

Designing a Workflow System using Coloured Petri Nets^{***}

Nick Russell¹, Wil M.P. van der Aalst^{1,2} and Arthur H.M. ter Hofstede²

¹Eindhoven University of Technology,
PO Box 513, 5600MB, Eindhoven, The Netherlands
{n.c.russell,w.m.p.v.d.aalst}@tue.nl

²Queensland University of Technology,
PO Box 2434, QLD, 4001, Australia
a.terhofstede@qut.edu.au

Abstract. Traditional workflow systems focus on providing support for the control-flow perspective of a business process, with other aspects such as data management and work distribution receiving markedly less attention. A guide to desirable workflow characteristics is provided by the well-known workflow patterns which are derived from a comprehensive survey of contemporary tools and modelling formalisms. In this paper we describe the approach taken to designing the *newYAWL* workflow system, an offering that aims to provide comprehensive support for the control-flow, data and resource perspectives based on the workflow patterns. The semantics of the *newYAWL* workflow language are based on Coloured Petri Nets thus facilitating the direct enactment and analysis of processes described in terms of *newYAWL* language constructs. As part of this discussion, we explain how the operational semantics for each of the language elements are embodied in the *newYAWL* system and indicate the facilities required to support them in an operational environment. We also review the experiences associated with developing a complete operational design for an offering of this scale using formal techniques.

Keywords: *newYAWL*, workflow technology, workflow patterns, business process management, coloured Petri nets

1 Introduction

There are a plethora of workflow systems on the market today providing organisations with various forms of automated support for their business processes. It is ironic however, that despite the rigour that workflow systems introduce

* An earlier version of this work was presented at PNDS'08, the International Workshop on Petri Nets and Distributed Systems [1].

** This research is conducted in the context of the *Patterns for Process-Aware Information Systems (P4PAIS)* project which is supported by the Netherlands Organisation for Scientific Research (NWO).

into the conduct of the processes that they coordinate, they themselves do not demonstrate the same rigour in the workflow languages that they enact. Indeed, it is a salient fact that, almost without exception, workflow languages are defined on an informal basis leaving their precise operation unclear to anyone other than the system developers. An additional shortcoming of existing workflow solutions is their focus on the control-flow aspects of business processes.

The *YAWL* Initiative sought to address the first of these issues. *YAWL* [2] is an acronym for *Yet Another Workflow Language*. It provides a comprehensive modelling language for business processes based on formal foundations. The content of the *YAWL* language is an adaptation of Petri Nets informed by the workflow patterns [3]. One of its major aims was to show that a relatively small set of constructs could be used to directly support most of the workflow patterns identified. It also sought to illustrate that they could coexist within a common framework. In order to validate that the language was capable of direct enactment, the *YAWL System*¹ was developed, which serves as a reference implementation of the language. Over time, the *YAWL* language and the *YAWL System* have increasingly become synonymous and have garnered widespread interest from both practitioners and the academic community alike².

Initial versions of the *YAWL System* focussed on the control-flow perspective and provided a complete implementation of 19 of the original 20 patterns. Subsequent releases incorporated limited support for selected data and resource aspects of processes, however this effort was hampered by the lack of a complete formal description of the requirements in these perspectives. Recent work conducted as part of the Workflow Patterns Initiative has identified the core elements in other process perspectives (data, resource, exception handling) and a recent review [4] of the control-flow perspective has identified 23 additional patterns which illustrate a number of commonly used control-flow constructs, many of which *YAWL* is unable to provide direct support for, including the partial join, transient and persistent triggers, iteration and recursion.

In an effort to manage the conceptual shortcomings of *YAWL* with respect to the range of workflow patterns that have now been identified, a substantial revision of the language — termed *newYAWL* is proposed — which aims to support the broadest range of the workflow patterns in the control-flow, data and resource perspectives. *newYAWL* synthesises this work to provide a fully formalised workflow language based on a comprehensive view of a business process. The validation of this proposal is to design (and ultimately build) the workflow system that embodies the workflow language. An interesting consequence of formalising the operational semantics for the language constructs in *newYAWL*, has been the establishment of the functional architecture for the system to be

¹ See <http://www.yawl-system.com> for further details of the *YAWL System* and to download the latest version of the software.

² Hereafter in this paper, we refer to the collective group of *YAWL* offerings developed to date — both the *YAWL* language as defined in [2] and also more recent *YAWL System* implementations of the language based on the original definition (up to and including release Beta 8.2) — as *YAWL*.

developed. This is based on a detailed consideration of the causal effects and data interactions required to support each of the language constructs and their behaviour in a broader operational environment. This paper outlines the approach taken to designing the *newYAWL* system. In this paper we not only describe the design of *newYAWL* using Coloured Petri Nets, but also reflect on the use of such a design approach from a software engineering standpoint.

The remainder of this paper proceeds as follows: Section 2 introduces the YAWL language from a functional perspective. Section 3 presents *newYAWL*, a significant extension to YAWL that encompasses the broad range of workflow patterns which identify desirable workflow functionality that have recently been identified. Section 4 describes the approach to designing a workflow system that can enact business processes described in terms of the *newYAWL* language. Section 5 overviews related work and Section 6 discusses the experiences of designing a workflow system using formal methods and concludes the paper.

2 YAWL: Yet Another Workflow Language

YAWL has its genesis in the workflow patterns which aimed to delineate desirable constructs in the control-flow perspective of workflow processes. Hence the initial version of the YAWL language focussed solely on control-flow aspects of processes. It aimed to show that the patterns could be operationalised in an integrated framework. Furthermore, it also aimed to show that this could be achieved in the context of a formal framework, providing both a syntax and semantics for language constructs to remove any potential for ambiguity or uncertainty in their interpretation.

The formal foundation for YAWL is based on hierarchical Petri nets and there is a striking similarity between the two graphical representations with tasks taking the place of classical Petri net transitions and conditions representing the various states between tasks in the same way that places typically serve as the inputs and outputs to transitions in Petri nets. However Petri nets only serves as a basis for the fundamental aspects of YAWL and it significantly extends its capabilities in a variety of ways.

1. YAWL allows for tasks to be directly connected by an arc in the situation where there would normally be a single condition between them (that was not connected to any other tasks);
2. YAWL provides for direct representation of AND-split, AND-join, XOR-split, XOR-join and OR-split constructs rather than requiring their explicit modelling in terms of fundamental language constructs. In conjunction with direct connections between tasks, this serves to significantly simplify process models;
3. The notion of the (inclusive) OR-join, which is frequently described in process modelling notations without any consideration of how it will actually be enacted, is directly available as a modelling construct in YAWL. Moreover, there is a complete formal treatment of its operationalisation described in [5];

4. Task concurrency, which is not a consideration in many process modelling formalisms, is directly represented via the notion of the multiple instance task. This denotes a task (or subprocess) which executes multiple times in parallel with some or all instances needing to be synchronised before the thread of control can pass to subsequent tasks;
5. Cancellation of individual tasks, portions of a process or even an entire process can be explicitly represented in YAWL process models through the notion of a cancellation region. Cancellation regions are attached to a specific task in a process and when it completes, any threads of control residing in conditions within the cancellation region are removed and any executing tasks within the cancellation region are terminated; and
6. A YAWL process model (whether it is the top-level net in a process or a subprocess definition) has a single start and endpoint denoted by specific input and output conditions. This provides a precise semantics for process enablement and termination and allows a range of verification techniques to be applied to YAWL process models.

