

XRL/Woflan: Verification of an XML/Petri-net-based language for inter-organizational workflows

H.M.W. VERBEEK*, W.M.P. VAN DER AALST

h.m.w.verbeek@tue.nl

Faculty of Technology Management, Eindhoven University of Technology, the Netherlands

A. KUMAR

akhil.kumar@colorado.edu

College of Business, CB 419, University of Colorado, Boulder, Co 80309, USA

Abstract. Internet-based technology, E-commerce, and the rise of networked virtual enterprises have fueled the need for inter-organizational workflows. Although XML allows trading partners to exchange information, it cannot be used to coordinate activities in different organizational entities. Therefore, we developed a workflow language named XRL (eXchangeable Routing Language) for supporting cross-organizational processes. XRL uses XML for the representation of process definitions and Petri nets for its semantics. Since XRL is instance-based, workflow definitions can be changed on the fly and sent across organizational boundaries. These features are vital for today's dynamic and networked economy. However, the features also enable subtle, but highly disruptive, cross-organizational errors. On-the-fly changes and one-of-a-kind processes are destined to result in errors. Moreover, errors of a cross-organizational nature are difficult to repair. In this paper, we present XRL/Woflan. XRL/Woflan is a software tool using state-of-the-art Petri-net analysis techniques for verifying XRL workflows. The tool uses *eXtensible Stylesheet Language Transformations (XSLT)* to transform XRL specifications to a specific class of Petri nets. The Petri-net representation is used to determine whether the workflow is correct. If the workflow is not correct, anomalies such as deadlocks and livelocks are reported.

Keywords: Workflow, Inter-organizational, verification, extensibility, XML, WF-net, XRL, Woflan

1. Introduction

Today's corporations often must operate across organizational boundaries. Phenomena such as E-commerce, extended enterprises, and the Internet stimulate cooperation between organizations. Therefore, the importance of workflows distributed over a number of organizations is increasing [3, 4, 19, 26]. Inter-organizational workflow offers companies the opportunity to re-shape business processes beyond the boundaries of their own organizations. However, inter-organizational workflows are typically subject to conflicting constraints. On the one hand, there is a strong need for coordination to optimize the flow of work in and between the different organizations. On the other hand, the organizations involved are essentially autonomous and have the freedom to create or modify workflows at any point in time. These conflicting constraints complicate the development of languages and tools for cross-organizational workflow support.

Recent development in Internet technology, and the emergence of the "electronic market makers", such as ChemConnect, Ariba, CommerceOne, Clarus, staples.com, Granger.com, VerticalNet, and mySAP.com have resulted in many XML-based standards for electronic commerce. The XML Common Business Library (xCBL) by CommerceOne, the Partner Interface Process (PIP) blueprints by RosettaNet, the Universal Description, Discovery and Integration (UDDI), the Electronic Business XML (ebXML) initiative by UN/CEFACT and OASIS, the Open Buying on the Internet (OBI) specification, the Open Application Group Integration Specification (OAGIS), and the BizTalk Framework are just some examples of the emerging standards based on XML. *These standards primarily focus on the exchange of data and not on the control flow among organizations.* Most of the standards provide standard *Document Type Definitions (DTDs)* or XML schemas for specific application domains (such as procurement). One of the few initiatives that also address the control flow is RosettaNet. The Partner Interface Process (PIP)

* Corresponding author

blueprints by RosettaNet do specify interactions using UML activity diagrams for the Business Operational View (BOV) and UML sequence diagrams for the Functional Service View (FSV) in addition to DTDs for data exchange. However, the PIP blueprints are not executable and need to be predefined. Moreover, like most of the standards, RosettaNet is primarily focusing on electronic markets with long-lasting pre-specified relationships between parties with one party (such as the market maker) imposing rigid business rules.

Looking at existing initiatives, it can be noted that:

- (1) process support for cross-organizational workflow has been neglected since lion's share of attention has gone to data and
- (2) only pre-specified standardized processes have been considered (such as, market places, procurement, and so on).

Based on these observations, we developed the *eXchangeable Routing Language* (XRL). The idea to develop a language like XRL was raised in [24] and the definition of the language was given in [8]. XRL uses the syntax of XML, but contains constructs that embed the semantics of control flow. Moreover, XRL supports highly dynamic one-of-a-kind workflow processes. For example, we consider the “first trade problem,” that is, the situation where parties have no prior trading relationship [27]. Clearly, the “first-trade problem” is the extreme case of highly dynamic one-of-a-kind workflow processes and therefore also the most difficult. To support highly dynamic one-of-a-kind workflow processes, XRL describes processes at the instance level. Traditional workflow modeling languages describe processes at the class or type level [21, 25]. An XRL routing schema describes the partial ordering of tasks for one specific instance. The advantages of doing so are that:

- (1) the workflow schema can be exchanged more easily,
- (2) the schema can be changed without causing any problems for other instances, and
- (3) the expressive power is increased.

Note that workflow-modeling languages typically have problems handling a variable number of parallel or alternative branches [6]. In our research on workflow patterns [6], we compared the expressive power of many contemporary workflow management systems including COSA, HP Changengine, Forté Conductor, I-Flow, InConcert, MQ Series Workflow, R/3 Workflow, Staffware, Verve, and Visual WorkFlo using a set of workflow patterns (See <http://www.tm.tue.nl/it/research/patterns/>). Based on the workflow patterns supported by these systems, and their relative use in practice, we carefully selected the most relevant constructs for XRL. Note that the expressive power of XRL far exceeds that of each of the workflow management systems mentioned above.

As was shown in [8], the semantics of XRL can be expressed in terms of Petri nets [32, 33]. Based on these semantics, we developed a workflow management system, named *XRL/Flower*, to support XRL. XRL/Flower benefits from the fact that it is based on both XML and Petri nets. Standard XML tools can be deployed to parse, check, and handle XRL documents. The Petri-net representation allows for a straightforward and succinct implementation of the workflow engine. XRL constructs are automatically transformed into Petri-net constructs. On the one hand, this allows for an efficient implementation. On the other hand, the system is easy to extend:

For supporting a new routing primitive, only the transformation to the Petri net format needs to be added and the engine itself does not need to change.

Unfortunately, the Petri-net-based semantics of XRL given in [8] results in Petri nets that do not fit into the class of so-called *Workflow nets* (WF-nets). WF-nets are a special subclass of Petri nets which possess an appealing correctness notion (the soundness property [1]), are based on strong theoretical results (such as, the link between soundness, liveness, and boundedness [1]), and are supported by powerful software (such as, the tool Woflan [38]). The semantics given in [8] does not allow for a direct use of these theoretical

results and tools. This limitation was recognized in [9]. In this paper, we present a direct transformation from XRL to WF-nets, that is, the semantics of XRL is given in terms of WF-nets. The transformation has been implemented in XSLT (eXtensible Stylesheet Language Transformations) and resulted in the tool XRL/Woflan.

XRL/Woflan builds on the workflow verification tool Woflan [37, 38]. Developers of contemporary workflow management systems have virtually neglected correctness issues. As a result, in most workflow management systems, it is possible to design workflows which suffer from anomalies such as deadlocks and livelocks without any form of warning. Few tools provide any form of workflow verification support. The tools Woflan [38] and Flowmake [36] are two noteworthy exceptions. To complicate matters, more and more workflow management systems are used to support inter-organizational business processes, for example, in the context of Business-To-Business (B2B) E-commerce. Especially for open E-commerce (that is, doing business among parties having no prior trading relationship), the workflow support should be trustworthy in the sense that trading partners who do not know each other, and may even come from different countries and cultures, may conduct business with the assurance that their interests will be protected in the event that “things go wrong”, whether by accident, negligence, or intentional fraud [27]. One of the prerequisites for this is the guarantee that the workflow process definitions do not contain any logical errors. Therefore, XRL/Woflan, the verification tool presented in this paper, is highly relevant for developers of inter-organizational workflows.

Figure 1 shows the architecture of the toolset involving XRL/Flower and XRL/Woflan. Using both the control flow data for the workflow case and the case specific data, the Petri-net engine computes the set of enabled tasks, that is, the set of *work items* that are *ready*. The engine sends this set to the work

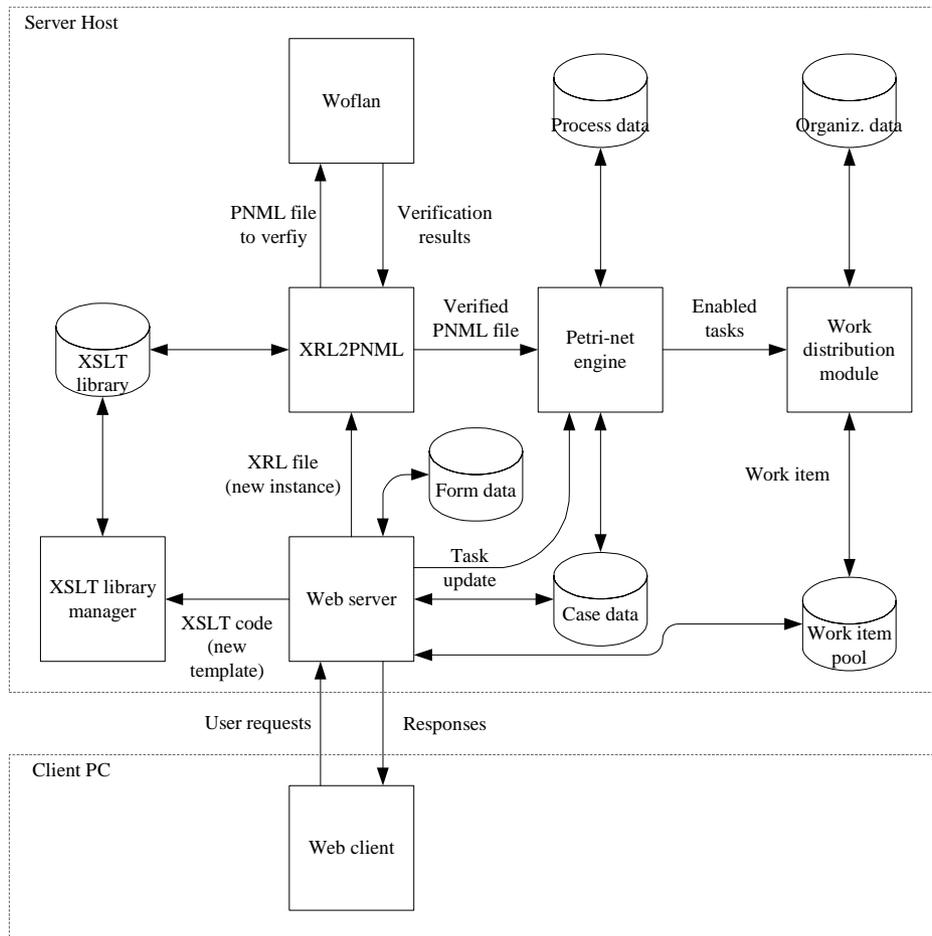


Figure 1. A detailed architecture for implementing inter-organizational workflows using XRL/Woflan.

distribution module. Based on information on the organizational roles and users, the work distribution module sends e-mails offering the work item to certain users who are qualified to work on it. A user would receive an e-mail notification with a URL pointing to new work item(s) waiting for her. By clicking on the URL, the user accepts the work item; thus, the *work item* becomes an *activity* for a specific user, and other users to whom the work item was also offered are notified that it has already been accepted and is no longer available to them. A user who has accepted an activity may perform work on it either at acceptance time or later. In order to enable a user to perform an activity, the web server fills the appropriate form template with the case specific data for the activity. The user indicates completion of an activity by, say, pressing a submit button. The web server stores the updated case data and signals the Petri-net engine that the activity has been completed. The Petri-net engine then recomputes a new set of work items that are ready. The user can also start an XRL instance by sending the corresponding XRL file to the web server. The web server forwards the XRL file to the XRL2PNML module that transforms XRL to PNML (Petri-Net Markup Language), which is a standard representation language for a Petri net in XML format [22].

Figure 2 shows how the XRL2PNML module makes the transformation from XRL to PNML. First, it transforms the XRL file to two PNML files: one for verification and one for enactment. The first PNML file (for verification) can be considerably smaller in size than the second (for enactment), but their soundness characteristics are the same. The first PNML file is verified using the XRL/Woflan tool. Based on the result of the verification, either the second PNML file is sent to the Petri-net engine for enactment, or the user is informed that the XRL instance contains flaws. In the latter case, the user may either abandon the new instance, or modify it to fix the errors. Of course, the fixed instance is also verified before it is enacted. If the expressive power of the current XSLT library does not satisfy the user's needs, she may decide to extend this library by adding a new template to it. Figure 3 shows how this is done. First, the user describes the DTD of the new pattern. Second, she writes the XSLT code that will transform the new template to the appropriate PNML code. After the new XSLT code is verified, it is incorporated into the XSLT library.

The remainder of this paper is organized as follows. Section 2 introduces XRL and Section 3 gives an example of how a workflow can be represented in XRL. Section 4 introduces WF-nets. Then Section 5 provides the formal semantics of XRL in terms of WF-nets. Based on these semantics, Section 6 discusses the soundness of XRL routes, while Section 7 proposes a verification procedure that exploits the structural properties of certain XRL constructs and Petri-net based reduction rules [32]. Section 8 demonstrates the

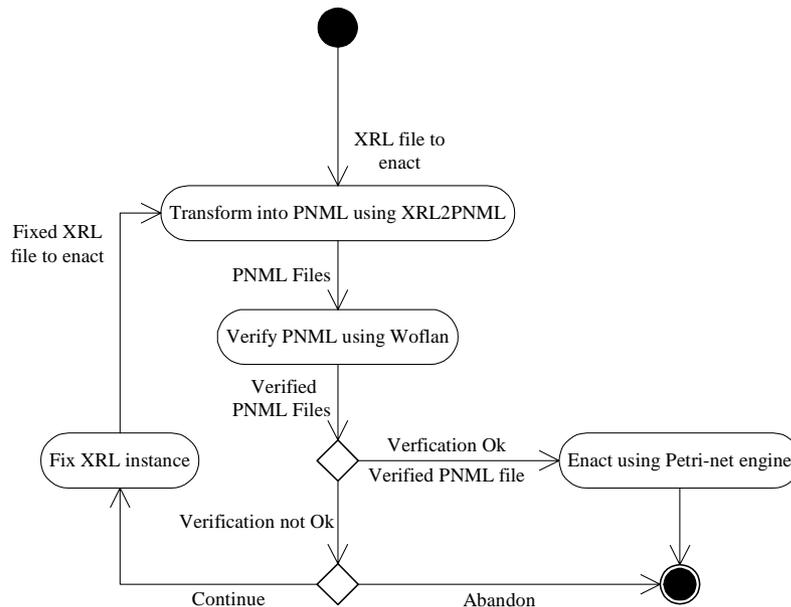


Figure 2. Transformation from XRL to PNML.

extensibility of XRL by showing how new constructs may be added to XRL. Section 9 presents our tool XRL/Woflan. Section 10 relates this paper to known research. Section 11 concludes the paper. Appendix A shows the DTD of XRL after the extensions from Section 8 have been added. Appendix B shows the XRL route for processing a customer order that is introduced in Section 3. Appendix C lists the complete XSLT specification of the extensions mentioned in Section 8, using the macros as defined in Appendix D. These macros were introduced to keep the verbose XSLT files as compact and readable as possible. The resulting ('macrofied') XSLT specifications are listed in Appendix E.

2. XRL: An XML based routing language

The focus of this paper is on verification and extensibility. Therefore, we limit ourselves to only a brief introduction to XRL and the workflow management system XRL/Flower. The syntax of XRL is completely specified by the DTD [15] shown in Figure 4. An *XRL route* is a consistent XML document, that is, a well-formed and valid XML file with top element *route* (see Figure 4).

The structure of any XML document forms a tree. In case, of XRL, the root element of that tree is the route. This route contains exactly one so-called *routing element*. A routing element (RE) is an important building block of XRL. It can either be simple (no child routing elements) or complex (one or more child routing elements). A complex routing element specifies *whether*, *when* and in *which order* the child routing elements are done.

XRL provides the following routing elements:

Task: Offer the given step to some resource, wait until the step has been performed, and afterwards set all events for which a child event element exists.

Sequence: Start the child routing elements in the given order and wait until all have been performed.

Any_sequence: Start the child routing elements in any order and wait until all have been performed.

Choice: Start one of the child routing elements and wait until it has been performed.

