS-2008, the International Conference on Automated Planning and Scheduling*, Sydney, Australia, 2008. AAAI Press.