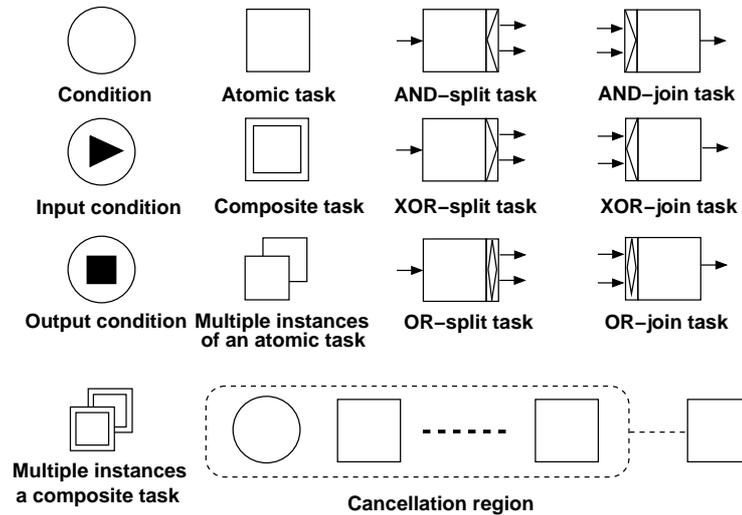


Fig. 1. YAWL symbology

The specific symbols used in a YAWL process are shown in Fig. 1. An example of a complete YAWL model using many of these symbols is depicted in Fig. 2. It denotes the sales fulfillment process for an ironmonger which sells and manufactures specialist metal fittings. An instance of the process is triggered when an order is received. This then initiates two distinct branches of activities (as signified by the outgoing AND-split). The first of these focusses on the financial aspects of the order. The other triggers the tasks associated with the actual assembly of the order for despatch, manufacturing and packing tasks. The first

branch involves a *credit check*. If the customer has sufficient funds, then their order can be invoiced. It progresses to the *despatch* stage when the order handling is complete. If there is insufficient credit, then a reminder is sent and there is a waiting period for the required payment to be received from the customer. If it is not received within 10 days, the order is cancelled (denoted by the *timeout* task which is linked to a cancellation region which encompasses all tasks which might be active in the process). The branch associated with order assembly involves a series of tasks. First the order is prepared for picking. This is a composite task involving a sequence of three distinct activities: reviewing the order contents, producing a picking slip for items available from the warehouse and reviewing the requirements for any parts that need to be specially cast. Having done this, one or both of the *picking* and *custom cast* tasks are triggered via the OR-split operator from the *prepare picking* task. The *custom cast* task is a multiple instance task and a separate instance of it is triggered for each component that requires manufacturing. Once all of the *picking* and *custom cast* tasks that were initiated have been completed, the order is packed. After invoicing and packing have been completed, the *despatch* task can run, followed by the archiving of the order and completion of the case.

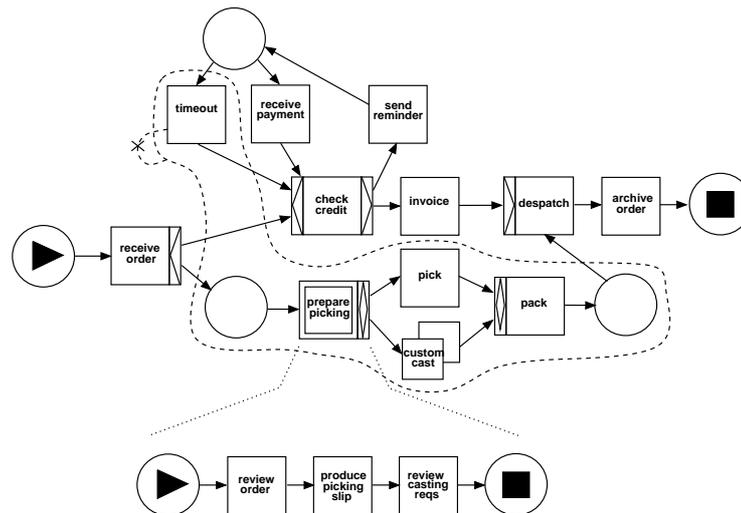


Fig. 2. Example of a YAWL process model: sales fulfillment

YAWL provides effective coverage of many commonly encountered control-flow constructs. However, a recent review [4] of the control-flow perspective identified a number of additional control-flow constructs that commonly arise in business processes. Moreover, there is also the need to consider other aspects of business processes, such as the requirements encountered in the data [6] and resource [7] perspectives as delineated by the data and resource patterns. In the

next section, we introduce a comprehensive extension to the YAWL language that addresses these issues.

3 *new*YAWL: Extending YAWL to Multiple Perspectives

*new*YAWL is a multi-perspective business process modelling and enactment language founded on the workflow patterns. It provides a comprehensive and integrated formal description of the workflow patterns, which to date have only partially been formalised. It has a complete abstract syntax which identifies the characteristics of each of the language elements together with an executable semantic model in the form of a series of Coloured Petri Nets which define the runtime semantics of each of the language constructs. The following sections provide an overview of the features of *new*YAWL in the control-flow, data and resource perspectives.

3.1 Control-Flow Perspective

The control-flow perspective of *new*YAWL is based on the revised workflow control-flow patterns [4] and serves to significantly extend the control-flow capabilities of the current YAWL language [2]. It retains all of the existing language elements in *YAWL* and they continue to perform the same functions. Several new constructs have been added based on the full range of workflow patterns that have now been identified. These are identified in Fig. 3. The specific capabilities provided by each of them are as follows:

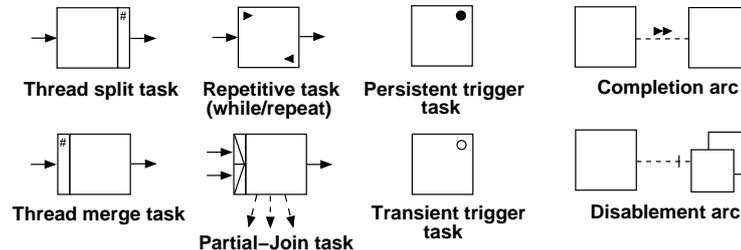


Fig. 3. Additional control-flow constructs in *new*YAWL

- the *Thread split* and *Thread merge* constructs, allow the thread of control to be split into multiple concurrent threads or several distinct threads to be merged into a single thread of control respectively. The number of threads being created/merged is specified for the construct in the design-time model. Fig. 4(a) illustrates these constructs. After the *make box* task, twelve threads of control are created ensuring that the *fill bottle* task runs 12 times before the *pack box* task can run (merging these threads before it commences);