Condition: If the given condition holds, start the child routing elements of all *true* child elements in parallel and wait until all have been performed. Otherwise, start the child routing elements of all *false* child

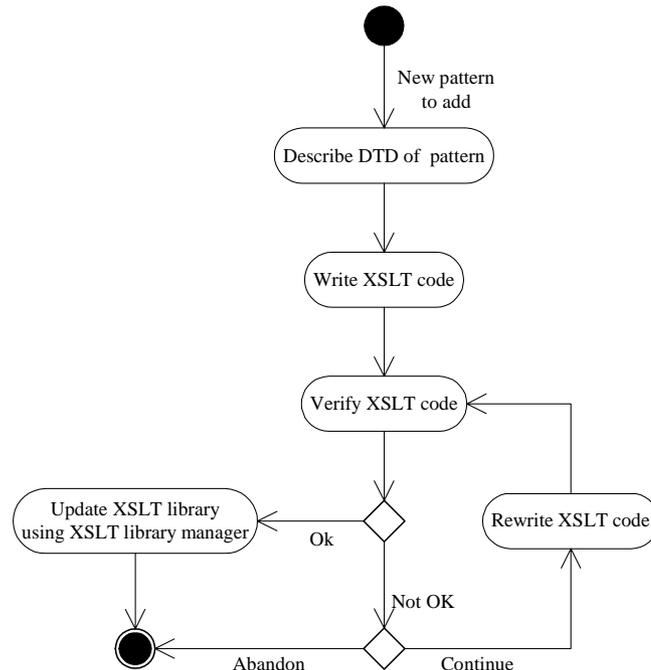


Figure 3. Adding a new template to the library.

```

<!ENTITY % routing_element
"task|sequence|any_sequence|choice|condition|parallel_sync|parall
el_no_sync|parallel_part_sync|parallel_part_sync_cancel|wait_all|
wait_any|while_do|terminate">
<!ELEMENT route ((%routing_element;), event*)>
<!ATTLIST route
  name ID #REQUIRED
  created_by CDATA #IMPLIED
  date CDATA #IMPLIED>
<!ELEMENT task (event*)>
<!ATTLIST task
  name ID #REQUIRED
  address CDATA #REQUIRED
  role CDATA #IMPLIED
  doc_read NMTOKENS #IMPLIED
  doc_update NMTOKENS #IMPLIED
  doc_create NMTOKENS #IMPLIED
  result CDATA #IMPLIED
  status (ready|running|enabled|disabled|aborted|null) #IMPLIED
  start_time NMTOKENS #IMPLIED
  end_time NMTOKENS #IMPLIED
  notify CDATA #IMPLIED>
<!ELEMENT event EMPTY>
<!ATTLIST event
  name ID #REQUIRED>
<!ELEMENT sequence ((%routing_element;|state)+)>
<!ELEMENT any_sequence ((%routing_element;)+)>
<!ELEMENT choice ((%routing_element;)+)>
<!ELEMENT condition ((true|false)*)>
<!ATTLIST condition
  condition CDATA #REQUIRED>
<!ELEMENT true (%routing_element;)>
<!ELEMENT false (%routing_element;)>
<!ELEMENT parallel_sync ((%routing_element;)+)>
<!ELEMENT parallel_no_sync ((%routing_element;)+)>
<!ELEMENT parallel_part_sync ((%routing_element;)+)>
<!ATTLIST parallel_part_sync
  number NMTOKEN #REQUIRED>
<!ELEMENT parallel_part_sync_cancel ((%routing_element;)+)>
<!ATTLIST parallel_part_sync_cancel
  number NMTOKEN #REQUIRED>
<!ELEMENT wait_all ((event_ref|timeout)+)>
<!ELEMENT wait_any ((event_ref|timeout)+)>
<!ELEMENT event_ref EMPTY>
<!ATTLIST event_ref
  name IDREF #REQUIRED>
<!ELEMENT timeout ((%routing_element;)?)>
<!ATTLIST timeout
  time CDATA #REQUIRED
  type (relative|s_relative|absolute) "absolute">
<!ELEMENT while_do (%routing_element;)>
<!ATTLIST while_do
  condition CDATA #REQUIRED>
<!ELEMENT terminate EMPTY>
<!ELEMENT state EMPTY>

```

Figure 4. The DTD of XRL

elements in parallel and wait until all have been performed. A condition may have any number (even none) of *true* and *false* child elements.

Parallel_sync: Start the child routing elements in parallel and wait until all have been performed.

Parallel_no_sync: Start the child routing elements in parallel but don't wait for any of them.

Parallel_part_sync: Start the child routing elements in parallel and wait until the given number of child routing elements has been performed.

Parallel_part_sync_cancel: Start the child routing elements in parallel, wait until the given number of child routing elements has been performed and cancel the remaining child routing elements if possible.

Wait_all: Wait until either all events for which an eventref child exists are set, or wait until the given deadline of some child timeout element has expired. If this timeout element has a child routing element, start it and wait until it has been performed.

Wait_any: Wait until either at least one of the events for which an eventref child exists is set, or wait until the given deadline of some child timeout element has expired. If this timeout element has a child routing element, start it and wait until it has been performed.

While_do: As long as the given condition holds, start the child routing element and wait until it has been performed.

Terminate: End this workflow instance.

As mentioned before, the routing elements of XRL are based on a thorough analysis of the workflow patterns supported by leading workflow management systems.

In the next section, we show how the routing elements of XRL can be used to design a real workflow, which is consistent with the DTD.

3. Example: An electronic bookstore

In this section we illustrate XRL using an example inspired by electronic bookstores, such as Amazon [12] and Barnes and Noble [13]. The activity diagram in Figure 5 shows a typical order flow. This figure gives the four parties or organizations involved (that is, customer, bookstore, publisher and shipper), and the steps performed by each one. The arrows show the sequence in which these steps are carried out. Some of the details of the real-world process are omitted from this diagram for clarity.

The workflow represented by the sequence diagram is described in XRL in Appendix B. The XRL rendition covers the typical order flow of Figure 5, and also some more details. First, the customer places an order (task *place_c_order*). This customer order is sent to and handled by the bookstore (task *handle_c_order*). The electronic bookstore is a virtual company that has no books in stock. Therefore, the bookstore transfers the order for the desired book to the first appropriate publisher (task *place_b_order*). We use the term “bookstore order” for the transferred order. The publisher evaluates the bookstore order (task *eval_b_order*). By shipping the XRL route back to the bookstore, the publisher informs the bookstore about the availability of the book. If the book is not available, the bookstore decides (task *decide*) to either search for an alternative publisher (task *alt_publ*) or to reject the customer (task *c_reject*). If the customer receives a negative answer (task *rec_decl*), the workflow terminates. If the book is available (task *c_accept*), the customer is informed (task *rec_acc*) and the bookstore continues processing the customer order. The bookstore sends a request to the shipper (task *s_request*), the shipper evaluates the request (task *eval_s_req*) and either accepts (task *s_accept*) or rejects (task *s_reject*). If the bookstore receives a negative answer, it searches for another shipper.

After a shipper has been found, the publisher is informed (task *inform_publ*), the publisher prepares the book for shipment (task *prepare_b*), and the book is sent from the publisher to the shipper (task *send_book*). The shipper prepares the shipment to the customer (task *prepare_s*) and actually ships the book to the customer (task *ship*). The customer receives the book (task *rec_book*) and the shipper notifies the bookstore (task *notify*). The bookstore sends the bill to the customer (task *send_bill*). After receiving

both the book and the bill (task *rec_bill*), the customer makes a payment (task *pay*). Then the bookstore processes the payment (task *handle_payment*) and the inter-organizational workflow terminates.

The XRL route shown in Appendix B just illustrates some of the XRL routing elements. The description is far from complete, for example, the detailed descriptions of tasks and conditions have not been included. Please note that, since an XRL route specifies the life cycle of a particular workflow instance (that is, work case), any instance can be modified without reference to some underlying workflow schema type.

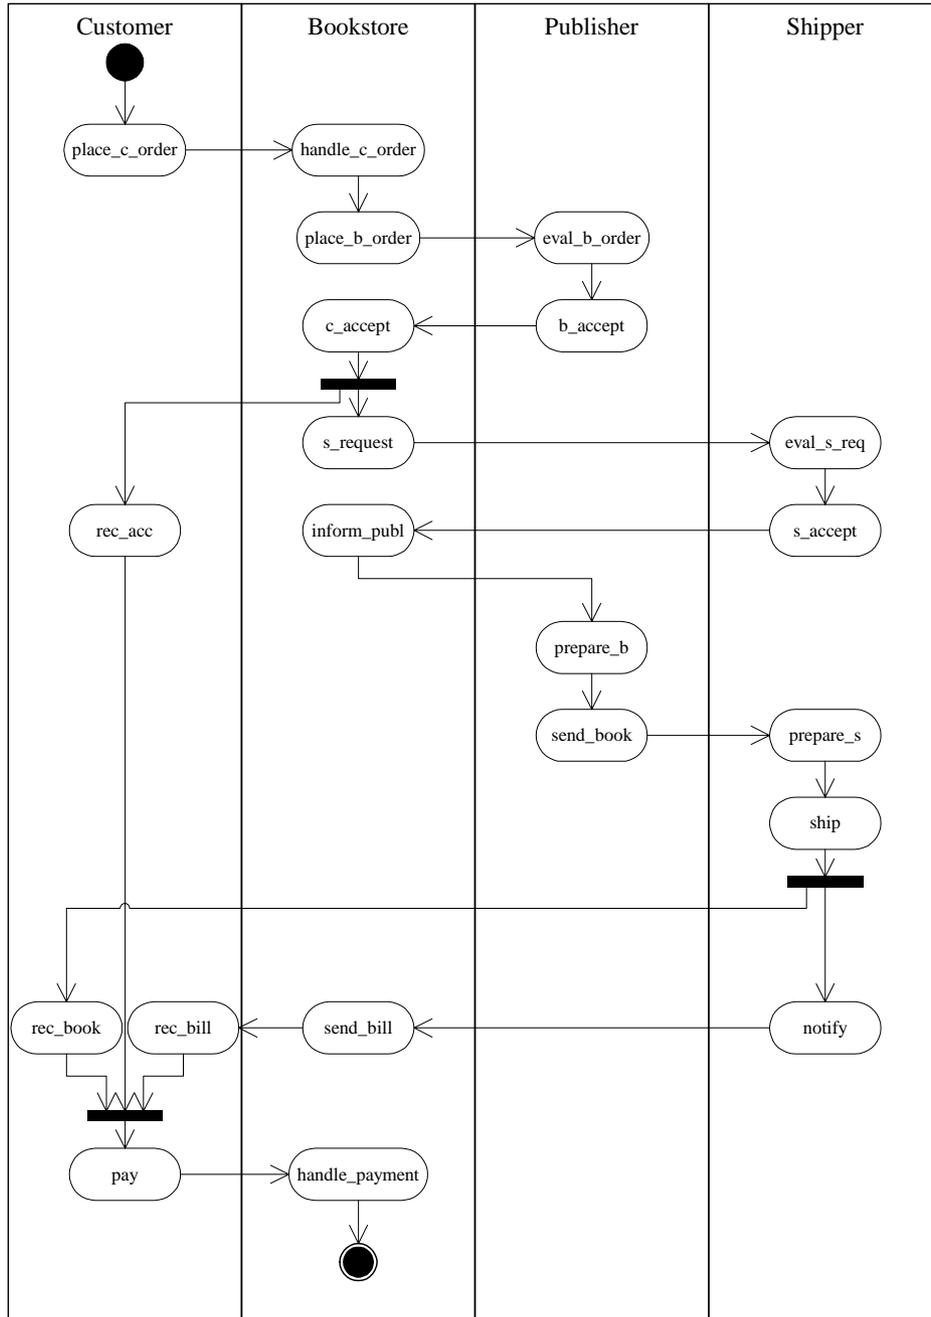


Figure 5. Typical order flow for an electronic bookstore.

4. Workflow nets

As is shown in [8], the semantics of XRL can easily be expressed in terms of Petri nets. XRL/Flower uses a Petri-net engine to interpret and execute XRL routes. Unfortunately, the transformation given in [8] does not result in so-called *Workflow nets* (WF-nets), as observed in [9]. Consequently, the strong theoretical results for WF-nets cannot be used. Moreover, it is difficult to deploy our workflow verification tool Woflan [37]. In this paper, we present a new approach. The semantics presented in [8] is modified considerably. The semantic transformation into Petri nets becomes more complex. However, the resulting Petri nets are WF-nets, which allows us to use the results presented in [1, 38]. Before we present this new transformation, we briefly introduce some of the concepts related to WF-nets. We assume some basic knowledge of Petri nets [32, 33].

A Petri net that models the control-flow dimension of a workflow, is called a *WF-net*. It should be noted that a WF-net specifies the dynamic behavior of a single case in isolation.

Definition 1. WF-net

A Petri net $PN = (P, T, F)$ is a WF-net (Workflow net) if and only if:

- (1) *There is one source place $i \in P$, that is, one place without any predecessors.*
- (2) *There is one sink place $o \in P$, that is, one place without any successors.*
- (3) *Every node $x \in P \cup T$ is on a path from i to o .*

A WF-net has one input place (i) and one output place (o) because any case handled by the procedure represented by the WF-net is created when it enters the workflow management system and is deleted once it is completely handled by the workflow management system, that is, the WF-net specifies the life-cycle of a case. The third requirement in Definition 1 has been added to avoid ‘dangling tasks’, that is, tasks that do not contribute to the processing of cases.

For sake of completeness, we mention that the original definition of WF-nets did not include arc weights (sometimes also called multiple arcs). However, as mentioned in [38], it is straightforward to extend WF-nets by allowing arc weights. For the remainder of this paper, we assume that arc weights are allowed in WF-nets.

The three requirements stated in Definition 1 can be verified statically, that is, they only relate to the structure of the Petri net. However, there is another requirement that should be satisfied:

For any case, the procedure will terminate eventually and upon termination there is a token in place o and all the other places are empty.

Moreover, there should be no dead tasks, that is, it should be possible to execute an arbitrary task by following the appropriate route through the WF-net. These two additional requirements correspond to the so-called *soundness property*.

Definition 2. Soundness

A procedure modeled by a WF-net $PN = (P, T, F)$ is sound if and only if:

- (1) *For every state M reachable from state i , there exists a firing sequence leading from state M to state o .*
- (2) *State o is the only state reachable from state i with a token in place o .*
- (3) *There are no dead transitions when starting in state i .*

Note that the soundness property relates to the dynamics of a WF-net. The first requirement in Definition 2 states that starting from the initial state (state i), it is always possible to reach the state with one token in place o (state o). The second requirement states that the moment a token is put in place o , all the other places should be empty. The last requirement states that there are no dead transitions (tasks) in the initial state i .

In [1], it is shown that there is an interesting relation between soundness and well-known Petri-net properties such as liveness and boundedness. A WF-net is sound if and only if the short-circuited net (that is, the net obtained by linking the sink place to the source place) is live and bounded. This result illustrates that standard Petri-net-based analysis techniques can be used to verify soundness.

5. Semantics of XRL in terms of WF-nets

The DTD shown in Figure 4 only describes the syntax of XRL and does not specify the semantics. To provide operational semantics of the routing elements we transform each routing element mentioned in the DTD into a Petri net. Such a transformation was given in [8]. However, as indicated earlier, this transformation does not necessarily yield WF-nets. Therefore, we have modified the transformation given in [8] such that XRL routes are transformed into WF-nets. First, we discuss the problems we encountered when trying to transform XRL routes into WF-nets. Second, we present for the route and for each routing element its Petri-net semantics. Third, we transform the example from Section 3 (see Appendix B for the XRL specification) into a WF-net.

In Section 2 we already observed that the structure of an XRL document forms a tree, with the route element as root. Many routing elements interface only with their parent element and their child elements. For this reason, we propose to ‘copy’ this tree structure to the resulting WF-net: Every routing element is replaced by some Petri-net fragment that interfaces with the Petri nets associated with its parent child elements nets. The exceptions to this rule are the terminate routing elements and the task, wait_all and wait_any routing elements (when events are involved). We deal with these exceptions later on; first we focus on the interface between a parent and a child element.

At first glance, only two places seem to be necessary for the communication between a parent and child: one from parent to child indicating that the child can be started, and one from child to parent indicating that the child has been performed. However, according to the descriptions given in Section 2, there are three routing elements that do not wait until all child elements have been performed: parallel_no_sync, parallel_part_sync, and parallel_part_sync_cancel. As a result, an instance might reach the point of completion while still, somewhere deep inside some subtree, elements still can be performed. Take for instance the simple example where the route contains a parallel_no_sync containing only one task. Because the parallel_no_sync does not wait until the task has been performed, the entire instance might reach the point of completion before the task has been accepted and started. Recall that soundness requires that the remainder of the entire Petri-net is empty when a token is put into the sink place, that is, a Petri-net fragment associated with any routing element has to be empty at that point. Therefore, before actually reaching completion, we have to wait until all these fragments are empty. For this reason, we introduce a third communication place: from child to parent indicating that the Petri-net fragment associated with the entire subtree of the child is now empty of tokens.

Figure 6 shows the basic routing element. A token in place *prev* indicates that the routing element can be started. A token in place *exec* indicates that the routing has started. However, in many routing elements this *exec* place is redundant, and therefore omitted. When the routing element has been performed, it puts a token in places *next* and *sig*. The token in place *next* informs the parent that this element has been performed, while the token in place *sig* indicates that the routing element is waiting until all descending

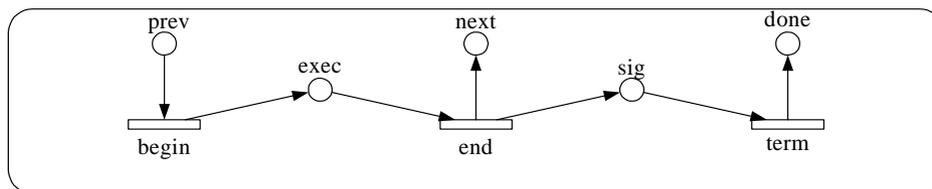


Figure 6. Basic routing element.

routing elements have been performed too, that is, until all subtrees are empty. If all are empty, it puts a token in place *done*, indicating that it is now empty of tokens except for the token in the *done* place.

When a terminate occurs, the entire instance, that is, the route itself, is to be completed. This clearly has an effect on the instance level. In Petri nets, it is hard to foresee all possible reachable states, and to add transitions such that from every reachable state we can reach state *o*. A simple observation alleviates this problem: If we bypass every task, wait_all, and wait_any, the instance automatically reaches completion! The tasks need to be bypassed because we cannot allow that a task is started after a terminate occurred. Both wait routing elements need to be bypassed because they are not allowed to wait any longer after a terminate occurred. Note that this solution assumes that running tasks are not preempted when a terminate occurs. As a result, we can treat a terminate almost in a similar way as we treat an event. Because all tasks need to have both places *terminate* and *nonterminate* present, and all waits need to have the place *terminate* present, we incorporate the global part of the terminate in the route.