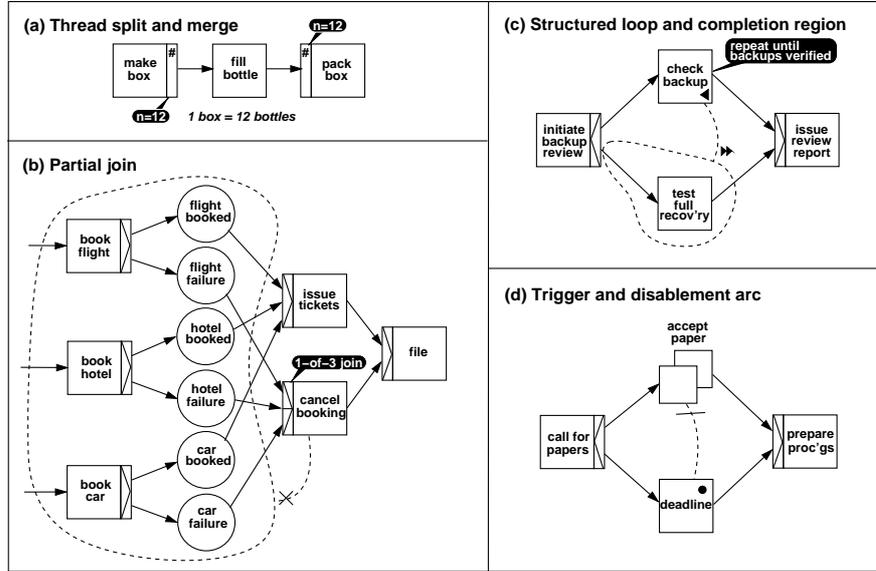


Fig. 4. Examples of *newYAWL* control-flow constructs

- the *Partial join* (also known as the *m-out-of-n join*) allows a series of incoming branches to be merged such that the thread of control is passed to the subsequent branch when m of the incoming n branches are enabled. The number of active threads required for the partial join to fire is specified in the design-time model. After firing, it resets (and can fire again) when all incoming branches have been enabled. In Fig. 4(b), the *cancel booking* task has a 1-out-of-3 join associated with it. If any of the incoming branches are enabled, then the *cancel booking* task is enabled (and any preceding tasks that are still executing in the associated cancellation region are withdrawn);
- the *Structured loop* (which supports while, repeat and combination loops) allows a task (or a sequence of tasks in the form of a subprocess) to execute repeatedly based on conditional tests at the beginning and/or end of each iteration. The loop is structured in form and it has a single entry and exit point. The entry and/or exit conditions are specified in the design-time process model. Fig. 4(c) illustrates a repeat loop for the *check backup* task which executes repeatedly until all backups have been verified (i.e. it is a post-tested loop);
- the *Completion region* supports the forced completion of tasks which it encompasses. In Fig. 4(c) the *test full recovery* task is forcibly completed once (all iterations of) the *check backup* task has finished. This allows the *issue review report* task to be immediately enabled;
- *Persistent triggers* and *Transient triggers* support the enablement of a task being contingent on a trigger being received from the operating environment. They are durable or transient in form respectively. Each trigger is associated

with a specific task and has a unique type so that incoming triggers can be differentiated. These details are captured in the design-time process model. Fig. 4(d) illustrates a persistent trigger (assumedly associated with some form of alarm) which allows the *deadline* task to be enabled when it is received. As this trigger is durable in form, it is retained for future use if it is received before the thread of control arrives at the *deadline* task;

- the *Disablement arc* allows a dynamic multiple instance task to be prevented from creating further instances but allows for each of the currently executing instances to complete normally. Fig. 4(d) has a disablement arc associated with the *deadline* task which prevents any further papers from being accepted once it has completed.

3.2 Data Perspective

Whilst the control-flow perspective has received considerable focus in many workflow initiatives, the data perspective is often only minimally supported with issues such as persistence, concurrency management and complex data manipulation being outsourced to third party products. In an effort to characterise the required range of data facilities in a workflow language, *newYAWL* incorporates a series of features derived from the data patterns. These include:

- Support for a variety of *distinct scopes* to which data elements can be bound. This allows the visibility and use of data elements to be restricted. The range of data scopes recognised include: *global* (available to all elements of all process instances), *folder* (available to the elements of process instances to which the folder is currently assigned), *case* (available to all elements in a given process instance), *block* (available to all elements of a specific process or subprocess definition for a given process instance), *scope* (available to a subset of the elements in a specific top-level process or subprocess definition for a given process instance), *task* (available to a given instance of a task) and *multiple-instance* (available to a specific instance of a multiple instance task);
- *Formal parameters* for specifying how data elements are transferred between process constructs (e.g. block to task, composite task to subprocess decomposition, block to multiple instance task). These parameters take a function-based approach to data transfer, thus providing the ability to support inline formatting of data elements and setting of default values. Parameters can be associated with tasks, blocks and processes;
- *Link conditions* for specifying conditions on outgoing arcs from OR-splits and XOR-splits that allow the determination of whether these branches should be activated;
- *Preconditions* and *postconditions* for tasks and processes. They are evaluated at the enablement or completion of the task or process with which they are associated. Unless they evaluate to true, the task or process instance with which they are associated cannot commence or complete execution; and

- *Locks* which allow tasks to specify data elements that they require exclusive access to (within a given process instance) in order to commence. Once these data elements are available, the associated task instance retains a lock on them until it has completed execution preventing any other task instances from using them concurrently. The lock is relinquished once the task instance completes.

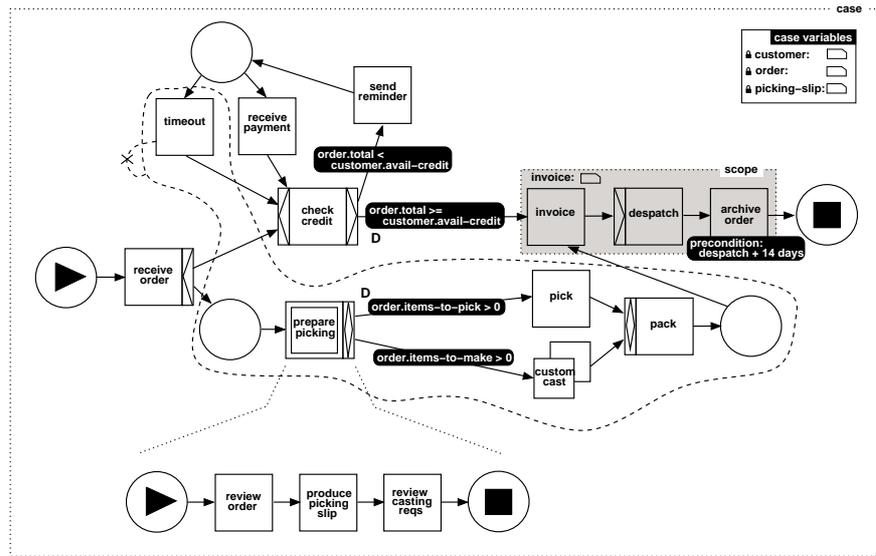


Fig. 5. Data perspective of sales fulfillment process

Figure 5 illustrates the main aspects of the data perspective for the sales fulfillment model (shown earlier in Fig. 2). The case variables *customer*, *order* and *picking-slip* are used throughout the tasks in the model. They are record-based in format and are passed on a reference basis between task instances. The lock beside each variable name indicates that when it is passed to a task instance an exclusive lock is applied to it whilst it is in use by that task instance to prevent problems arising from concurrent usage. In contrast to the case variables, the *invoice* variable is bound to a specific scope involving only three tasks and, although also passed by reference, does not require a lock as it is only updated by the first task in the scope and cannot be used concurrently by several tasks in the scope. The *credit check* task has an XOR-split associated with it and the (disjoint) conditions on outgoing branches are illustrated. Similarly, the *prepare picking* task has an OR-split associated with it and the two outgoing conditions are also shown although in this case, there is no requirement for them to be

disjoint since one or both outgoing branches can be enabled. For both split constructs, the default branch is indicated with a D and this branch is enabled if none of the conditions specified evaluate to true. The *archive order* task has a precondition associated with it which ensures it can only commence 14 days after the *despatch* task has completed. This is to allow for any returns or damage claims that might arise during transport.

3.3 Resource Perspective

The resource perspective in *newYAWL* provides a variety of means of controlling and optimising the way in which work is distributed to users and the manner in which it is progressed through to ultimate completion. For each task, a specific *interaction strategy* can be specified which precisely describes the way in which the work item will be communicated to the user, how their commitment to executing it will be established and how the time of its commencement will be determined. Similarly, a detailed *routing strategy* can be defined which determines the range of potential users that can undertake the work item. The routing strategy can nominate the potential users in a variety of ways — they can be directly specified by name, in terms of roles that they perform or the decision as to possible users can be deferred to runtime. There is also provision for determining the range of potential users based on capabilities that individual users possess, the organisational structure in which the process operates or the recorded execution history. The routing strategy can be further refined through the use of constraints that restrict the potential user population. Indicative constraints may include: *retain familiar* (i.e. route to a user that undertook a previous work item) and *four eyes principle* (i.e. route to a different user than one who undertook a previous work item). Allocation directives can also be used where a single user need to be selected from a group of potential users to whom a task can be allocated. Candidate allocation directives include *random allocation* (route to a user at random from the range of potential users), *round robin allocation* (route to a user from the potential population on an equitable basis such that all users receive a similar number of work items over time) and *shortest queue allocation* (route the work item to the user with the shortest work queue).

newYAWL also supports two advanced operating modes that are designed to expedite the throughput of work by imposing a defined protocol on the way in which the user interacts with the system and work items are allocated to them. These modes are: *pled execution* where all work items corresponding to a given task are routed to the same user and *chained execution* where subsequent work items in a process instance are routed to the same user once they have completed a preceding work item. Finally, there is also provision for specifying a range of user privileges, both at process and individual task level, that restrict or augment the range of interactions that they can have with the workflow engine when they are undertaking work items.

Figure 6 illustrates the resource perspective for the sales fulfillment process. Each task is annotated with the basic distribution strategy (DS) and interaction strategy (IS) for the task. The distribution strategy indicates which users and

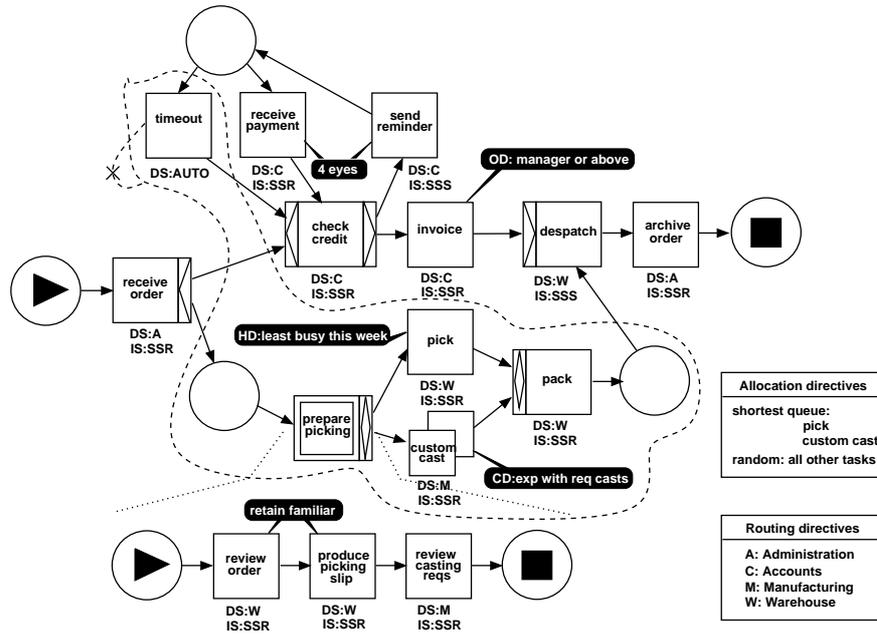


Fig. 6. Resource perspective of sales fulfillment process

roles the task will be routed to at runtime. For the purposes of this model, routing is either to specific roles (e.g. A is the role for Administration users) or AUTO where tasks are done automatically without requiring resource support. Extended routing directives apply to the *invoice*, *pick* and *custom cast* tasks. These operate in conjunction with the basic distribution strategy and further refine the specification of the user(s) to whom a task may be distributed. An organisational distribution directive applies to the *invoice* task requiring it be distributed to a member of the Accounts role that is at least a manager in organisational seniority. A historical distribution directive applies to the *pick* task requiring it be distributed to the member of the Warehouse role that has been least busy in the past week. A capability distribution directive applies to the *custom cast* task requiring it be distributed to the member of the Manufacturing role that has experience with the required casts that need to be manufactured for this order.

The interaction strategy indicates whether the system or an individual resource are responsible for triggering the offer, allocation and commencement of a task. For this model, the SSR and SSS interaction strategies are utilised. The former involves the system allocating the task to a specific user but the user being able to nominate the time at which they commence it, the latter (also known as “heads down processing”) involves the system allocating a task to a user and indicating when they should start it. In both cases, allocation directives are employed to select the individual user to whom a task should be allocated,

this is generally based on random selection of a user although the *pick* and *custom cast* tasks are allocated to user on the basis of who has the shortest current work queue. Distribution constraints exist between various pairs of tasks in the model. There is a four eyes constraint between the *receive payment* and *send reminder* tasks indicating they should not be allocated to the same user in a given case. There is also a retain familiar constraint between the *review order* and *produce picking slip* tasks indicating that they should be undertaken by the same user in a given case.

This section has focussed on providing a comprehensive introduction to the various language elements that make up *newYAWL* from a conceptual standpoint. In the next section, we discuss the design of a system that is able to operationalise these constructs.

4 *newYAWL*: The System

A workflow system encompasses a number of distinct functions as illustrated by the diagram in Fig. 7. Generally the business process that is to be automated is captured in the form of a *process model*. A *workflow management system* is responsible for coordinating the execution of instances of the process model. It comprises a number of discrete components. First the *workflow engine* is responsible for managing the control-flow and data elements that are associated with each process instance. As the thread of control flows through a process, it results in the triggering of individual tasks that make up the process model. The enabling of a task results in the creation of a new work item which needs to be executed by a human resource. However, in order for this to occur, the identity of one or more suitable resources needs to be determined. This activity is the responsibility of the *work item routing* component and is based on the interpretation of task routing information associated with each task in the context of the current state of the process instance.

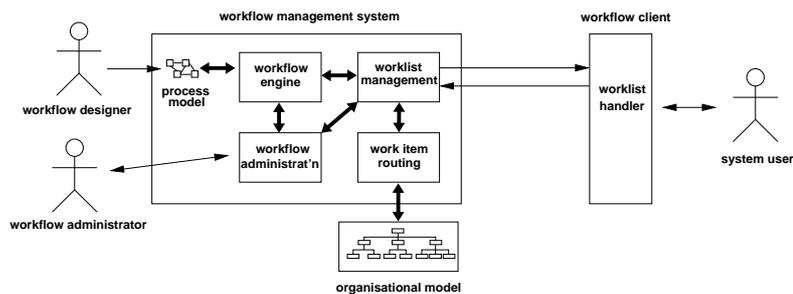


Fig. 7. Outline of major workflow system components

Once a set of suitable resources have been determined for a work item, it is necessary to advise them of the pending work item. This function is undertaken by the *worklist management* component which places the work item on the worklist of each resource to whom it is to be routed. Thus the workflow management system retains a centralised view of the state of all work items and also provides *workflow administration* facilities should it be necessary to intervene in the normal conduct of this process. However despite the consistent centralised view of pending work maintained by the workflow management system, there is another layer of complexity in managing the actual distribution and conduct of work items across the range of resources coordinated by the workflow system. This stems from the fact that resources typically operate independently of the workflow system. They retain a distinct view of the work that they are conducting which is accessed via a *worklist handler* (which typically takes the form of a software client running at a distinct location to that of the workflow management system). The worklist handler operates on a client-server basis with respect to the workflow management system. It is generally disconnected from the workflow system, connecting only when it wishes to refresh its view of the current work allocation or to advise the workflow system of a change of state in the work items it has been allocated.

Clearly a workflow system involves a relatively complex set of software components and interactions. In order to provide a precise definition of how a business process should actually be enacted in an operational environment, it is necessary not only to provide an operational semantics for the workflow language that describes the business process, but also to define the overall architecture and operation of the workflow system. This has been done for *newYAWL* using a series of interrelated Coloured Petri Nets developed using the CPN Tools environment [8]. This approach to formalising the system offers the dual benefits of establishing a precise definition of the operation of each of the language constructs which comprise *newYAWL* and also providing a means of describing exactly how an instance of a *newYAWL* specification should be executed. There are 55 distinct CPNs which make up the *newYAWL* system description. These are illustrated in Fig. 8 along with the relationships between them. The correspondence between the functional workflow system components identified in Fig. 7 and each of the CPNs is also delineated. An indication of the complexity of individual nets is illustrated by the p and t values included for each of them which indicate the number of places and transitions that they contain. Clearly it is not possible to discuss the operation of all of these nets in the confines of this paper, however some of them (indicated by the shaded boxes and cross-references) are discussed in further detail in subsequent sections. A comprehensive description of the 55 CPNs which comprise the *newYAWL* system can be found in [9]. In the following sections, we will outline the operation of three of these areas, illustrated by the shaded boxes in Fig. 8. These provide an overview of the *workflow engine*, *worklist management* and *worklist handler* components of the workflow system.