Figure 7 shows the semantics of the route element, containing the terminate at the instance level. For sake of clarity, a stub (drawn dotted) replaces the top routing element (the only child routing element of the route element). Initially, place *input* (which corresponds to place *i* in Definition 1) contains one token, indicating that the instance has not started yet. Transition *begin* starts the instance and

- (1) starts the top child routing element,
- (2) enables the terminate,
- (3) enables all events, and
- (4) enables every *parallel_part_sync* or *parallel_part_sync_cancel*.

Note that items (3) and (4) are explained later on and not shown in this net. Item 3 is shown in Figure 8 and item 4 is shown in Figures 16 and 17.

After the top routing element and all its descendants have completed, transition *end* initiates the completion phase of the instance. First, all events and the terminate are reset (which happens while place *almost* contains a token). Second and last, transition *done* completes the instance and

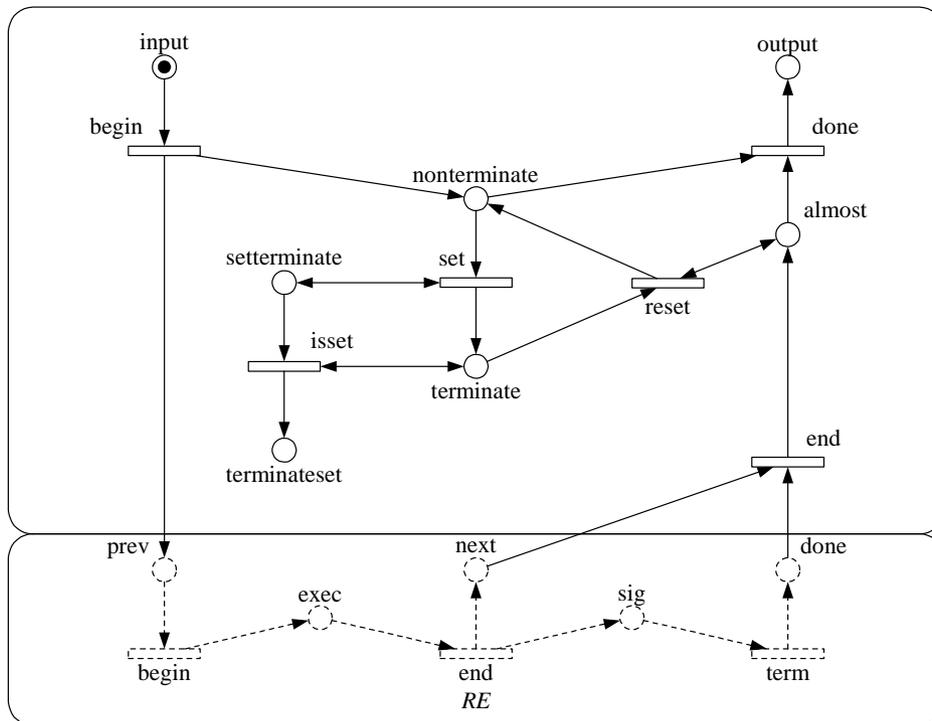


Figure 7. Semantics of route.

- (1) removes the token from place *almost*,
- (2) disables the terminate,
- (3) disables all events, and
- (4) disables every *parallel_part_sync* and *parallel_part_sync_cancel*.

Again, items (3) and (4) are explained later on and not shown in this net.

Of course, if no terminates are present in the XRL route, the part concerning the terminate is discarded from the semantics of the route.

Like terminate, events are defined on the top level of the instance, that is, the route level, not on some local level deeply nested in some subtree: If some task in some subtree sets a certain event, than some wait in some other subtree might be affected. For this reason, *events are handled on the instance level*: For every event, we add a Petri-net fragment that manages the event. Such a fragment interfaces only with the Petri-nets associated with the route itself (for enabling and disabling the event), tasks (for setting the event), and both waits (for testing the event). On the instance level, two places are introduced for every event: *event* and *nonevent*. Only one of them may contain a token: Either the event has occurred (*event* contains a token) or not (*nonevent* contains a token). When the instance is started, each event is enabled by putting a token into place *nonevent*, and when the instance completes, it is disabled by removing the token from that place. When the event has been set, transition *reset* can move the token from *event* to *nonevent* when the instance is completing. A task sets the event by putting a token in place *setevent* and waiting until a token is put into place *eventset*. The transitions *set* and *isset* take care that the event is set when this place contains a token. They also take care that this token is moved to place *eventset* when the event has been set. Figure 8 shows an event on the instance level.

At this point, we have dealt with the exceptions and can introduce the Petri-net semantics of every routing element.

Figure 9 shows the semantics of the task routing element. When a task is started, what happens depends on whether or not a terminate has occurred. If a terminate has occurred, the transitions *bbegin* and *bend* bypass the normal execution of the task. Otherwise, transition *begin* starts the execution by offering the task to some resources. After the task has been performed by some resource, transition *events* has the appropriate events set (could be none), and the task waits until these events are set. After all have been set, transition *end* signals that the task has been performed. Because there are no child routing elements, this automatically results in an empty subtree, which is signaled too.

Figure 10 shows the semantics of the sequence routing element. Transition *begin* starts the sequence, and puts a token in the place *prev* of the first child routing element. Doing so, this routing element can be started. When this routing element has been performed, it puts a token in its place *next*. Transition *next_j* moves this token to the place *prev* of the next child routing element, and so on. When the last child routing

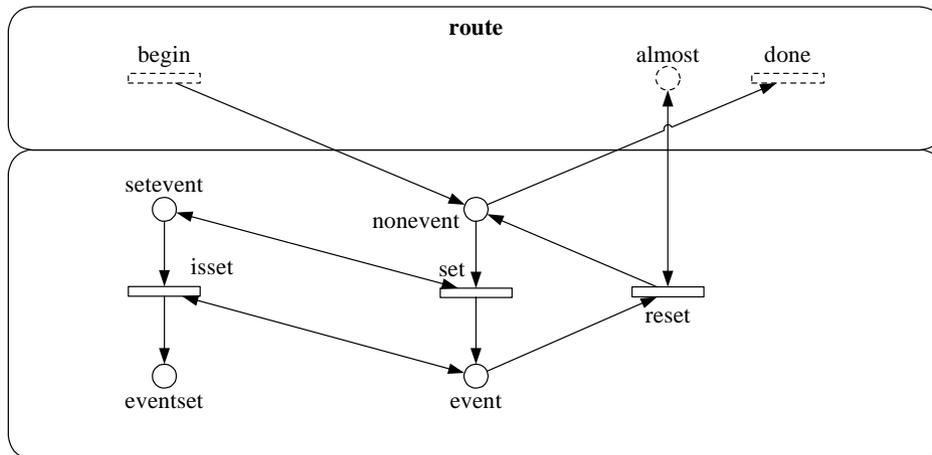


Figure 8. Event on global level.

element has been performed, transition *end* puts a token in place *next*, signaling that the entire sequence has been performed. After firing *end*, the sequence waits until all descending routing elements have been performed. If all have been performed and hence their subtrees are empty, transition *term* puts a token in place *done*.

Figure 11 shows the semantics of the *any_sequence* routing element. Transition *begin* starts the *any_sequence*, and puts tokens in the place *prev* for every child routing element, and a token in the place *exec*. The token in place *exec* guarantees the mutual exclusion of the child routing elements. A child routing element can only start if place *exec* contains a token, and when it starts, it removes that token. When the child routing element has been performed and transition *end* fires, the token is put back. After all child routing elements have been performed, the transition *end* puts a token in place *next*, signaling that the entire *any_sequence* has been performed. After firing *end*, the *any_sequence* waits until all descending

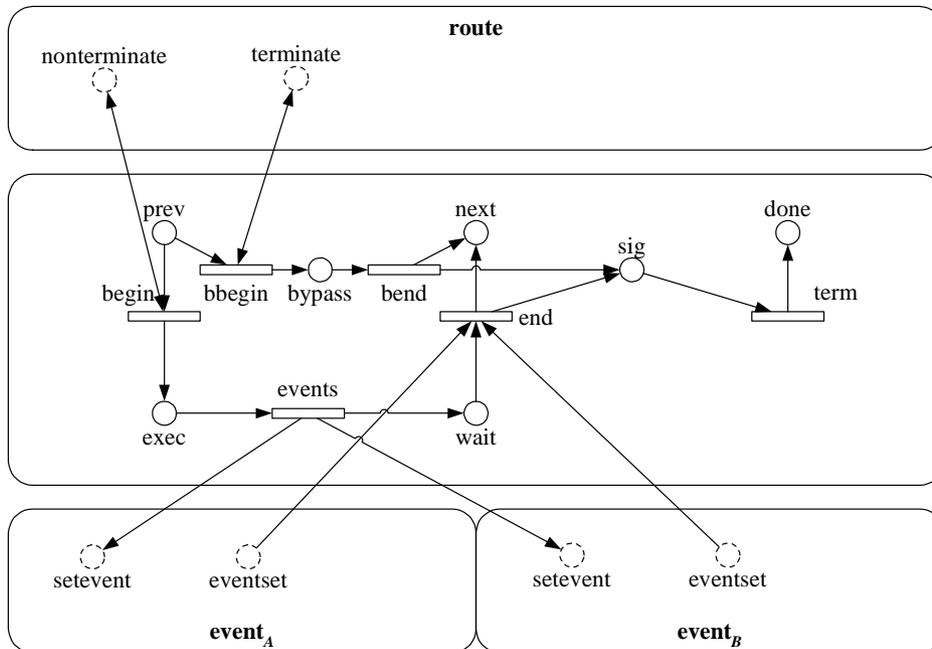


Figure 9. Semantics of task.

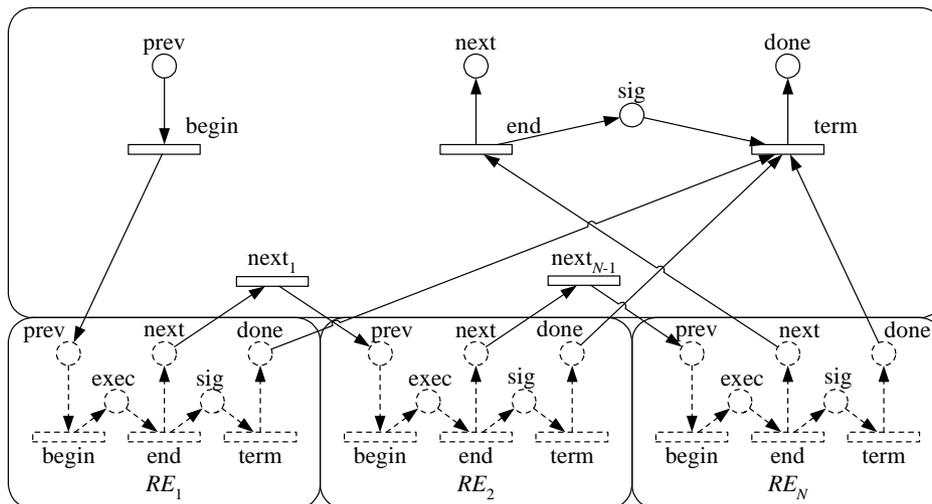


Figure 10. Semantics of sequence.

routing elements have performed. If all have been performed and hence their subtrees are empty, it puts a token in place *done*.

Figure 12 shows the semantics of the choice routing element. Transition *begin* starts the *choice*, and puts a token in every child's *prev* place. The first child routing element that fires its transition *begin*, disables the other child routing elements. After this first child routing element has been performed, its transition *end* fires, which puts a token in place *join*. As a result, transition *end* fires, signaling that the choice has been performed. After firing *end*, the choice waits until the first child routing element and all its descending routing elements have been performed. If all have been performed and hence their subtrees are empty, the child's transition *term* fires, followed by the choice's transition *term*, which puts a token in its place *done*.

Figure 13 shows the semantics of the condition routing element. Transition *begin* starts the condition, and puts a token in place *split*. At this point, the condition is evaluated. If the condition is evaluated to true, transition *tbegin* fires, otherwise transition *fbegin* fires. Assume without loss of generality that transition *tbegin* fires. This transition puts a token in place *texec* and enables the child routing element of every *true* child element. After all these 'grandchild' routing elements have been performed, transition *tend* fires, which puts tokens in places *join* and *tsig*. At this point, transition *end* fires, signaling that the condition has

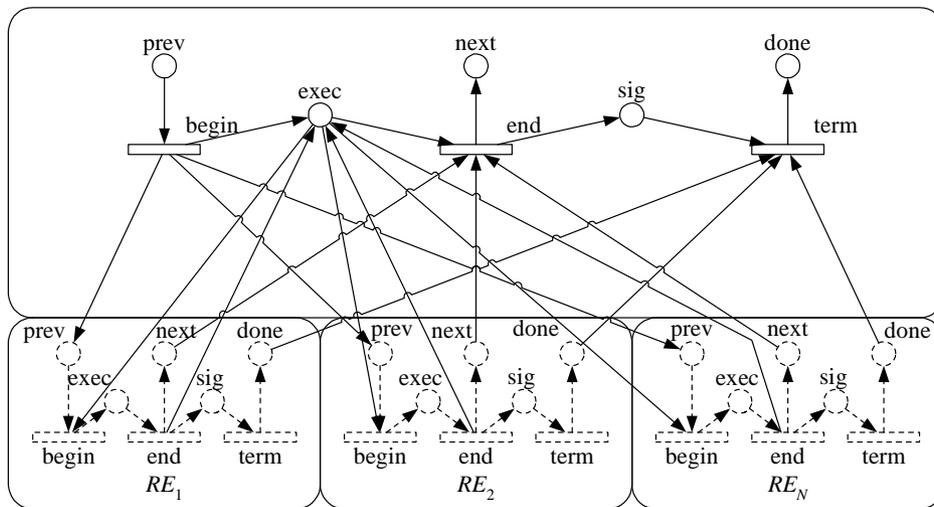


Figure 11. Semantics of any_sequence.

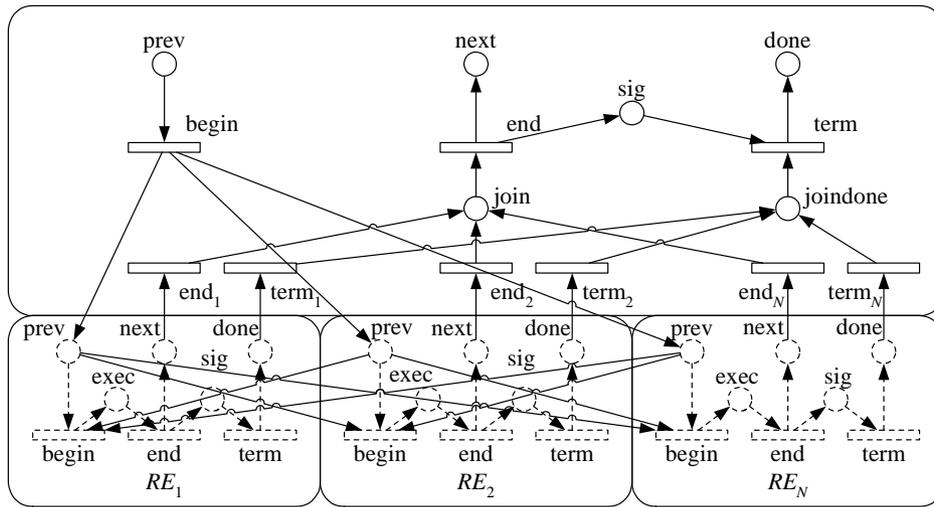


Figure 12. Semantics of choice.

been performed. After all descending routing elements of the *true* child elements have also have been performed, transition *tterm* fires, followed by transition *term*, which puts a token in place *done*.

Figure 14 shows the semantics of the *parallel_sync* routing element. Transition *begin* starts the *parallel_sync*, and puts a token in the place *prev* for every child routing element. After all child routing elements have been performed, transition *end* fires, which puts a token in place *next*. After firing *end*, the *parallel_sync* waits until all descending routing elements have been performed. If all have been performed, transition *term* fires, which puts a token in place *done*.

Figure 15 shows the semantics of the *parallel_no_sync* routing element. Transition *begin* starts the *parallel_no_sync*, and puts a token in place *prev* for every child routing element and a token in place *exec*. As a result of the token in place *exec*, transition *end* can fire immediately, signaling that the *parallel_no_sync* has been performed. After firing *end*, the *parallel_no_sync* waits until all descending child routing elements have been performed. If all have been performed, transition *term* fires, which puts a token in place *done*.

Figure 16 shows the semantics of the *parallel_part_sync* routing element. Transition *begin* starts the *parallel_part_sync*, and puts a token in place *prev* for every child routing element and a token in place *exec*.

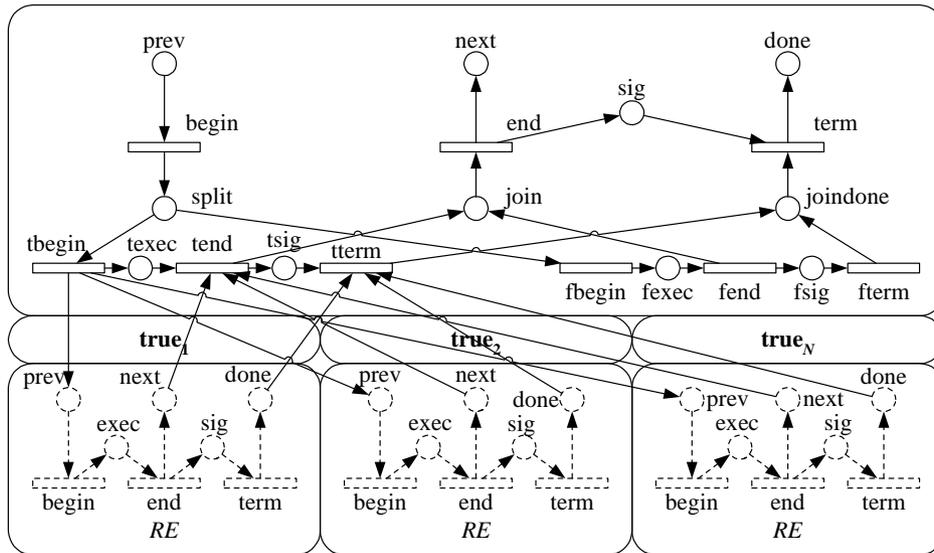


Figure 13. Semantics of condition.

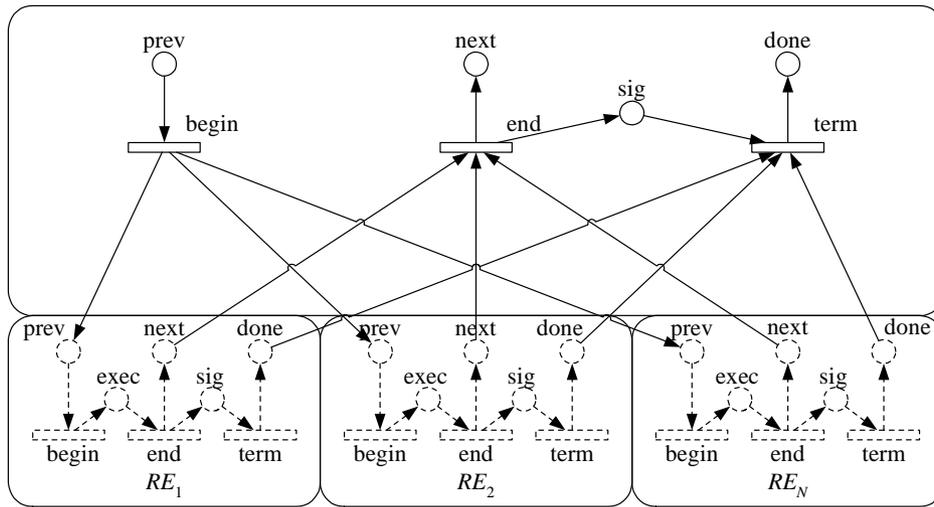


Figure 14. Semantics of *parallel_sync*.

After any child routing element has been performed, the appropriate transition $count_i$ fires, which moves the token to the place $count$. After K tokens have been moved this way, that is, after at least K child routing elements have been performed, transition end fires, signaling that the `parallel_part_sync` has been performed, and puts a token in place $next$. Note that the place $exec$ is not redundant for this routing element: if we had removed the place $exec$, transition end could fire $\lfloor N/K \rfloor$ times. For instance, if there were 10 child routing elements ($N = 10$) and the given number were 3 ($K = 3$), transition end could fire 3 times, signaling 3 times that the `parallel_part_sync` has been performed, which is clearly an error. After firing end , the `parallel_part_sync` waits until the remaining child routing elements and all descending routing elements have been performed. If all have been performed, transition $term$ fires, which puts a token in place $done$.

To avoid problems with recurrent `parallel_part_sync`s, that is, if a `parallel_part_sync` is embedded in a `while_do`, we prevent the `parallel_part_sync` to be recurrent by introducing the place $free$. As indicated in

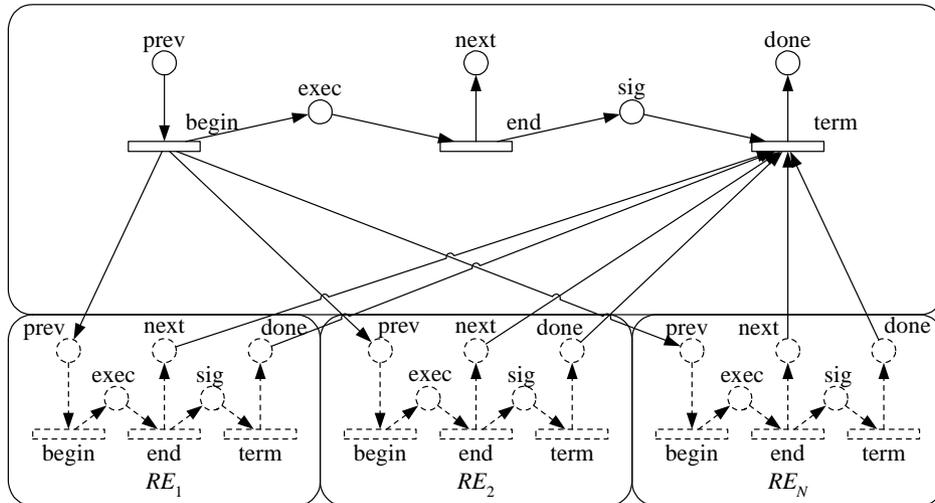


Figure 15. Semantics of `parallel_no_sync`.

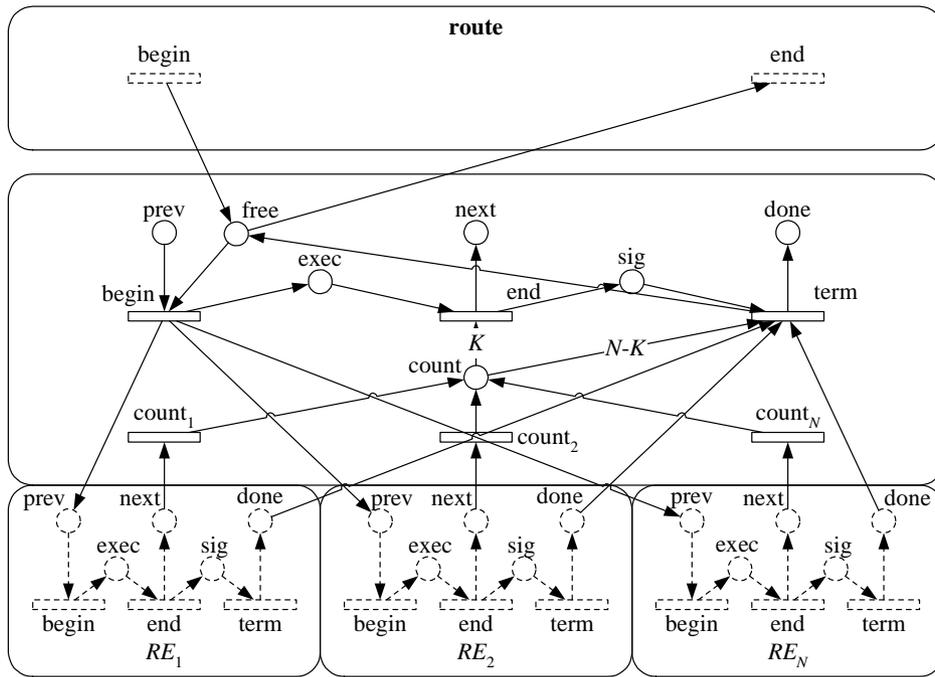


Figure 16. Semantics of `parallel_part_sync`.

Figure 16, transition *begin* of the route puts a token in place *free*, and transition *end* of the route removes the token. Note that transition *begin* has to fire before any routing element is started, and that *end* can only fire after all routing elements have been performed.

Figure 17 shows the semantics of the *parallel_part_sync_cancel* routing element. This semantics is an extension of the semantics of the *parallel_part_sync*: Transitions *cancel_i* have been added. After it has been performed, any child routing element that has not started yet can be cancelled by firing the appropriate transition *cancel_i*. When transition *cancel_i* fires, it removes the token from the place *prev* of the *i*-th child routing element, and puts tokens in its places *next* and *done*.

Figure 18 shows the semantics of the *wait_all* routing element. Transition *begin* starts the *wait_all*, and puts a token in place *wprev*. After firing *begin*, the *wait_all* waits until

- (1) all events referred to have occurred,
- (2) a terminate has occurred, or
- (3) any of deadlines associated with the timeouts are exceeded.

If all events referred to have occurred, transition *wait* fires, which puts tokens in places *wnext* and *wdone*. If a terminate has occurred, transition *twait* fires, which also puts tokens in places *wnext* and *wdone*. If a timeout occurs, the child routing element of this timeout is started. After this child routing element has been performed, a token is put in place *wnext*. After all descending routing elements of the timeout have been performed, a token is put in place *wdone*. When place *wnext* contains a token, transition *end* fires, signaling that the *wait_all* has been performed, and puts a token in place *next*. After firing *end*, the *wait_all* waits until place *wdone* contains a token. If this place contains a token, transition *term* fires and puts a token in place *done*.

Figure 19 shows the semantics of the *wait_any* routing element. Transition *begin* starts the *wait_any*, and puts a token in place *wprev*. After firing *begin*, the *wait_any* waits until either

- (1) any event referred to has occurred,
- (2) a terminate has occurred, or
- (3) any of the timeouts occurs.

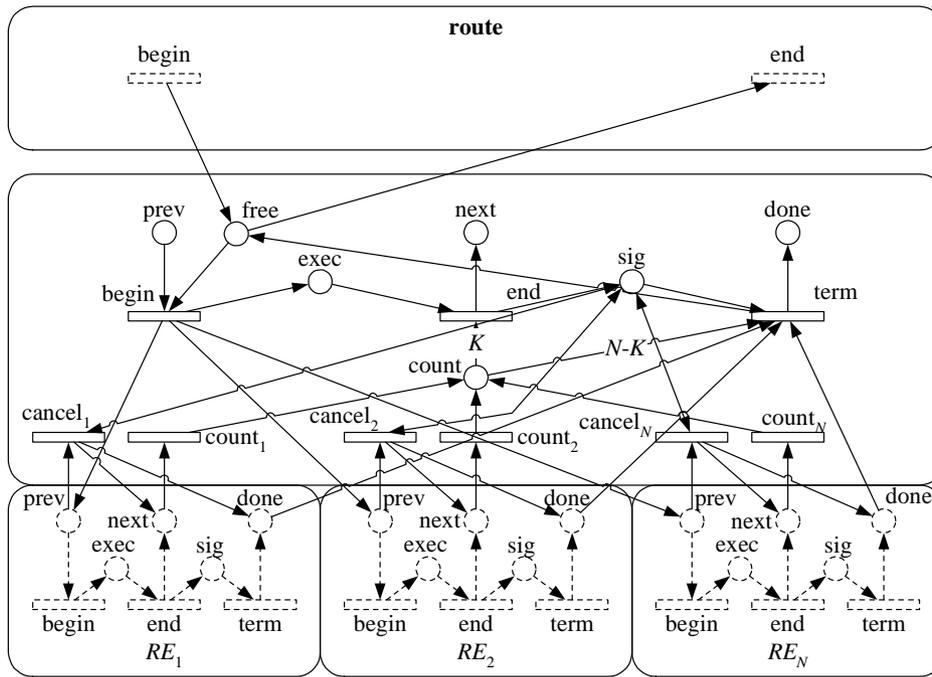


Figure 17. Semantics of *parallel_part_sync_cancel*.

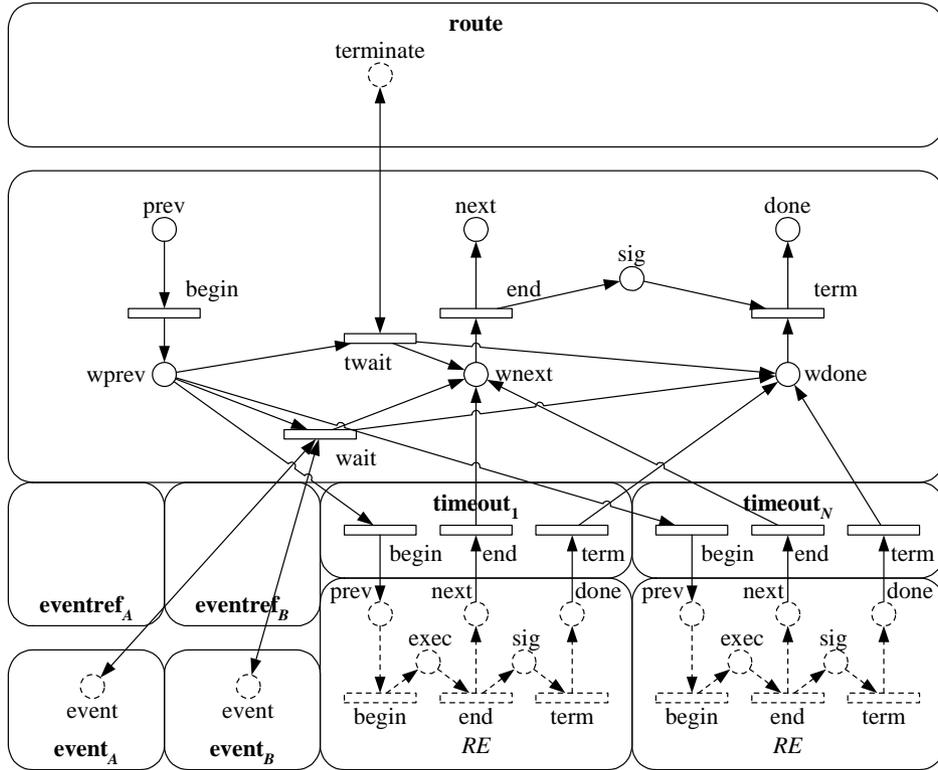


Figure 18. Semantics of wait_all.

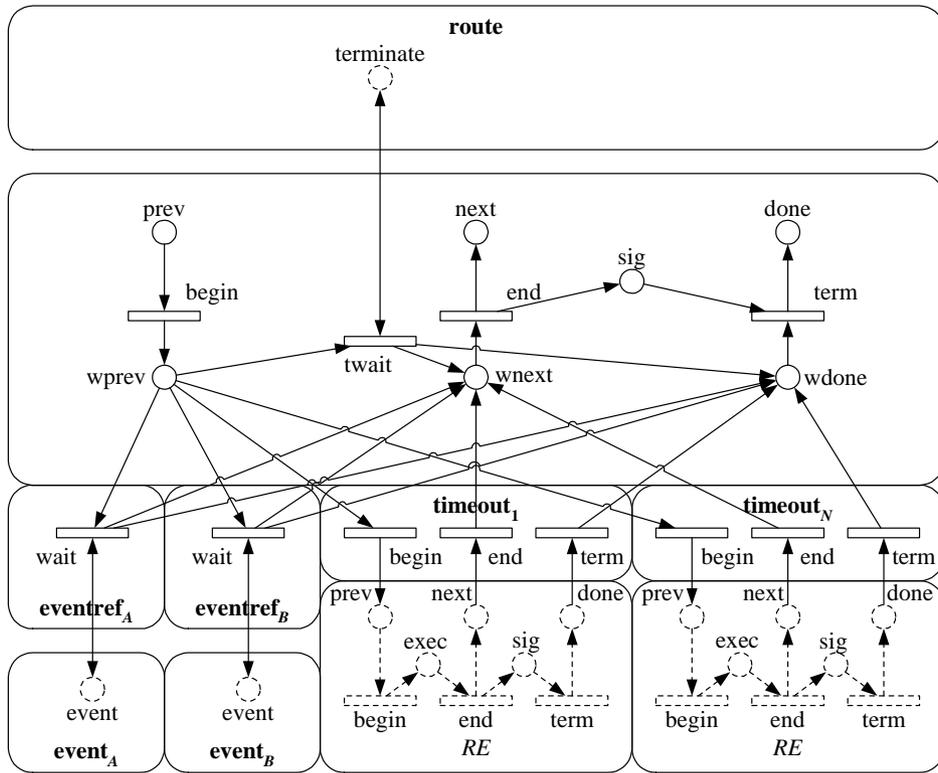


Figure 19. Semantics of wait_any.

If any event referred to has occurred, its transition *wait* fires, which puts tokens in places *wnext* and *wdone*. If a terminate has occurred, transition *twait* fires, which also puts tokens in places *wnext* and *wdone*. If a timeout occurs, the child routing element of this timeout is started. After this child routing element has been performed, a token is put in place *wnext*. After all descending routing elements of the timeout have been performed, a token is put in place *wdone*. When place *wnext* contains a token, transition *end* fires, signaling that the *wait_any* has been performed, and puts a token in place *next*. After firing *end*, the *wait_any* waits until place *wdone* contains a token. If this place contains a token, transition *term* fires and puts a token in place *done*.

Figure 20 shows the semantics of the *while_do* routing element. Transition *begin* starts the *while_do*, and puts a token in place *next* of its child routing element, and X tokens in the place *done*, where X is a positive number that serves as a parameter to limit the number of concurrent instances. After firing *begin*, the condition is evaluated (by the enactment server). If the condition evaluates to true, the transition *true* fires, which starts another iteration of the child routing element. If the condition evaluates to false, transition *end* fires, signaling that the *while_do* has been performed and puts a token in place *next*. After firing *end*, the *while_do* waits until all descending routing elements corresponding to all iterations have been performed. If all have, transition *term* fires, which puts a token in place *done*. Note that a new iteration can only start, when at most $X - 1$ previously started iterations are still running. Also note that X can be set to a very large value when enacting the entire XRL route, but for verification purposes X should be as small as possible: The larger X is, the larger the state space of the entire resulting Petri net becomes.

Figure 21 shows the semantics of the *terminate* routing element. Recall that the instance level part is incorporated in the semantics of the route element, cf. Figure 7. Transition *begin* starts the *terminate*, and has the *terminate* set. Transition *end* fires if the *terminate* has been set, signaling that the *terminate* has been performed, and puts a token in place *next*. After firing *end*, transition *term* puts a token in place *done*.

At this point, all XRL routing elements can be transformed into Petri-nets. By starting with the XRL route and recursively transforming each child routing element into its corresponding Petri-net semantics, one obtains a WF-net. As a result, we can check the important *soundness property* introduced in the previous section. This will be discussed in the next section.

The example shown in Appendix B results in a WF-net containing 303 places and 275 transitions when applying the rules presented in this section. These numbers indicate that the dynamic behavior of the XRL route presented in Section 3 is complex. However, this complexity is hidden from both the designer and the user.