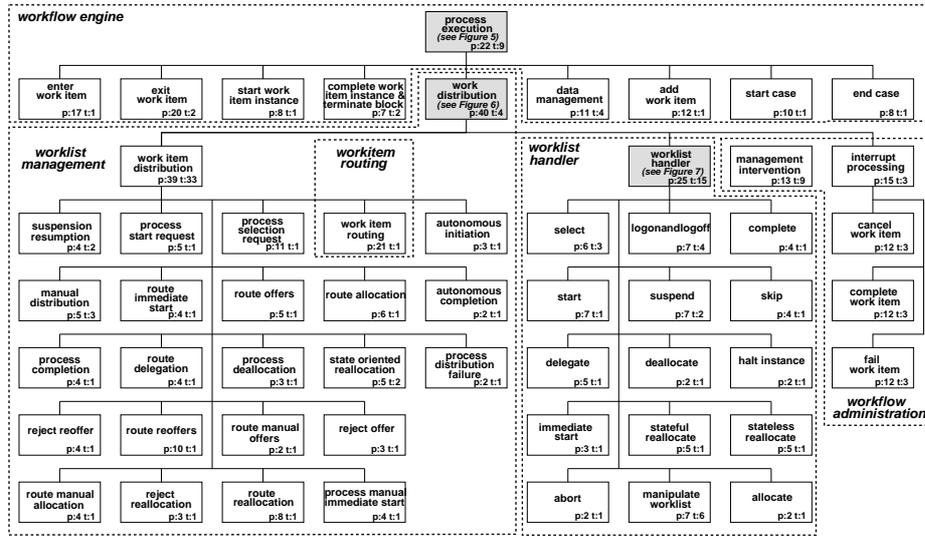


Fig. 8. *newYAWL* system CPN model hierarchy

4.1 Workflow Engine

Figure 9, which is the topmost net in the *newYAWL* model, provides a useful summary of the operation of a workflow engine. The various aspects of control-flow, data management and work distribution information which make up a static *newYAWL* specification are encoded in the CPN model as tokens in individual places. The top level view of the lifecycle of a process instance is indicated by the transitions in this diagram connected by the thick black line. First a new process instance is started, then there is a succession of **enter**→**start**→**complete**→**exit** transitions which fire as individual task instances are enabled, the work items associated with them are started and completed and the task instances are finalised before triggering subsequent tasks in the process model. Each atomic work item needs to be routed to a suitable resource for execution, an act which occurs via the **work distribution** transition. This cycle repeats until the last task instance in the process is completed. At this point, the process instance is terminated via the **end case** transition. There is provision for data interchange between the process instance and the environment via the **data management** transition. Finally, where a process model supports task concurrency via multiple work item instances, there is provision for the dynamic addition of work items via the **add** transition.

The major data items shared between the activities which facilitate the process execution lifecycle are shown as shared places in this diagram. Not surprisingly, this includes both *static* elements which describe characteristics of

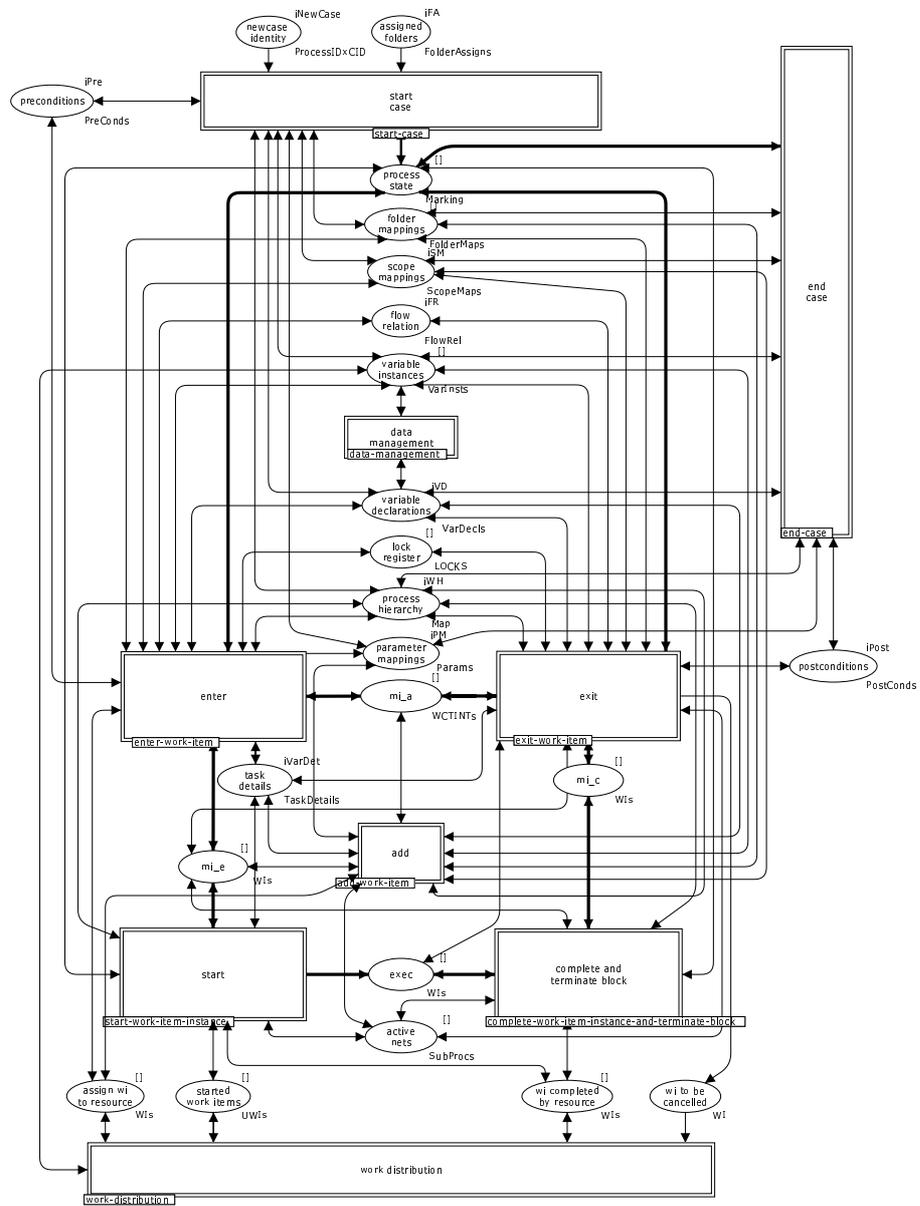


Fig. 9. Overview of the *newYAWL* workflow engine

individual processes such as the flow relation, task details, variable declarations, parameter mappings, preconditions, postconditions, scope mappings and the hierarchy of process and subprocess definitions which make up an overall process model, all of which remain unchanged during the execution of particular instances of the process. It also includes *dynamic* elements which describe how an individual process instance is being enacted at any given time. These elements are commonly known as the *state* of a process instance and include items such as the current marking of the place in the flow relation, variable instances and their associated values, locks which restrict concurrent access to data elements, details of subprocesses currently being enacted, folder mappings (identifying shared data folders assigned to a process instance) and the current execution state of individual work items (e.g. *enabled*, *started* or *completed*).