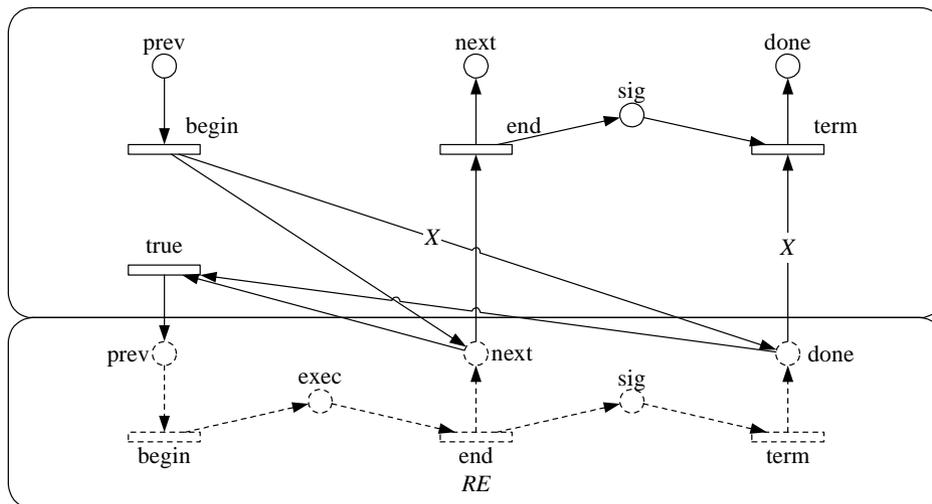


Figure 20. Semantics of *while_do*.

6. Soundness of XRL

This section discusses the soundness of an XRL route, that is, the soundness of the WF-net it is transformed into. Recall that for soundness three requirements should hold. The first requirement states that the final state o should be reachable, that is, proper completion is possible. The second requirement states that completion is always proper, that is, no tokens are left after completion. The third requirement states that for every transition there is a way to fire it from the initial state i . First, we show that completion of an XRL route is always proper. Second, we show that only certain XRL routing elements can invalidate the other two soundness requirements.

Lemma 1. *Coverage of a single routing element.*

Let RE be a routing element and let PN be its Petri-net semantics as described in Section 5. Then all places in PN can be covered by two place invariants P_{RE} and Q_{RE} , where P_{RE} contains the places $prev$ and $next$ with identical weights and Q_{RE} contains the places $prev$ and $done$ with identical weights.

Proof. By induction. Most of this proof is straightforward. For this reason, we will restrict ourselves to two routing elements that serve as examples: *task* and *parallel_part_sync_cancel*. For sake of simplicity, we assume that the invariants of the child routing elements are recalibrated such that the places *next* and *done* have weight 1. As a result, some weights might be fractions instead of natural numbers, but for the proof this is irrelevant. For sake of completeness, the invariants for all routing elements are listed in Appendix F. For the task (*tk* for short), which serves as example for the induction base, the invariants are as follows:

$$P_{tk} = prev + exec + wait + bypass + next,$$

$$Q_{tk} = prev + exec + wait + bypass + sig + done.$$

For the *parallel_part_sync* (*ppsc* for short), which serves as example for the induction step, the invariants are as follows:

$$P_{ppsc} = prev + exec + next,$$

$$Q_{ppsc} = 2N \times (prev + done) + free + \left(\sum_{i=1}^N P_{RE_i} + Q_{RE_i} \right) + exec + count + (K + 1) \times sig.$$

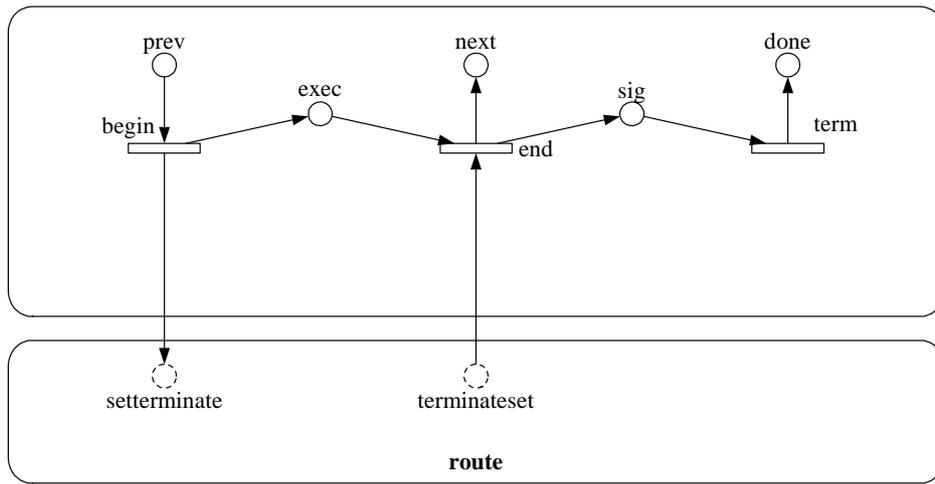


Figure 21. Semantics of terminate.

Lemma 2. *Coverage of top routing element.*

Let R be an XRL route, let ME be the top routing element in R , and let PN_R and PN_{RE} be the Petri-net semantics of R and RE as described in Section 5. Then there exists a place invariant P_1 containing (i) the places *input* and *output* with identical weight and (ii) all places in PN_{RE} .

Proof. By construction. $P_1 = \text{input} + \text{almost} + \text{output} + P_{RE} + Q_{RE}$.

Lemma 3. *Coverage of a single event.*

Let R be an XRL route, let E be an event and let PN_R and PN_E be their Petri-net semantics as described in Section 5. Then there exists a place invariant P_E containing (i) the place *input* and *output* with identical weight and (ii) all places of PN_E .

Proof. By construction. The places *event* and *nonevent* are easily covered. However, to cover *setevent* and *eventset* we need to a variant of P_1 . Let P_1^E be equal to P_1 , except for the tasks where event E is set. For these tasks, the occurrence of place *wait* is replaced by the expression $E_{\text{setevent}} + E_{\text{setevent}}$. It is straightforward to check that P_1^E is a place invariant containing places *input* and *output* with identical weights. Then $P_E = P_1^E + \text{input} + \text{almost} + \text{output} + E_{\text{event}} + E_{\text{nonevent}}$.

Lemma 4. *Coverage of all events.*

Let R be an XRL route and let S be the set of events in R . Then there exists a place invariant P_2 containing (i) the places *input* and *output* with identical weight and (ii) all places corresponding to the events in S .

Proof. By construction. $P_2 = \sum_{E \in S} P_E$.

Lemma 5. *Coverage of complete XRL route.*

Let R be an XRL route and let PN_R be its Petri-net semantics as described in Section 5. Then there exists a place invariant P containing (i) the places *input* and *output* with identical weight and (ii) all places in PN_R .

Proof. By construction. To cover the places *terminate* and *nonterminate* we introduce another variant of P_1 : P_1^T . In P_1^T all occurrences of any terminate's *exec* place is replaced by the expression *terminate* + *nonterminate*. It is straightforward to check that P_1^T is a place invariant containing places *input* and *output* with identical weights. Then $P = P_1 + P_1^T + P_2$.

Theorem 6. *Completion of an XRL route is always proper.*

Proof. By Lemma 5, place invariant P contains (i) places *input* and *output* with identical weights and (ii) all places of the Petri-net semantics of the XRL route. Because initially only *input* is marked with one token, no other place can be marked when *output* is marked.

Theorem 7. *If every wait_all and every wait_any in an XRL route contains a timeout, then completion of the XRL route is guaranteed.*

Proof. By induction. It is straightforward to check that only a wait_all or wait_any that does not contain a timeout can get stuck. Note that we assume that tasks are completed eventually.

Theorem 8. *If every wait_all and every wait_any in an XRL route contains a timeout, and if the XRL route contains no terminate, then no transition in the Petri-net semantics of the route is dead.*

Proof. It is straightforward to check that only a wait_all or wait_any that does not contain a timeout can get stuck. It is also straightforward to check that some transitions can only get by-passed by a terminate construct.

Corollary 1. *There are two possible causes for violation of soundness: (i) a wait_all or wait_any gets stuck, or (ii) a terminate occurs.*

In case a terminate occurs, the short-circuiting transition in the short-circuited net will be live: A terminate guarantees completion. In case a wait_all or wait_any gets stuck, the short-circuiting transition (that is, the transition linking the place *output* to *input* in the short-circuited net) will not be live.

7. Verification of XRL

With the semantics specified in terms of WF-nets, described in Section 5, the theory and tools for WF-nets can be deployed in a straightforward manner. This allows us to use Woflan for verifying the correctness of an XRL route using criteria such as the soundness property. Unfortunately, XRL routes with a lot of parallelism tend to have a large state space, thus complicating verification from a computational point of view. Therefore, we propose a verification procedure that consists of two optimization steps. In the first step, the XSL transformation, which transforms the XRL route into a WF-net, reduces the WF-net by using structural properties of XRL. In the second step, Woflan reduces the WF-net by applying the well-known liveness and boundedness preserving reduction rules for Petri nets [32].

The first step is the reduction by the XSL transformation based on structural properties of XRL. Figures 9 to 21 show a place named *done* to accommodate the situation where completion of a routing element does not automatically yield completion of its descendants. This situation can only occur if the routing element ascends some parallel_no_sync, parallel_part_sync, or parallel_part_sync_cancel routing element. In all other cases, there is no need to model things related to these done places. For instance, assuming the routing element in Figure 20 has no done place allows us to remove the upper half of the Petri net (that is, RE_{done} , *sig*, *term*, and *done*). Similar simplifications are possible if no events are used. Moreover, we can apply the result presented in [9]: a routing element without any event, wait_all, wait_any, or terminate is sound and can therefore be replaced by the basic routing element shown in Figure 6. Note that this is consistent with the results in Section 6. When these reduction rules are applied, the XRL route shown in Appendix B is transformed into a WF-net that contains only 108 places and 105 transitions. Compared to the original WF-net, the reduced WF-net is considerably smaller and less complex. Note that several routing elements can be abstracted from, and that the resulting WF-net need not contain any done places.

The second step is the reduction of the resulting WF-net by Woflan based on liveness and boundedness preserving reduction rules. Fragments of various routing elements are connected by transitions. This introduces a lot of transitions that are not relevant for the verification but introduce transient states. These and other parts of the WF-net can be reduced enormously without losing information relevant for the verification. In Section 4, it was pointed out that soundness corresponds to liveness and boundedness [1]. This allows us to apply the well-known liveness and boundedness preserving reduction rules for Petri nets [32]. These rules are shown in Figure 22. Note that not all rules are relevant for reducing a WF-net derived from an XRL route: For instance the fifth rule will not be applied, because the only marked place in the WF net has no input arcs. After these reduction rules are applied, the reduced WF-net mentioned under Step 1 contains only 21 places and 18 transitions and is shown in

Figure 23. Table 1 shows how the reductions affect the size (in number of places and transitions) of the example WF-net.

After making the appropriate changes, the soundness results from Section 6 still hold upon applying these reduction rules. It is straightforward to check that the *wait* transition in a *wait_all* or *wait_any* will not be reduced. Nor will be the *terminate*, if present. Note that we do not apply the six WF-net-based reduction rules on the short-circuited net, but on the original WF-net. As a result, the short-circuited

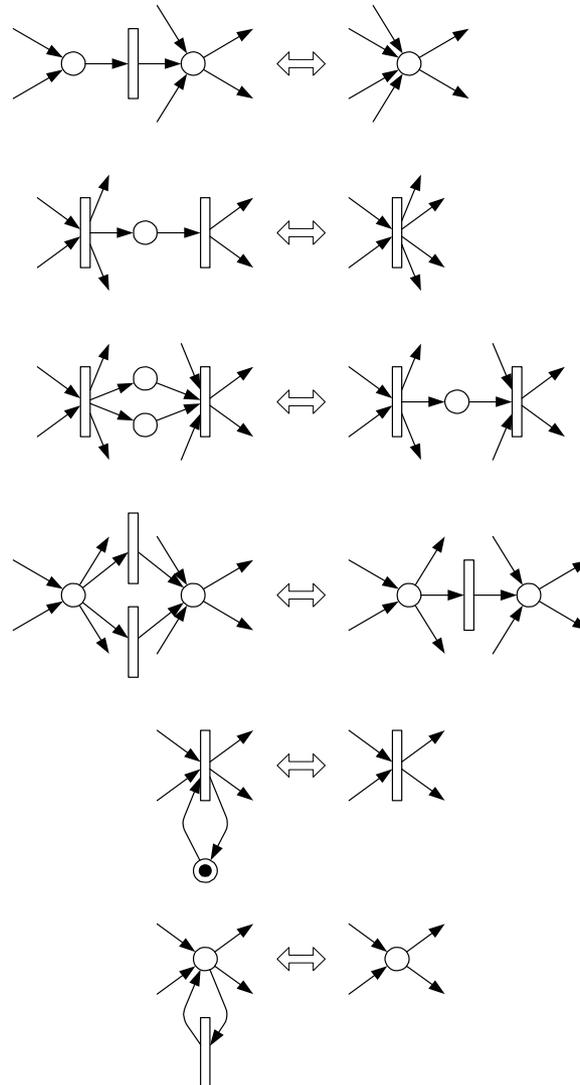


Figure 22. Liveness and boundedness preserving transformation rules.

WF-net	Number of places	Number of Transitions
Original	303	275
After XRL-based reduction	108	105
After both XRL-based and Petri-net-based reduction	21	18

Table 1. The effects of both reductions on the size of the example WF-net.

transition will still be present after the WF-net has been reduced. As mentioned in Section 6, this transition can be very useful when diagnosing the net.

Using standard Petri-net-based analysis tools, or dedicated tools such as Woflan, it is easy to show that Figure 23 is sound. Therefore, the XRL route shown in Section 3 is correct, that is, free of deadlocks, livelocks and other anomalies. Note that Figure 23 is obtained after applying both types of reduction.

XRL/Woflan is based on our workflow verification tool Woflan [37, 38]. Woflan (See <http://www.tm.tue.nl/it/woflan>) is designed as a workflow-management-system-independent analysis tool. In principle, it can interface with many workflow management systems. At present, Woflan can interface with the workflow products COSA (Thiel Logistic AG/Software Ley), METEOR (LSDIS), and Staffware (Staffware), and the BPR-tool Protos (Pallas Athena). Furthermore, Woflan can read Petri Net Markup Language (PNML) files. PNML is a Petri net file format based on XML [22]. Therefore, it is natural to use XSLT to automatically transform an XRL route into a PNML representation that can be diagnosed using Woflan. We have implemented this transformation in XSLT and the two types of reduction rules presented in this section. These two types of reduction rules allow us to verify large XRL routes containing hundreds of tasks.

We propose a verification procedure that consists of three steps. In the first step, the XRL route is transformed into a WF-net, taking into account the three reduction rules based on the structure of XRL. In the second step, the resulting WF-net is reduced even further using the six reduction rules based on the structure of the WF-net. In the third step, we use Woflan to verify and diagnose the reduced WF-net.

A final note on the terminate. If terminates are present, there might be tasks and waits that have to complete before any terminate can occur. In such a task or wait, the bypassing transitions are dead. However, these dead transitions do not correspond to an error in the XRL route, they only befoul the diagnostic information. To get rid of these dead transitions, we can add two additional routing elements to the route: A `parallel_sync` and a `terminate`. The `parallel_sync` will be the new top routing element and will have the old top routing element and the additional `terminate` as child routing elements. As a result, a `terminate` can occur immediately after the instance has started, so the set of tasks and waits that have to complete before any `terminate` can occur will be empty.

8. Extensibility of XRL

Extensibility is an important feature of our approach. Therefore, the architecture of Figure 1 has been designed for extensibility. This means that an end-user can add his or her own new routing elements to the

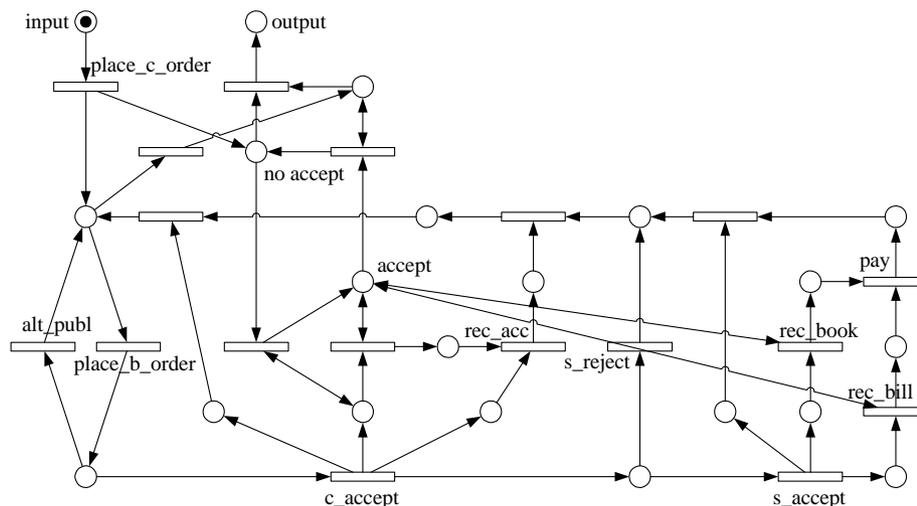


Figure 23. The WF-net corresponding to the example XRL route in Section 3 after both reductions.

DTD and implement them using this toolset. This gives the end-user a powerful capability. We first describe the general approach and then illustrate it with three examples.

The basic approach is as follows:

- (1) Add a new routing element to the DTD of XRL.
- (2) Specify the semantics of the new routing element in XSLT.
- (3) Verify the XSLT specification of the new routing element.
- (4) Add the XSLT specification to the library used in the transformation step.

Thus, for any new routing element it suffices to add it to the DTD and specify its corresponding Petri-net semantics in XSLT. However, XSLT is a very verbose format, and editing XSLT specifications directly is cumbersome. For this reason, we use a set of macros while editing XSLT specifications. These macros, which are transformed into XSLT using a C preprocessor, are listed in Appendix D.