There is relatively tight coupling between the places and transitions in Fig. 9, illustrating the close integration that is necessary between the various aspects of the control-flow and data perspectives in order to enact a process model. The coupling between these places and the **work distribution** transition however is much looser. There are no static aspects of the process that are shared with other transitions in the model (i.e. the transitions underpinning **work distribution**) and other than the places which serve to communicate work items being distributed to resources for execution (and being started, completed or cancelled), the **variable instances** place is the only aspect of dynamic data that is shared with the work distribution subprocess. This reflects the functional independence of the *workflow engine*, *work item routing* and *worklist management* components. The next section looks at the issue of worklist management in more detail.

4.2 Worklist Management

The main motivation for workflow systems is achieving more effective and controlled distribution of work. Hence the actual distribution and management of work items are of particular importance. The process of managing the distribution of work items to resources is summarised by Fig. 10. It coordinates the interaction between the *workflow engine*, *work item routing*, *worklist handler* and *workflow administration* components.

The correspondences between these components and the transitions in Fig. 10 can be summarised as follows:

- the *worklist management* component is facilitated by the **work item distribution** transition, which handles the overall management of work items through the distribution and execution process (note that it subsumes the *work item routing* component);
- the *worklist handler* component corresponds to the **worklist handler** transition, which is the user-facing client software that advises users of work items requiring execution and manages their interactions with the main **work item distribution** transition in regard to committing to execute specific work items, starting and completing them;

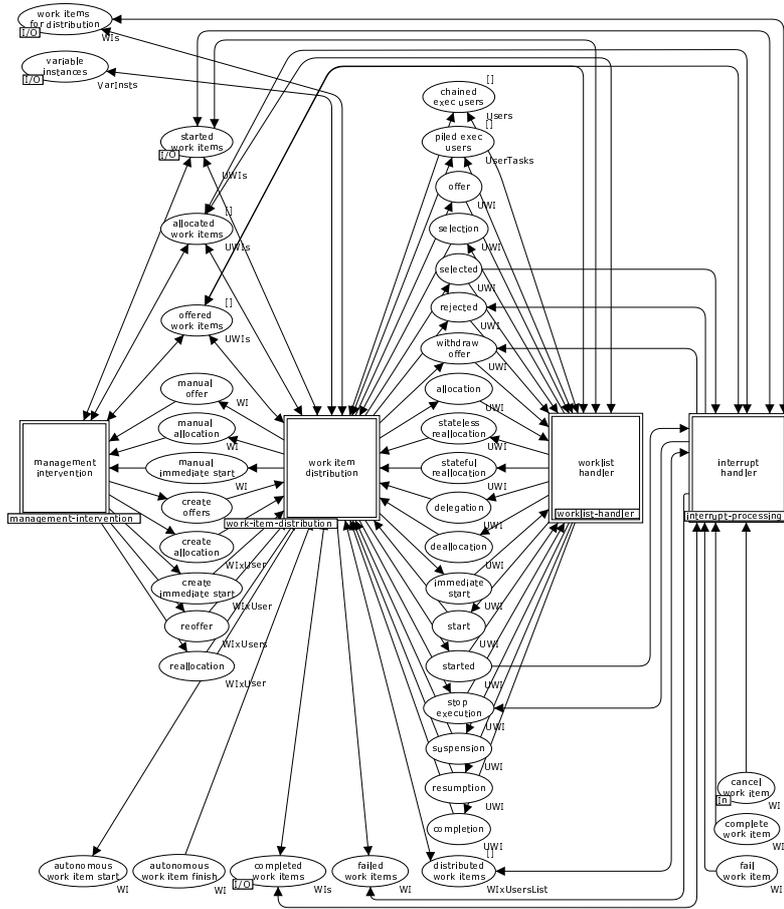


Fig. 10. Top level view of the *worklist management* component

- the *workflow administration* component is facilitated via two distinct transitions: the **management intervention** transition, that provides the ability for a workflow administrator to intervene in the **work distribution** process and manually reassign work items to users where required; and the **interrupt handler** transition that supports the cancellation, forced completion and forced failure of work items as may be triggered by other components of the workflow engine (e.g. the control-flow process, exception handlers).

Work items that are to be distributed are communicated between the *workflow engine* and the *worklist management* components via the **work items for distribution** place. This then prompts the **work item distribution** transition to determine how they should be routed for execution. This may involve the services of the workflow administrator in which case they are sent to the

`management intervention` transition or alternatively they may be sent directly to one or more resources via the `worklist handler` transition. The various places between these three transitions correspond to the range of requests that flow between them. In the situation where a work item corresponds to an *automatic* task, it is sent directly to the `autonomous work item start` place and no further distribution activities take place. An automatic task is considered complete when a token is inserted in the `autonomous work item finish` place.

A common view of work items in progress is maintained for the `work item distribution`, `worklist handler`, `management intervention` and `interrupt handler` transitions via the `offered work items`, `allocated work items` and `started work items` places (although obviously this information is only available to the *worklist handler* when it is actually connected to the workflow management system). There is also shared information about users in advanced operating modes that is recorded in the `pled exec users` and `chained exec users` places. Although there is significant provision for shared information about the state of work items, the determination of when a work item is actually complete rests with the `work item distribution` transition and when this occurs, it inserts a token in the `completed work items` place. Similarly, work item failures are notified via the `failed work items` place. The only exception to these arrangements are for work items that are subject to some form of interrupt (e.g. an exception being detected and handled). The `interrupt handler` transition is responsible for managing these occurrences on the basis of cancellation, forced completion and failure requests received in the `cancel work item`, `complete work item` and `fail work item` places respectively. All of the activities in the *worklist management* component are illustrated by substitution transitions indicating that each of them are defined in terms of significantly more complex subprocesses. It is not possible to present each of them in this paper. Finally we focus on one other significant component: the *worklist handler*.

4.3 Worklist Handler

The *worklist handler* component is illustrated in Fig. 11 and describes how the user-facing workflow interface (typically a worklist handler software client) operates and interacts with the *worklist management* component. The main path through this process is indicated by the thick black arcs. There are various transitions that make up the process, these correspond to actions that individual users can request in order to alter the current state of a work item to more closely reflect their current handling of it. These actions may simply be requests to start or complete it or they may be “detour” requests to reroute it to other users e.g. via `delegation` or `deallocation`. The manner in which these requests operate is illustrated by the shared places in Fig. 10. Typically the inclusion of a request in one of these shared places results in a message flowing between the *worklist handler* and *worklist management* components which ultimately causes the relative states of the two components to be synchronised.