Now, consider the architecture shown in Figure 1. Note that, in the context of this architecture, the engine of XRL/Flower does not change after extending XRL in this manner. To demonstrate the extensibility of XRL we discuss the effort it takes to add two new routing elements.

The first new routing element we want to extend XRL with is a generalization of both the `any_sequence` as the `parallel_sync` and is called `restricted_parallel_sync`. Basically, the `restricted_parallel_sync` is a `parallel_sync` restricted by a maximum number of branches that can execute in parallel. If this number equals the number of branches, the `restricted_parallel_sync` resembles a `parallel_sync`, if the number equals 1, it resembles an `any_sequence`. The `restricted_parallel_sync` can be used when the parallel branches share a limited number of identical resources. For example, a database might be involved in all branches, and only a number of connections can be made to that database. First, we add the following lines to the DTD shown in Figure 4:

```
<!ELEMENT restricted_parallel_sync ((%routing_element;)+)>
<!ATTLIST restricted_parallel_sync number NMTOKEN #REQUIRED>.
```

Note that the second line is added to specify the restrictive number of this element. The first line of the DTD is also changed to add this element to the list of routing elements. Second, the XSLT code of this routing element in terms of Petri nets is specified. The actual macro specification can be found in Append C, the complete XSLT specification in Appendix E, and Figure 24 shows the equivalent Petri net of the `restricted_parallel_sync`. The following macro snippet shows the core of the `restricted_parallel_sync`:

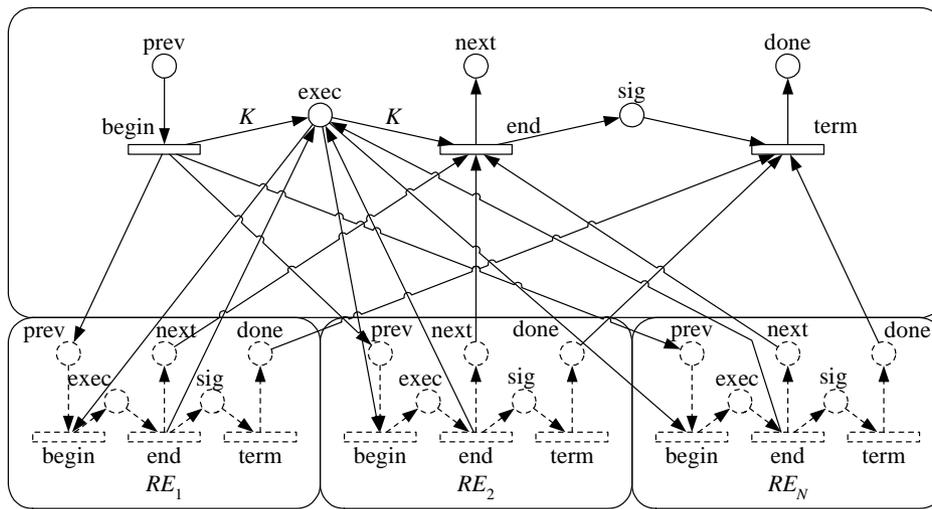


Figure 24. Semantics of `restricted_parallel_sync`.

```

rememberNumber()
forEveryChildRE()
...
switch()
caseNumber()
  arc(".",arc2_<xsl:number value="position()"/>,"..",begin,"..",exec)
  arc(".",arc3_<xsl:number value="position()"/>,"..",exec,"..",end)
endCase()
endSwitch()
endFor()

```

If K is the value of the number attribute, then the first K child routing elements add K arcs from the parent's *begin* transition to the parent's *exec* place, and K arcs from that *exec* place to the parent's *end* transition. Third, we verify the *restricted_parallel_sync*. This can be done in numerous ways. For example, we could take a representative set of existing sound XRL instances and replace every *parallel_sync* by an *restricted_parallel_sync*. After transforming these instances to PNML, we can verify them using Woflan. If all are sound, verification is complete. Fourth and last, we add the *restricted_parallel_sync* to the XSLT transformation library.

The second new routing element we want to extend XRL with is called *fast_sequence*, and demonstrates that tasks need not be atomic in XRL. This routing element executes a sequence of tasks such that when two tasks are in a *fast_sequence*, the second task can start before the preceding task is completed, provided the preceding task has started. Moreover, the second task can only complete if the preceding task is completed. The *fast_sequence* is frequently seen for long-lived activities that can span days or months. For instance, let the first task in the *fast_sequence* represent an approval process, while the second represents the preparation for the actual construction. We may want to specify that the preparation for the construction can start any time after the approval process has been started and that it cannot be completed before the approval process has completed, because only then all things are fixed. To add this routing element, we follow the steps as described above. First, we first add the following line to the DTD shown in Figure 4:

```
<!ELEMENT fast_sequence ((%routing_element;)+)>.
```

The first line of the DTD is also changed to add this element to the list of routing elements. Second, the XSLT code for this routing element is specified in terms of Petri nets. The actual macro specification can be found in Appendix C, the complete XSLT specification in Appendix E, and Figure 25 shows the

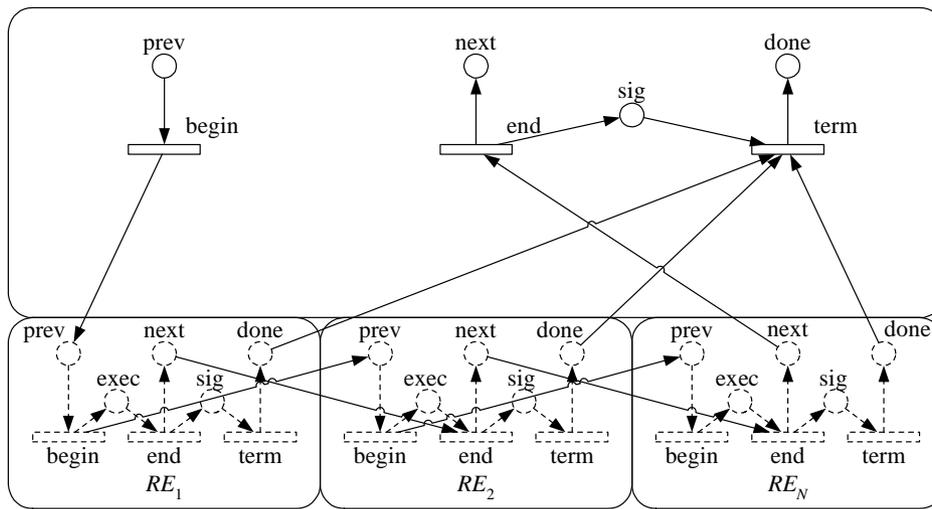


Figure 25. Semantics of *fast_sequence*.

equivalent Petri-net of the `fast_sequence`. The following macro snippet shows the core of the `fast_sequence`:

```

forEveryNonLastChildRE()
  arc("..",arc5_<xsl:number value="position()" />,".",begin,nextRE(),prev)
  arc("..",arc6_<xsl:number value="position()" />,".",next,nextRE(),end)
  ...
endFor()

```

For every child routing element except the last, arcs are added that connect its `begin` transition to the `prev` place of the next routing element, and its `next` place to the `end` transition of that next routing element. After verifying the `fast_sequence`, we add it to the XSLT transformation library.

As a third and final example, we extend XRL in such a way that an event can also be `reset`. Note that, although this extends the functionality of XRL, this is not really a new routing element. First, we extend the DTD with a `type` attribute for events:

```

<!ATTLIST event name ID #REQUIRED type (set|reset) "set">.

```

Possible values for the `type` attribute are `set` and `reset`, with `set` being the default value. This ensures backward compatibility: if unspecified, `set` is assumed. Second, we update the appropriate semantics. Figure 26 shows the updated semantics of the task. Depending on the `type` attribute, either the places `setevent` and `eventset` are connected to transition `events`, or the places `resetevent` and `eventreset`. Figure 27 shows the updated event on the global level. When the place `setevent` contains a token, the event is set, and the token is moved to `eventset`. When the place `resetevent` contains a token, the event is reset and the token is moved to `eventreset`. Note that it is possible to set and reset an event simultaneously, in which case it is impossible to tell whether the event is set or reset at the end. After verifying that the updates were specified correctly, the new specifications replace the old ones in the library. Note that, after adding this extension, the transformation from XRL to Petri nets does not necessarily result in a WF-net: If a certain event is only `reset` in a route, then the corresponding `event` place will be an additional source place. Because a route containing an event that can never be set is evidently erroneous, we don't regard this a problem. It even

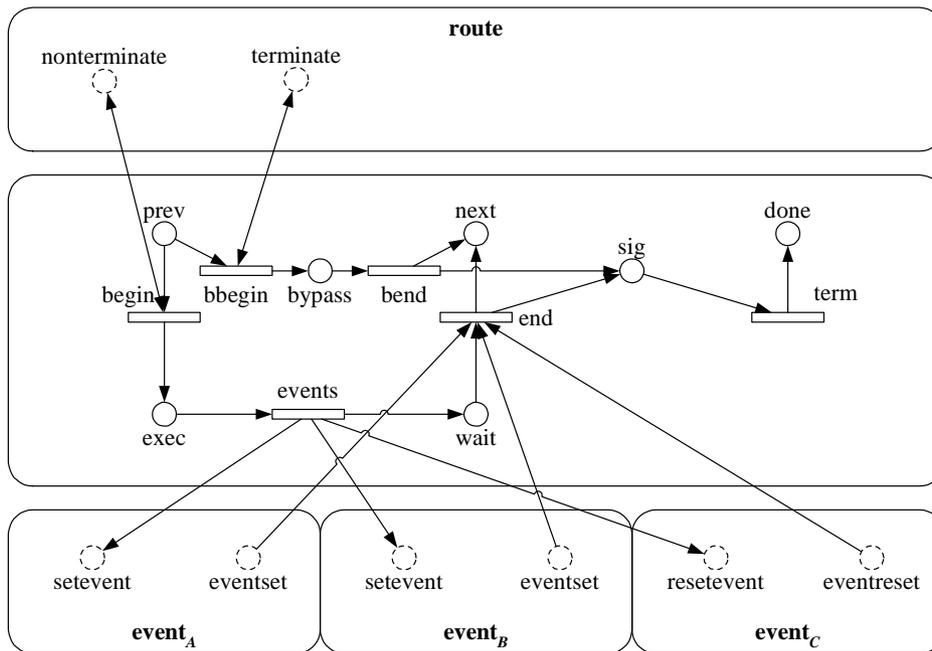


Figure 26. Semantics of task after extending event.

helps verifying the XRL route: If some *event* place happens to be an additional source place, then that event is only reset in the route.

After adding the XSLT specifications for both new XRL routing elements to the XSLT transformation library, and replacing the old specifications by the new ones for the third extension, both the enactment service (XRL/Flower) and the verification tool (XRL/Woflan) support the new or updated routing elements.

Similarly, other new application-specific routing elements may be added on-the-fly in this manner. Therefore, it is possible to create new application-specific workflow patterns by writing XSLT routines that describe the semantics of the pattern. The patterns can then be incorporated into the DTD of XRL after they have been tested and verified with Woflan. We are currently developing a complete methodology for workflow extensibility.

9. Verification tool

XRL/Woflan is based on our workflow verification tool Woflan [37, 38]. Woflan (<http://www.tm.tue.nl/it/woflan>) is designed as a workflow-management-system-independent analysis tool. In principle, it can interface with many workflow management systems. At present, Woflan can also interface with the workflow products COSA (Thiel Logistics AG/Software Ley), METEOR (LSDIS), and Staffware (Staffware), and the BPR tool Protos (Pallas Athena).

We have implemented the transformation from an XRL route into a WF-net using an XSLT specification. This XSLT specification produces a *Petri Net Markup Language* (PNML) file. PNML is the standard Petri net file format based on XML [22]. A second XSLT specification has been implemented that enables Woflan to read these PNML files.

The reduction rules based on the XRL structure have been implemented as an option in the first XSLT specification. The reduction rules based on the structure of the WF-net have been implemented as an option in Woflan.

PNML requires that every place, transition, and arc has a unique id. For this reason, the identification of these objects is an important implementation issue. For diagnostic purposes, it is vital that the names of the places and transitions are meaningful to the developer of the XRL route. We use the fact that a route is structured as a tree. For instance, we could identify a transition as transition *begin* of the second sequence of the first parallel_sync of the third true of the condition of the second while_do of the sequence of the route named *route*. To avoid long names as much as possible, we use abbreviations to identify the different routing elements. Table 2 lists these abbreviations. Figure 28 shows how a sequence is identified. The mode *id* is used to obtain the identification of any routing element. First, the sequence requests the id of its parent. Second, it adds “/s”. Third, it adds its rank among the sibling sequences. The example transition

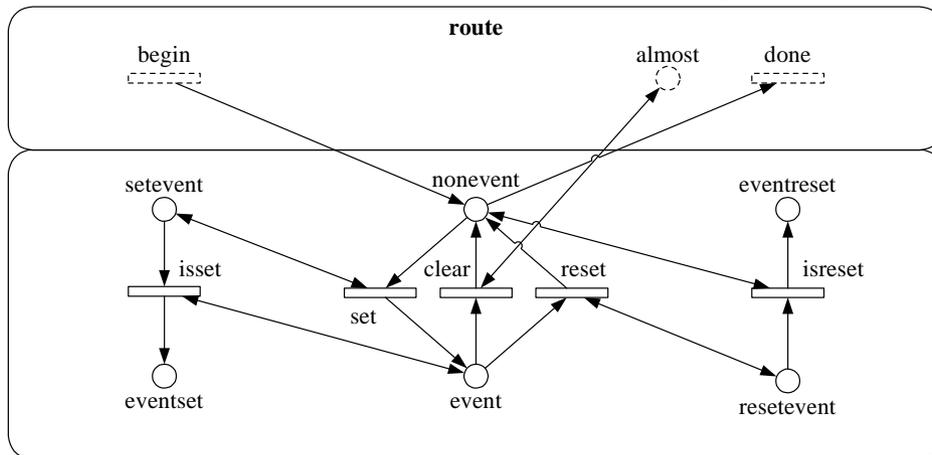


Figure 27. Semantics of event on global level after extending event.

mentioned is now identified by “route/s1/wd2/cn1/t3/ps1/s2/begin”. For most routing elements, this is fine, but we make an exception for tasks and events. These elements have explicit XRL names, and we want to use these names to identify them. Suppose the transition mentioned happens to be the transition *begin* of the task named *task*. Then it is identified by “route/task/begin”. Figure 29 shows how this is done. First, the task requests the name of the top element, that is, of the route. Second, it adds “/”. Third, it adds the value of its attribute *name*. Likewise, the transition *set* of the event named *event* is identified by “route/event/set”.

Figure 30 shows a screendump of Woflan after it diagnosed the example XRL route with all reductions applied. Clearly, after the transformation and both reductions, the example XRL route corresponds to a sound WF-net. As a result, we may conclude that the XRL route itself meets the soundness requirements. Therefore, we can take it into production safely.

10. Related work

The semantics of XRL is expressed in terms of Petri nets. Petri nets have been proposed for modeling workflow process definitions long before the term “workflow management” was coined and workflow management systems became readily available. Consider for example the work on Information Control Nets, a variant of the classical Petri nets, in the late seventies [16, 17]. The readers interested in the application of Petri nets to workflow management may refer to the two recent workshops on workflow management held in conjunction with the annual International Conference on Application and Theory of Petri Nets [7, 31] and an elaborate paper on workflow modeling using Petri nets [1].

Only a few papers in the literature focus on the verification of workflow process definitions. In [20] some verification issues have been examined and the complexity of selected correctness issues has been identified, but no concrete verification procedures have been suggested. In [1] and [11] concrete

Routing element	Abbreviation	Routing element	Abbreviation
Any_sequence	as	Parallel_sync	ps
Choice	ce	Restricted_parallel_sync	rps
Condition	cn	Sequence	s
Event	e	True	t
Event_ref	er	Task	tk
False	f	Timeout	to
Fast_sequence	fs	Terminate	tt
Parallel_no_sync	pns	Wait_all	wl
Parallel_part_sync	pps	Wait_any	wy
Parallel_part_sync_cancel	ppsc	While_do	wd

Table 2. Abbreviations.

```
<xsl:template match="sequence" mode="id">
  <xsl:apply-templates select=".." mode="id"/>/s<xsl:number/>
</xsl:template>
```

Figure 28. Identifying a sequence.

```
<xsl:template match="task" mode="id">
  <xsl:apply-templates select="/" mode="id"/>/<xsl:value-of
select="@name"/>
</xsl:template>
```

Figure 29. Identifying a task.

verification procedures based on Petri nets have been proposed. This paper builds upon the work presented in [1] where the concept of a sound WF-net was introduced (see Section 4). The technique presented in [11] has been developed for checking the consistency of transactional workflows including temporal constraints. However, the technique is restricted to acyclic workflows and only gives necessary conditions (that is, not sufficient conditions) for consistency. In [36] a reduction technique has been proposed. This reduction technique uses a correctness criterion which corresponds to soundness and the class of workflow processes considered are in essence acyclic free-choice Petri nets. Based on this reduction technique the analysis tool *FlowMake* [35] has been developed. Flowmake can interface with the IBM MQSeries Workflow product. Note that, although completely different techniques are used, Flowmake is very similar to the tool Woflan used in this paper [38].

This paper differs from the above approaches because the focus is on interorganizational workflows. Only a few papers explicitly focus on the problem of verifying the correctness of interorganizational workflows [2, 23]. In [2] the interaction between domains is specified in terms of message sequence charts and the actual overall workflow is checked with respect to these message sequence charts. A similar, but more formal and complete, approach is presented by Kindler, Martens, and Reisig in [23]. The authors give local criteria, using the concept of scenarios (similar to runs or basic message sequence charts), to guarantee the absence of certain anomalies at the global level.