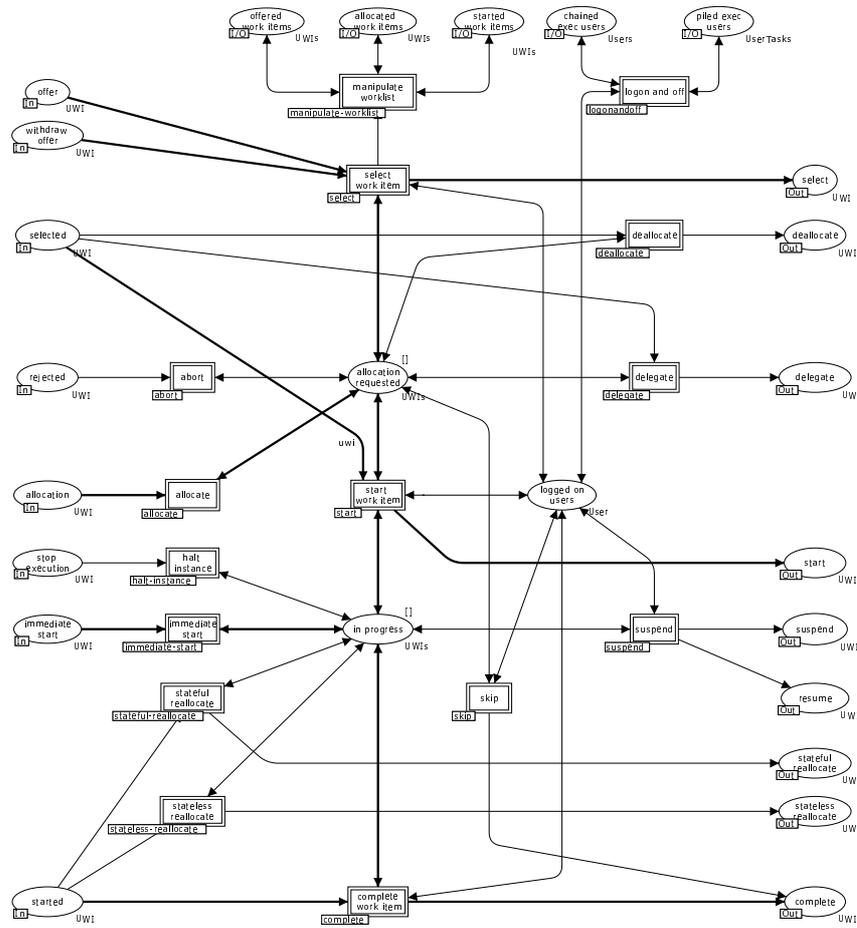


Fig. 11. Worklist handler component

5 Related Work

There have been numerous papers advocating approaches to workflow and business process modelling based on Petri Nets (cf. [10,11,12,13]), however these tend to either focus on a single aspect of the domain (e.g. the control-flow perspective) or they are based on a relatively simplistic language. There have also been attempts to provide formal semantics using Petri Nets for many of the more widely used approaches to business process modelling including EPCs [14], UML 2.0 Activity Diagrams [15] and BPMN [16], although in each case arriving at a complete semantics has been hampered by inherent ambiguities in the informal descriptions for each of the formalisms. There has been minimal work on formalisation of the other workflow perspectives, one exception is [17] which

investigates mechanisms for work distribution in workflows and presents CPN models for a number of the workflow resource patterns.

Historically, the modelling and enactment of processes have often been treated distinctly and it is not unusual for separate design and runtime models to be utilised by systems. Approaches to managing the potential disparities between these models have included the derivation of executable process descriptions from design-time models [18] and the direct animation of design-time models for requirements validation [19]. The latter of these approaches which uses a strategy based on Coloured Petri Nets [8] and CPN Tools [20] as an enablement vehicle is one of a number of initiatives that have successfully used the CPN Tools offering as a means of executing various design-time modelling formalisms including Protos models [21], sequence diagrams [22] and task descriptions [23].

There has been a significant body of work that describes software architectures for workflow management systems. Significant examples of such systems include MOBILE [24], WIDE [25], CrossFlow [26] and WAMO [27] amongst many others however none of these systems offer a fully formalised description both of their language elements and the overall operation of the workflow system.

6 Experiences and Conclusion

The selection of Coloured Petri Nets as the conceptual foundation for *newYAWL* proved to be a fortuitous choice. Being state-based and graphical, the formalism delivered immediate modelling benefits as a consequence of its commonalities with the domain it was used to represent. The availability of a means of integrating the handling of the data-related aspects of the *newYAWL* language into the model (i.e. using “colour”) and partitioning the model on a hierarchical basis into units of related functionality meant that a more compact means of representing the overall *newYAWL* model was possible. The most significant advantage of this design choice however proved to be the availability of an interactive modelling and execution environment in the form of CPN Tools. Indeed, it is only with the aid of an interactive modelling environment such as CPN Tools that developing a formalisation of this scale actually becomes viable.

Although there are other candidates for developing large-scale system designs, none of them deliver the benefits inherent in the Coloured Petri Nets and CPN Tools combination. Conceptual foundations such as π -calculus and process algebra as well as software-oriented specification formalisms such as Z and VDM lack a graphical representation meaning that the visualisation and assessment of specific design choices is difficult. In contrast, lighter-weight approaches to business process modelling such as those embodied in offerings such as Protos, ARIS and other business process modelling tools do provide an intuitive approach to specifying business process that is both graphical and state-based, however they lack a complete formal semantics. Moreover, they only allow for the specification of a specific candidate model and do not provide a means of capturing an arbitrary range of business processes in a single model as is required for the *newYAWL* language. This shortcoming stems from the fact that the

modelling formalisms employed in these tools are control-flow centric and lack a fully fledged data perspective. High-level CASE tools (e.g. Rational Rose) share similar shortcomings and their generalist nature means that they do not provide any specific support for business process modelling and enactment initiatives.

One of the major advantages of the approach pursued in developing *newYAWL* is that it provided a design that is executable. This allowed fundamental design decisions to be evaluated and tested much earlier than would ordinarily be the case during the development process. Where suboptimal design decisions were revealed, the cost of rectifying them was significantly less than it would have been later in the development lifecycle. There was also the opportunity to test alternate solutions to design issues with minimal overhead before a final decision was settled on. A particular benefit afforded by this approach to formalisation was that the CPN hierarchy established during the design process provided an excellent basis on which to make subsequent architectural and development decisions.

Whilst complete, the resultant model *newYAWL* system model³ is extremely complex. *It incorporates 55 distinct pages of CPN diagrams and encompasses 480 places, 138 transitions and in excess of 1500 lines of ML code.* It took approximately six months to develop. The size of the model gives an indication of the relative complexity of formally specifying a comprehensive business process modelling language such as *newYAWL*. The original motivations for this research initiative were twofold: (1) to establish a fully formalised business process modelling language based on the synthesis of the workflow patterns and (2) to demonstrate that the language was not only suitable for conceptual modelling of business processes but that it also contained sufficient detail for candidate models to be directly enacted. *newYAWL* achieves both of these objectives and *directly supports 118 of the 126 workflow patterns* that have been identified. It is interesting to note however that whilst the development of a system model of this scale offers some extremely beneficial insights into the overall problem domain and provides a software design that can be readily utilised as the basis for subsequent programming activities, it also has its limitations. Perhaps the most significant of these is that the scale and complexity of the model obviates any serious attempts at verification. Even on a relatively capable machine (P4 2.1Ghz dual-core, 2Gb RAM), it takes almost 4 minutes just to load the model. Moreover the potentially infinite range of business process models that the *newYAWL* system can encode, rules out the use of techniques such as state space analysis. This raises the question as to how models of this scale can be comprehensively tested and verified.

Notwithstanding these considerations however, the development of the new-YAWL system model delivered some salient insights into areas of newYAWL that needed further consideration during the design activity. These included:

- the introduction of a deterministic mechanism for recording status changes in the work item execution lifecycle in order to ensure that the views of

³ This model is available at www.yawl-system.com/newYAWL.

- these details maintained by the *worklist management* and *worklist handler* components are consistent;
- the establishment of a coherence protocol to ensure that reallocation of work items to alternate resources either by resources themselves or the workflow administrator are handled in a consistent manner in order to ensure that potential race conditions arising during reallocation do not result in the workflow engine, workflow administrator or the initiating resource (i.e. *worklist handler*) having irreconcilable views of the current state of work item allocations;
 - the introduction of a consistent approach for handling the evaluation of any functions associated with a *newYAWL* specification e.g. for outgoing links in an XOR-split, pre/postconditions, pre/post tests for iterative tasks etc. This issue was ultimately addressed by mapping any necessary function calls to ML functions and establishing a standard approach to encoding the invocation of these functions and the passing of any necessary parameters and the return of associated results;
 - adoption of a standard strategy for characterising parameters to functions in order to ensure that they could be passed in a uniform way to the associated ML functions that evaluated them;
 - the introduction of a locking strategy for data elements to prevent inadvertent side-effects of concurrent data usage; and
 - recognition that when a self-cancelling task completes: (1) it should process the cancellation of itself last of all in order to prevent the situation where it cancels itself before all other cancellations have been completed and (2) it needs to establish whether it is cancelling itself before it can make the decision to put tokens in any relevant output places associated with the task.

The *newYAWL* system model provides a complete description of an operational environment for the *newYAWL* language. It is sufficiently detailed to be directly useful for system design and development activities. It will serve as the design blueprint for upcoming versions of the open-source YAWL System. In fact the resource management component of the *newYAWL* language has already been incorporated in the YAWL System.

Acknowledgement The authors would like to thank the anonymous reviewers for their constructive comments and suggestions.

References

1. Russell, N., ter Hofstede, A.H.M., van der Aalst, W.M.P.: *newYAWL*: Specifying a workflow reference language using Coloured Petri Nets. In: Proceedings of the Eighth Workshop and Tutorial on Practical Use of Coloured Petri Nets and the CPN Tools. Number DAIMI PB-584, Department of Computer Science, University of Aarhus, Denmark (2007) 107–126

2. van der Aalst, W.M.P., ter Hofstede, A.H.M.: YAWL: Yet another workflow language. *Information Systems* **30**(4) (2005) 245–275
3. van der Aalst, W.M.P., ter Hofstede, A.H.M., Kiepuszewski, B., Barros, A.: Workflow patterns. *Distributed and Parallel Databases* **14**(3) (2003) 5–51
4. Russell, N., ter Hofstede, A.H.M., van der Aalst, W.M.P., Mulyar, N.: Workflow control-flow patterns: A revised view. Technical Report BPM-06-22 (2006) <http://www.BPMcenter.org>.
5. Wynn, M., Edmond, D., van der Aalst, W.M.P., ter Hofstede, A.H.M.: Achieving a general, formal and decidable approach to the OR-join in workflow using Reset nets. In Ciardo, G., Darondeau, P., eds.: *Proceedings of the 26th International Conference on Application and Theory of Petri nets and Other Models of Concurrency (Petri Nets 2005)*. Volume 3536 of LNCS., Miami, USA, Springer-Verlag (2005) 423–443
6. Russell, N., ter Hofstede, A.H.M., Edmond, D., van der Aalst, W.M.P.: Workflow data patterns: Identification, representation and tool support. In Delcambre, L., Kop, C., Mayr, H., Mylopoulos, J., Pastor, O., eds.: *Proceedings of the 24th International Conference on Conceptual Modeling (ER 2005)*. Volume 3716 of LNCS., Klagenfurt, Austria, Springer (2005) 353–368
7. Russell, N., van der Aalst, W.M.P., ter Hofstede, A.H.M., Edmond, D.: Workflow resource patterns: Identification, representation and tool support. In Pastor, O., Falcão e Cunha, J., eds.: *Proceedings of the 17th Conference on Advanced Information Systems Engineering (CAiSE'05)*. Volume 3520 of LNCS., Porto, Portugal, Springer (2005) 216–232
8. Jensen, K.: *Coloured Petri Nets. Basic Concepts, Analysis Methods and Practical Use*. Volume 1, Basic Concepts. Monographs in Theoretical Computer Science. Springer-Verlag, Berlin, Germany (1997)
9. Russell, N., ter Hofstede, A.H.M., Edmond, D., van der Aalst, W.M.P.: *newYAWL: achieving comprehensive patterns support in workflow for the control-flow, data and resource perspectives*. Technical Report BPM-07-05 (2007) <http://www.BPMcenter.org>.
10. van der Aalst, W.M.P.: The application of Petri nets to workflow management. *Journal of Circuits, Systems and Computers* **8**(1) (1998) 21–66
11. Ellis, C., Nutt, G.: Modelling and enactment of workflow systems. In Marsan, M.A., ed.: *Proceedings of the 14th International Conference on Application and Theory of Petri Nets*. Volume 691 of LNCS., Chicago, IL, USA, Springer (1993) 1–16
12. Adam, N., Atluri, V., Huang, W.: Modeling and analysis of workflows using Petri nets. *Journal of Intelligent Information Systems* **10**(2) (1998) 131–158
13. Moldt, D., Rölke, H.: Pattern based workflow design using reference nets. In van der Aalst, W., ter Hofstede, A., Weske, M., eds.: *Proceedings of the Business Process Management Conference 2003*. Volume 2678 of LNCS., Eindhoven, The Netherlands, Springer (2003) 246–260
14. van der Aalst, W.M.P.: Formalization and verification of event-driven process chains. *Information and Software Technology* **41**(10) (1999) 639–650
15. Störrle, H., Hausmann, J.: Towards a formal semantics of UML 2.0 activities. In Liggesmeyer, P., Pohl, K., Goedicke, M., eds.: *Proceedings of the Software Engineering 2005, Fachtagung des GI-Fachbereichs Softwaretechnik*. Volume 64 of *Lecture Notes in Informatics.*, Essen, Germany, Gesellschaft für Informatik (2005) 117–128
16. Dijkman, R., Dumas, M., Ouyang, C.: Semantics and analysis of business process models in BPMN. *Information and Software Technology* **50**(12) (2008) 1281–1294

17. Pesic, M., van der Aalst, W.M.P.: Modelling work distribution mechanisms using colored Petri nets. *International Journal on Software Tools for Technology Transfer* **9**(3) (2007) 327–352
18. Di Nitto, E., Lavazza, L., Schiavoni, M., Tracanella, E., Trombetta, M.: Deriving executable process descriptions from UML. In: *ICSE '02: Proceedings of the 24th International Conference on Software Engineering*, New York, NY, USA, ACM Press (2002) 155–165
19. Machado, R., Lassen, K., Oliveira, S., Couto, M., Pinto, P.: Requirements validation: Execution of UML models with CPN tools. *International Journal on Software Tools for Technology Transfer* **9**(3) (2007) 353–369
20. Jensen, K., Kristensen, L., Wells, L.: Coloured Petri nets and CPN tools for modelling and validation of concurrent systems. *International Journal of Software Tools for Technology Transfer* **9**(3) (2007) 213–254
21. Gottschalk, F., van der Aalst, W., Jansen-Vullers, M., Verbeek, H.: Protos2CPN: Using colored Petri nets for configuring and testing business processes. In Jensen, K., ed.: *Proceedings of the 7th Workshop and Tutorial on Practical Use of Coloured Petri Nets and the CPN Tools*. Volume PB-579 of DAIMI Reports., Aarhus, Denmark (2006) 137–155
22. Ribeiro, O., Fernandes, J.: Some rules to transform sequence diagrams into coloured Petri nets. In Jensen, K., ed.: *Proceedings of the 7th Workshop and Tutorial on Practical Use of Coloured Petri Nets and the CPN Tools*. Volume PB-579 of DAIMI Reports., Aarhus, Denmark (2006) 137–155
23. J.B. Jørgensen, K.L., van der Aalst, W.M.P.: From task descriptions via coloured Petri nets towards an implementation of a new electronic patient record. In Jensen, K., ed.: *Proceedings of the 7th Workshop and Tutorial on Practical Use of Coloured Petri Nets and the CPN Tools*. Volume PB-579 of DAIMI Reports., Aarhus, Denmark (2006) 137–155
24. Jablonski, S., Bussler, C.: *Workflow Management: Modeling Concepts, Architecture and Implementation*. Thomson Computer Press, London, UK (1996)
25. Ceri, S., Grefen, P., Sanchez, G.: WIDE: a distributed architecture for workflow management. In: *Proceedings of the Seventh International Workshop on Research Issues in Data Engineering (RIDE'97)*, Birmingham, England, IEEE Computer Society Press (1997)
26. Ludwig, H., Hoffner, Y.: Contract-based cross-organisational workflows - the Cross-Flow project. In Grefen, P., Bussler, C., Ludwig, H., Shan, M., eds.: *Proceedings of the WACC Workshop on Cross-Organisational Workflow Management and Coordination*, San Francisco (1999)
27. Eder, J., Liebhart, W.: The workflow activity model (WAMO). In Laufmann, S., Spaccapietra, S., Yokoi, T., eds.: *Proceedings of the Third International Conference on Cooperative Information Systems (CoopIS-95)*, Vienna, Austria, University of Toronto Press (1995) 87–98