Much work has been done on workflow transactions in the context of cross-organizational workflows, for example [18, 19, 34]. However, this work typically considers correctness issues at the task level rather than the process level. For example, the coordination model and the service model presented in [18] are not

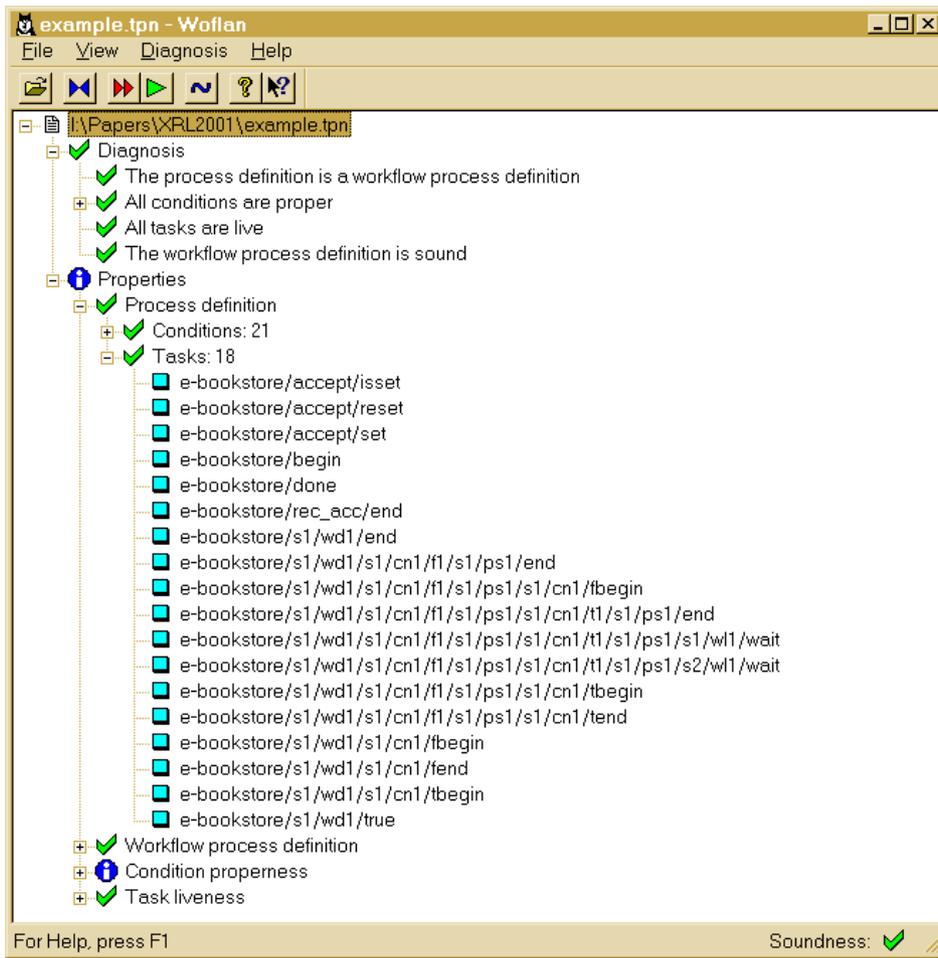


Figure 30. Woflan screendump.

explicitly addressing control flow problems resulting from causal relations (or the absence of such relations). The work conducted in projects such as CrossFlow [19], WISE [26], OSM [29], and COSMOS [30] is highly relevant for the enactment of interorganizational workflows. However, these projects do not consider the correctness issues tackled in this paper. Consider for example the Common Open Service Market (COSM) infrastructure proposed in [29, 30]. This infrastructure proposes mobile agents. The control-flow within each agent is managed by a Petri-net-based workflow engine. Unfortunately, this work does not address correctness issues at the process level.

In the last couple of years, many XML-based standards for electronic commerce have been proposed. The XML Common Business Library (xCBL) by CommerceOne, the Partner Interface Process (PIP) blueprints by RosettaNet, the Universal Description, Discovery and Integration (UDDI), the Electronic Business XML (ebXML) initiative by UN/CEFACT and OASIS, the Open Buying on the Internet (OBI) specification, the Open Application Group Integration Specification (OAGIS), and the BizTalk Framework are just some examples of the emerging standards based on XML. These standards primarily focus on the exchange of data and not on the control flow among organizations. Most of the standards provide standard DTDs or XML schemas for specific application domains (such as procurement). As mentioned in the introduction, one of the few initiatives which also addresses the control flow is RosettaNet. The Partner Interface Process (PIP) blueprints by RosettaNet do specify interactions using UML activity diagrams for the Business Operational View (BOV) and UML sequence diagrams for the Functional Service View (FSV) in addition to DTDs for data exchange. However, the PIP blueprints are not executable and need to be predefined. Like all of the standards, RosettaNet is primarily focusing on electronic markets with long-lasting prespecified relationships with one party (e.g. the market maker) imposing rigid business rules. XRL is not restricted to data-oriented static market places. XRL focuses on the process aspect and also consider others network structures (for example, the “first trading problem” [14, 27, 28]).

Finally, we would like to refer to two existing approaches toward interorganizational workflows based on Petri-nets. The first approach uses Documentary Petri Nets (DPN's), that is, a variant of high-level Petri nets with designated places for message exchange, to model and enact trade procedures [14, 27, 28]. The Interprocs system is based on these nets. The main difference between the Interprocs language and XRL is that XRL is instance based and supports less structured and more dynamic processes. Another approach combining Petri nets and interorganizational workflows is the P2P approach described in [10]. This approach uses inheritance to align local workflows. In [10] this approach is used to design an interorganizational workflow for a fictitious electronic bookstore similar to amazon.com or bn.com. A predecessor of the P2P approach has also been applied to an interorganizational workflow in the Telecom industry [5]. An interesting topic for future research is to see how the inheritance concepts used in [10, 5] translate to XRL. We would also like to develop a more formal methodology for verification of new workflow patterns like the ones we introduced in Section 8.

11. Conclusion

XRL is an XML-based language for describing workflow enactments. Woflan is a tool for the verification of Petri-net workflows. In this paper, we show how these two technologies can be combined together to create a powerful toolset for designing, verifying and implementing extensible workflows.

We have presented a novel way to verify the correctness of XRL routes. XRL routes are automatically transformed into WF-nets using XSLT. As a result, Woflan can be used to verify the correctness of the XRL route. The analysis procedure is optimized by exploiting dynamic properties of XRL routing elements and by using standard reduction rules at the Petri-net level [32]. We consider these verification capabilities essential for inter-organizational workflows. As was argued in the introduction, contemporary workflows need to be changed on the fly and sent across organizational boundaries. Unfortunately, the features also enable subtle, but highly disruptive, cross-organizational errors. On-the-fly changes and one-of-a-kind processes are destined to result in errors. Moreover, errors of a cross-organizational nature are

difficult to repair. Therefore, a language such as XRL (that is, a language with formal semantics) and verification tools such as XRL/Woflan are highly relevant for today's dynamic and networked economy.

References

- [1] W.M.P. van der Aalst. The Application of Petri Nets to Workflow Management. *The Journal of Circuits, Systems and Computers*, 8(1):21—66, 1998.
- [2] W.M.P. van der Aalst. Interorganizational Workflows: An Approach based on Message Sequence Charts and Petri Nets. *Systems Analysis - Modelling - Simulation*, 34(3):335—367, 1999.
- [3] W.M.P. van der Aalst. Loosely Coupled Interorganizational Workflows: Modeling and Analyzing Workflows Crossing Organizational Boundaries. *Information and Management*, 37(2):67—75, March 2000.
- [4] W.M.P. van der Aalst. Process-oriented Architectures for Electronic Commerce and Interorganizational Workflow. *Information Systems*, 24(8):639—671, 2000.
- [5] W.M.P. van der Aalst and K. Anyanwu. Inheritance of Interorganizational Workflows to Enable Business-to-Business E-commerce. In A. Dognac, E. van Heck, T. Saarinen, and et. al., editors, *Proceedings of the Second International Conference on Telecommunications and Electronic Commerce (ICTEC'99)*, pages 141—157, Nashville, Tennessee, October 1999.
- [6] W.M.P. van der Aalst, A.H.M. ter Hofstede, B. Kiepuszewski, and A.P. Barros. Advanced Workflow Patterns. In O. Etzion and P. Scheuermann, editors, *7th International Conference on Cooperative Information Systems (CoopIS 2000)*, volume 1901 of *Lecture Notes in Computer Science*, pages 18—29. Springer-Verlag, Berlin, 2000.
- [7] W.M.P. van der Aalst, G. De Michelis, and C.A. Ellis, editors. *Proceedings of Workflow Management: Net-based Concepts, Models, Techniques and Tools (WFM'98)*, Lisbon, Portugal, June 1998. UNINOVA, Lisbon.
- [8] W.M.P. van der Aalst and A. Kumar. XML Based Schema Definition for Support of Interorganizational Workflow. *Information Systems Research* (forthcoming).
- [9] W.M.P. van der Aalst, H.M.W. Verbeek, and A. Kumar. Verification of XRL: An XML-based Workflow Language. In W. Shen, Z. Lin, J.-P. Barthès, and M. Kamel, editors, *Proceedings of the Sixth International Conference on CSCW in Design (CSCWD 2001)*, pages 427—432, London, Ontario, Canada, July 2001.
- [10] W.M.P. van der Aalst and M. Weske. The P2P approach to Interorganizational Workflows. In K.R. Dittrich, A. Geppert, and M.C. Norrie, editors, *Proceedings of the 13th International Conference on Advanced Information Systems Engineering (CAiSE'01)*, volume 2068 of *Lecture Notes in Computer Science*, pages 140—156. Springer-Verlag, Berlin, 2001.
- [11] N.R. Adam, V. Atluri, and W. Huang. Modeling and Analysis of Workflows using Petri Nets. *Journal of Intelligent Information Systems*, 10(2):131—158, 1998.
- [12] Amazon.com, Inc. Amazon.com. <http://www.amazon.com>, 1999.
- [13] Barnes and Noble. bn.com. <http://www.bn.com>, 1999.
- [14] R.W.H. Bons, R.M. Lee, and R.W. Wagenaar. Designing trustworthy interorganizational trade procedures for open electronic commerce. *International Journal of Electronic Commerce*, 2(3):61—83, 1998.
- [15] T. Bray, J. Paoli, C.M. Sperberg-McQueen, and E. Maler. eXtensible Markup Language (XML) 1.0 (Second Edition). <http://www.w3.org/TR/REC-xml>, 2000.
- [16] C.A. Ellis. Information Control Nets: A Mathematical Model of Office Information Flow. In *Proceedings of the Conference on Simulation, Measurement and Modeling of Computer Systems*, pages 225—240, Boulder, Colorado, 1979. ACM Press.
- [17] C.A. Ellis and G.J. Nutt. Modelling and Enactment of Workflow Systems. In M. Ajmone Marsan, editor, *Application and Theory of Petri Nets 1993*, volume 691 of *Lecture Notes in Computer Science*, pages 1—16. Springer-Verlag, Berlin, 1993.

- [18] D. Georgakopoulos, H. Schuster, A. Cichocki, and D. Baker. Managing process and service fusion in virtual enterprises. *Information Systems*, 24(6):429—456, 1999.
- [19] P. Grefen, K. Aberer, Y. Hoffner, and H. Ludwig. CrossFlow: Cross-organizational Workflow Management in Dynamic Virtual Enterprises. *International Journal of Computer Systems, Science, and Engineering*, 15(5):277—290, 2001.
- [20] A.H.M. ter Hofstede, M.E. Orłowska, and J. Rajapakse. Verification Problems in Conceptual Workflow Specifications. *Data and Knowledge Engineering*, 24(3):239—256, 1998.
- [21] S. Jablonski and C. Bussler. *Workflow Management: Modeling Concepts, Architecture, and Implementation*. International Thomson Computer Press, London, UK, 1996.
- [22] M. Jungel, E. Kindler, and M. Weber. The Petri Net Markup Language. In S. Philippi, editor, *Proceedings of AWPN 2000 - 7th Workshop Algorithmen und Werkzeuge für Petrinetze*, pages 47—52. Research Report 7/2000, Institute for Computer Science, University of Koblenz, Germany, 2000.
- [23] E. Kindler, A. Martens, and W. Reisig. Inter-Operability of Workflow Applications: Local Criteria for Global Soundness. In W.M.P. van der Aalst, J. Desel, and A. Oberweis, editors, *Business Process Management: Models, Techniques, and Empirical Studies*, volume 1806 of *Lecture Notes in Computer Science*, pages 235—253. Springer-Verlag, Berlin, 2000.
- [24] A. Kumar and J.L. Zhao. Workflow Support for Electronic Commerce Applications. *Decision Support Systems* (forthcoming).
- [25] P. Lawrence, editor. *Workflow Handbook 1997, Workflow Management Coalition*. John Wiley and Sons, New York, 1997.
- [26] A. Lazcano, G. Alonso, H. Schuldt, and C. Schuler. The WISE Approach to Electronic Commerce. *International Journal of Computer Systems, Science, and Engineering*, 15(5):345—357, 2001.
- [27] R.M. Lee. Distributed Electronic Trade Scenarios: Representation, Design, Prototyping. *International Journal of Electronic Commerce*, 3(2):105—120, 1999.
- [28] R.M. Lee and R.W.H. Bons. Soft-Coded Trade Procedures for Open-edi. *International Journal of Electronic Commerce*, 1(1):27—49, 1996.
- [29] M. Merz, B. Liberman, and W. Lamersdorf. Using Mobile Agents to Support Interorganizational Workflow-Management. *International Journal on Applied Artificial Intelligence*, 11(6):551—572, 1997.
- [30] M. Merz, B. Liberman, and W. Lamersdorf. Crossing Organisational Boundaries with Mobile Agents in Electronic Service Markets. *Integrated Computer-Aided Engineering*, 6(2):91—104, 1999.
- [31] G. De Michelis, C. Ellis, and G. Memmi, editors. *Proceedings of the Second Workshop on Computer-Supported Cooperative Work, Petri nets and Related Formalisms*, Zaragoza, Spain, June 1994.
- [32] T. Murata. Petri Nets: Properties, Analysis and Applications. *Proceedings of the IEEE*, 77(4):541—580, April 1989.
- [33] W. Reisig and G. Rozenberg, editors. *Lectures on Petri Nets I: Basic Models*, volume 1491 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 1998.
- [34] A. Reuter and F. Schwenkreis. Contracts - a low-level mechanism for building general-purpose workflow management-systems. *Data Engineering Bulletin*, 18(1):4—10, 1995.
- [35] W. Sadiq and M.E. Orłowska. FlowMake Product Information, Distributed Systems Technology Centre, Queensland, Australia. <http://www.dstc.edu.au/Research/Projects/FlowMake/productinfo/index.html>.
- [36] W. Sadiq and M.E. Orłowska. Analyzing Process Models using Graph Reduction Techniques. *Information Systems*, 25(2):117—134, 2000.
- [37] H.M.W. Verbeek and W.M.P. van der Aalst. Woflan 2.0: A Petri-net-based Workflow Diagnosis Tool. In M. Nielsen and D. Simpson, editors, *Application and Theory of Petri Nets 2000*, volume 1825 of *Lecture Notes in Computer Science*, pages 475—484. Springer-Verlag, Berlin, 2000.
- [38] H.M.W. Verbeek, T. Basten, and W.M.P. van der Aalst. Diagnosing Workflow Processes Using Woflan. *The Computer Journal*, 44(4):246—279. British Computer Society, 2001.

Appendix A. The extended DTD of XRL

This appendix shows the DTD of XRL *after* all extensions have been added. For sake of clarity, the differences with the original DTD (see Figure 4) are emphasized.

```
<!ENTITY % routing_element
"task|sequence|any_sequence|choice|condition|fast_sequence|parallel_sync|parallel_no_s
ync|parallel_part_sync|parallel_part_sync_cancel|restricted_parallel_sync|wait_all|wai
t_any|while_do|terminate">
<!ELEMENT route ((%routing_element;), event*)>
<!ATTLIST route
  name ID #REQUIRED
  created_by CDATA #IMPLIED
  date CDATA #IMPLIED>
<!ELEMENT task (event*)>
<!ATTLIST task
  name ID #REQUIRED
  address CDATA #REQUIRED
  role CDATA #IMPLIED
  doc_read NMTOKENS #IMPLIED
  doc_update NMTOKENS #IMPLIED
  doc_create NMTOKENS #IMPLIED
  result CDATA #IMPLIED
  status (ready|running|enabled|disabled|aborted|null) #IMPLIED
  start_time NMTOKENS #IMPLIED
  end_time NMTOKENS #IMPLIED
  notify CDATA #IMPLIED>
<!ELEMENT event EMPTY>
<!ATTLIST event
  name ID #REQUIRED
  type (set/reset) "set">
<!ELEMENT sequence ((%routing_element;|state)+)>
<!ELEMENT any_sequence ((%routing_element;)+)>
<!ELEMENT choice ((%routing_element;)+)>
<!ELEMENT condition ((true|false)*)>
<!ATTLIST condition
  condition CDATA #REQUIRED>
<!ELEMENT true (%routing_element;)>
<!ELEMENT false (%routing_element;)>
<!ELEMENT fast_sequence ((%routing_element;)+)>
<!ELEMENT parallel_sync ((%routing_element;)+)>
<!ELEMENT parallel_no_sync ((%routing_element;)+)>
<!ELEMENT parallel_part_sync ((%routing_element;)+)>
<!ATTLIST parallel_part_sync
  number NMTOKEN #REQUIRED>
<!ELEMENT parallel_part_sync_cancel ((%routing_element;)+)>
<!ATTLIST parallel_part_sync_cancel
  number NMTOKEN #REQUIRED>
<!ELEMENT restricted_parallel_sync ((%routing_element;)+)>
<!ATTLIST restricted_parallel_sync
  number NMTOKEN #REQUIRED>
<!ELEMENT wait_all ((event_ref|timeout)+)>
<!ELEMENT wait_any ((event_ref|timeout)+)>
<!ELEMENT event_ref EMPTY>
<!ATTLIST event_ref
  name IDREF #REQUIRED>
<!ELEMENT timeout ((%routing_element;)?)>
<!ATTLIST timeout
  time CDATA #REQUIRED
```

```

    type (relative|s_relative|absolute) "absolute">
<!ELEMENT while_do (%routing_element;)>
<!ATTLIST while_do
    condition CDATA #REQUIRED>
<!ELEMENT terminate EMPTY>
<!ELEMENT state EMPTY>

```

Appendix B. The E-bookstore example

This appendix shows the XRL route for processing a customer order by the e-bookstore. Note that most attributes have been omitted for brevity.

```

<!DOCTYPE route SYSTEM "xrl.dtd">
<route name="e-bookstore" created_by="H.M.W. Verbeek" date="June 11, 2001">
  <sequence>
    <task name="place_c_order" address="customer"/>
    <task name="handle_c_order" address="bookstore"/>
    <while_do condition="No publisher found yet">
      <sequence>
        <task name="place_b_order" address="bookstore"/>
        <task name="eval_b_order" address="publisher"/>
        <condition condition="No publisher found yet">
          <true>
            <sequence>
              <task name="decide" address="publisher"/>
              <condition condition="Try alternative publisher">
                <true>
                  <task name="alt_publ" address="publisher"/>
                </true>
                <false>
                  <sequence>
                    <task name="b_reject" address="publisher"/>
                    <task name="c_reject" address="bookstore"/>
                    <task name="rec_decl" address="customer"/>
                  </sequence>
                </false>
              </condition>
            </sequence>
          </true>
          <false>
            <sequence>
              <task name="b_accept" address="publisher"/>
              <task name="c_accept" address="bookstore"/>
              <parallel_sync>
                <task name="rec_acc" address="customer">
                  <event name="accept"/>
                </task>
              </parallel_sync>
            </sequence>
          <while_do condition="No shipper found yet">
            <sequence>
              <task name="s_request" address="bookstore"/>
              <task name="eval_s_req" address="shipper"/>
            </sequence>
          </while_do>
          <condition condition="Shipper found">
            <true>
              <sequence>
                <task name="s_accept" address="shipper"/>
                <task name="inform_publ" address="bookstore"/>
              </sequence>
            </true>
          </condition>
        </condition>
      </sequence>
    </while_do>
  </sequence>

```

```

        <task name="prepare_b" address="publisher"/>
        <task name="send_book" address="publisher"/>
        <task name="prepare_s" address="shipper"/>
        <task name="ship" address="shipper"/>
        <parallel_sync>
            <sequence>
                <task name="notify" address="shipper"/>
                <task name="send_bill" address="bookstore"/>
                <wait_all>
                    <event_ref name="accept"/>
                </wait_all>
                <task name="rec_bill" address="customer"/>
            </sequence>
            <sequence>
                <wait_all>
                    <event_ref name="accept"/>
                </wait_all>
                <task name="rec_book" address="customer"/>
            </sequence>
        </parallel_sync>
        <task name="pay" address="customer"/>
        <task name="handle_payment" address="bookstore"/>
    </sequence>
</true>
<false>
    <task name="s_reject" address="shipper"/>
</false>
</condition>
</sequence>
</parallel_sync>
</sequence>
</false>
</condition>
</sequence>
</while_do>
</sequence>
</route>

```

Appendix C. Extensions, using macros

This appendix shows the macro specifications of the `restricted_parallel_sync`, the `fast_sequence`, and the event *after* all three extensions have been added.

restricted_parallel_sync

```

<xsl:template match="restricted_parallel_sync">
    countSuspects()
    countSpawners()
    localTransition(begin)
    localTransition(end)
    localPlace(prev)
    localPlace(next)
    localPlace(exec)
    localArc(arc1,prev,begin)
    localArc(arc4,end,next)
    switch()
        caseAnySuspects()
            <xsl:apply-templates/>

```

```

rememberNumber()
forEveryChildRE()
  countSpawnersChild()
  arc("..",arc5_<xsl:number value="position()"/>,"..",begin,".",prev)
  arc("..",arc6_<xsl:number value="position()"/>,"..",exec,".",begin)
  arc("..",arc7_<xsl:number value="position()"/>,".",next,"..",end)
  arc("..",arc8_<xsl:number value="position()"/>,".",end,"..",exec)
  switch()
    caseAnySpawnersChild()
      arc("..",arc9_<xsl:number value="position()"/>,".",done,"..",term)
    endCase()
  endSwitch()
  switch()
    caseNumber()
      arc("..",arc2_<xsl:number value="position()"/>,"..",begin,"..",exec)
      arc("..",arc3_<xsl:number value="position()"/>,"..",exec,"..",end)
    endCase()
  endSwitch()
endFor()
endCase()
endSwitch()
switch()
  caseAnySpawners()
    localTransition(term)
    localPlace(done)
    localPlace(sig)
    localArc(arc10,end,sig)
    localArc(arc11,sig,term)
    localArc(arc12,term,done)
  endCase()
endSwitch()
</xsl:template>

<xsl:template match="restricted_parallel_sync" mode="id">
  <xsl:apply-templates select=".." mode="id"/>/rps:rankRE()
</xsl:template>

```

fast_sequence

```

<xsl:template match="fast_sequence">
  countSuspects()
  countSpawners()
  localTransition(begin)
  localTransition(end)
  localPlace(prev)
  localPlace(next)
  localArc(arc1,prev,begin)
  localArc(arc2,end,next)
  switch()
    caseAnySuspects()
      <xsl:apply-templates/>
      arc(".",arc3,".",begin,firstChildRE(),prev)
      arc(".",arc4,lastChildRE(),next,".",end)
    forEveryNonLastChildRE()
      arc("..",arc5_<xsl:number value="position()"/>,".",begin,nextRE(),prev)
      arc("..",arc6_<xsl:number value="position()"/>,".",next,nextRE(),end)
    switch()
      caseAnyTerminates()
        checkSelfTask()

```

```

        switch()
        caseSelfTask()
            arc("..",arc7_<xsl:number value="position()"/>,".",bbegin,nextRE(),prev)
        endCase()
    endSwitch()
    checkSuccessorTask()
    switch()
        caseSuccessorTask()
            arc("..",arc8_<xsl:number value="position()"/>,".",next,nextRE(),bend)
        endCase()
    endSwitch()
    endCase()
endSwitch()
endFor()
forEveryChildRE()
    countSpawnersChild()
    switch()
        caseAnySpawnersChild()
            arc("..",arc12_<xsl:number value="position()"/>,".",done,"..",term)
        endCase()
    endSwitch()
endFor()
endCase()
default()
    localPlace(exec)
    localArc(arc10,begin,exec)
    localArc(arc11,exec,end)
endDefault()
endSwitch()
switch()
    caseAnySpawners()
        localTransition(term)
        localPlace(done)
        localPlace(sig)
        localArc(arc7,end,sig)
        localArc(arc8,sig,term)
        localArc(arc9,term,done)
    endCase()
endSwitch()
</xsl:template>

<xsl:template match="fast_sequence" mode="id">
    <xsl:apply-templates select=".." mode="id"/>/fs[rankRE()]
</xsl:template>

```

event

```

<xsl:template match="event">
    rankEvent()
    switch()
        caseReset()
            mixedArc(".", "id", arc1, "..", "id", events, ".", "eid", resetevent)
            mixedArc(".", "id", arc2, ".", "eid", eventreset, "..", "id", end)
        endCase()
    default()
        mixedArc(".", "id", arc1, "..", "id", events, ".", "eid", setevent)
        mixedArc(".", "id", arc2, ".", "eid", eventset, "..", "id", end)
    endDefault()

```

```

endSwitch()
switch()
  caseEventRankedFirst()
    countSetsAndResets()
    eventTransition(clear)
    eventPlace(event)
    eventPlace(nonevent)
    eventArc(arc1,event,clear)
    eventArc(arc2,clear,nonevent)
    mixedArc(".", "eid", arc3, "/", "id", almost, ".", "eid", clear)
    mixedArc(".", "eid", arc4, ".", "eid", clear, "/", "id", almost)
    mixedArc(".", "eid", arc5, ".", "eid", nonevent, "/", "id", done)
    mixedArc(".", "eid", arc6, "/", "id", begin, ".", "eid", nonevent)
  ifSets()
    eventTransition(set)
    eventTransition(isset)
    eventPlace(setevent)
    eventPlace(eventset)
    eventArc(arc7, setevent, set)
    eventArc(arc8, setevent, isset)
    eventArc(arc9, nonevent, set)
    eventArc(arc10, set, event)
    eventArc(arc11, event, isset)
    eventArc(arc12, isset, event)
    eventArc(arc13, set, setevent)
    eventArc(arc14, isset, eventset)
  endIf()
  ifResets()
    eventTransition(reset)
    eventTransition(isreset)
    eventPlace(resetevent)
    eventPlace(eventreset)
    eventArc(arc17, resetevent, reset)
    eventArc(arc18, resetevent, isreset)
    eventArc(arc19, nonevent, reset)
    eventArc(arc20, reset, event)
    eventArc(arc21, event, isreset)
    eventArc(arc22, isreset, event)
    eventArc(arc23, reset, resetevent)
    eventArc(arc24, isreset, eventreset)
  endIf()
  endCase()
endSwitch()
</xsl:template>

<xsl:template match="event" mode="id">
  <xsl:apply-templates select=".." mode="id"/>e<xsl:number/>
</xsl:template>

<xsl:template match="event" mode="eid">
  <xsl:apply-templates select="/" mode="id"/><xsl:value-of select="@name"/>
</xsl:template>

```

Appendix D. Macro definitions

This appendix shows the macro definitions *after* the `restricted_parallel_sync`, the `fast_sequence`, and the event extensions have been added.

```
#define parentTransition(self) \
```

```

<transition><xsl:attribute name="id"><xsl:apply-templates select=".." mode="id"/>/self</xsl:attribute></transition>

#define localTransition(self) \
<transition><xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/self</xsl:attribute><xsl:apply-templates select="attribute::condition"/></transition>

#define eventTransition(self) \
<transition><xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/self</xsl:attribute></transition>

#define localNamedTransition(self) \
<transition><xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/self</xsl:attribute><xsl:apply-templates select="attribute::role"/><xsl:apply-templates select="attribute::address"/><name><value>&quot;<xsl:value-of select="@name"/>/self&quot;</value></name></transition>

#define place(parent,self) \
<place><xsl:attribute name="id"><xsl:apply-templates select=parent mode="id"/>/self</xsl:attribute></place>

#define localPlace(self) \
<place><xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/self</xsl:attribute></place>

#define eventPlace(self) \
<place><xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/self</xsl:attribute></place>

#define localNamedPlace(self) \
<place><xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/self</xsl:attribute><name><value>&quot;<xsl:value-of select="@name"/>/self&quot;</value></name></place>

#define localMarkedPlace(self) \
<place><xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/self</xsl:attribute><initialMarking><value>1</value></initialMarking></place>

#define arc(aparent,aself,sparent,sself,tparent,tself) \
<arc><xsl:attribute name="id"><xsl:apply-templates select=aparent mode="id"/>/aself</xsl:attribute><xsl:attribute name="source"><xsl:apply-templates select=sparent mode="id"/>/sself</xsl:attribute><xsl:attribute name="target"><xsl:apply-templates select=tparent mode="id"/>/tself</xsl:attribute></arc>

#define mixedArc(aparent,aid,aself,sparent,sid,sself,tparent,tid,tself) \
<arc><xsl:attribute name="id"><xsl:apply-templates select=aparent mode=aid/>/aself</xsl:attribute><xsl:attribute name="source"><xsl:apply-templates select=sparent mode=sid/>/sself</xsl:attribute><xsl:attribute name="target"><xsl:apply-templates select=tparent mode=tid/>/tself</xsl:attribute></arc>

#define localArc(aself,sself,tself) \
<arc><xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/aself</xsl:attribute><xsl:attribute name="source"><xsl:apply-templates select="." mode="id"/>/sself</xsl:attribute><xsl:attribute name="target"><xsl:apply-templates select="." mode="id"/>/tself</xsl:attribute></arc>

#define eventArc(aself,sself,tself) \
<arc><xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/aself</xsl:attribute><xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/

```

```

>/sself</xsl:attribute><xsl:attribute name="target"><xsl:apply-templates select="."
mode="eid"/>/tself</xsl:attribute></arc>

#define rememberNumber() \
<xsl:variable name="rememberedNumber"><xsl:value-of select="@number"/></xsl:variable>

#define caseNumber() \
<xsl:when test="position() &lt;= $rememberedNumber">

#define rememberSelf() \
<xsl:variable name="rememberedSelf"><xsl:apply-templates select="." mode="id"/></
xsl:variable>

#define rememberedSelf() \
"$rememberedSelf"

#define countEvents() \
<xsl:variable name="numberOfEvents"><xsl:value-of
select="count(child::*[self::event])"/></xsl:variable>

#define caseAnyEvents() \
<xsl:when test="$reduce = 'No' or $numberOfEvents &gt; 0">

#define countSuspects() \
<xsl:variable name="numberOfSuspects"><xsl:value-of select="count(descen-
dant::*[self::event or self::wait_all or self::wait_any or self::terminate])"/></
xsl:variable>

#define caseAnySuspects() \
<xsl:when test="$reduce = 'No' or $numberOfSuspects &gt; 0">

#define countSpawners() \
<xsl:variable name="numberOfSpawners"><xsl:value-of select="count(descen-
dant::*[self::parallel_part_sync or self::parallel_part_sync_cancel or
self::parallel_no_sync])"/></xsl:variable>

#define caseAnySpawners() \
<xsl:when test="$reduce = 'No' or $numberOfSpawners &gt; 0">

#define countSpawnersChild() \
<xsl:variable name="numberOfSpawnersChild"><xsl:value-of select="count(descendant-or-
self::*[self::parallel_part_sync or self::parallel_part_sync_cancel or
self::parallel_no_sync])"/></xsl:variable>

#define caseAnySpawnersChild() \
<xsl:when test="$reduce = 'No' or $numberOfSpawnersChild &gt; 0">

#define countTerminates() \
<xsl:variable name="numberOfTerminates"><xsl:value-of select="count(descen-
dant::*[self::terminate])"/></xsl:variable>

#define caseAnyTerminates() \
<xsl:when test="$reduce = 'No' or $numberOfTerminates &gt; 0">

#define rankEvent() \
<xsl:variable name="nameOfEvent"><xsl:value-of select="@name"/></xsl:vari-
able><xsl:variable name="rankOfEvent"><xsl:value-of select="count(preced-
ing::event[@name=$nameOfEvent]|ancestor::event[@name=$nameOfEvent])"/></xsl:variable>

```

```

#define caseEventRankedFirst() \
<xsl:when test="$rankOfEvent = 0">

#define switch() \
<xsl:choose>

#define caseNoReduction() \
<xsl:when test="$reduce = 'No'">

#define endCase() \
</xsl:when>

#define default() \
<xsl:otherwise>

#define endDefault() \
</xsl:otherwise>

#define endSwitch() \
</xsl:choose>

#define rankRE() \
<xsl:value-of select="1 + count(preceding-sibling::*[self::sequence or
self::any_sequence or self::fast_sequence or self::choice or self::condition or
self::parallel_sync or self::parallel_no_sync or self::parallel_part_sync or
self::parallel_part_sync_cancel or self::restricted_parallel_sync or self::wait_all or
self::wait_any or self::while_do or self::task or self::terminate])"/>

#define firstChildRE() \
"child::*[self::sequence or self::any_sequence or self::fast_sequence or self::choice
or self::condition or self::parallel_sync or self::parallel_no_sync or
self::parallel_part_sync or self::parallel_part_sync_cancel or
self::restricted_parallel_sync or self::wait_all or self::wait_any or self::while_do or
self::task or self::terminate][1]"

#define lastChildRE() \
"child::*[self::sequence or self::any_sequence or self::fast_sequence or self::choice
or self::condition or self::parallel_sync or self::parallel_no_sync or
self::parallel_part_sync or self::parallel_part_sync_cancel or
self::restricted_parallel_sync or self::wait_all or self::wait_any or self::while_do or
self::task or self::terminate][last()]"

#define nextRE() \
"following-sibling::*[self::sequence or self::any_sequence or self::fast_sequence or
self::choice or self::condition or self::parallel_sync or self::parallel_no_sync or
self::parallel_part_sync or self::parallel_part_sync_cancel or
self::restricted_parallel_sync or self::wait_all or self::wait_any or self::while_do or
self::task or self::terminate][1]"

#define forEveryChildRE() \
<xsl:for-each select="child::*[self::sequence or self::any_sequence or
self::fast_sequence or self::choice or self::condition or self::parallel_sync or
self::parallel_no_sync or self::parallel_part_sync or self::parallel_part_sync_cancel
or self::restricted_parallel_sync or self::wait_all or self::wait_any or self::while_do
or self::task or self::terminate]">

#define forEveryEventRef() \
<xsl:for-each select="child::*[self::event_ref]">

```

```

#define forEveryTimeout() \
<xsl:for-each select="child::*[self::timeout]">

#define forEveryPrecedingTask() \
<xsl:for-each select="preceding-sibling::*[self::task]">

#define forEveryFollowingTask() \
<xsl:for-each select="following-sibling::*[self::task]">

#define forEveryNonLastChildRE() \
<xsl:for-each select="child::*[self::sequence or self::any_sequence or
self::fast_sequence or self::choice or self::condition or self::parallel_sync or
self::parallel_no_sync or self::parallel_part_sync or self::parallel_part_sync_cancel
or self::restricted_parallel_sync or self::wait_all or self::wait_any or self::while_do
or self::task or self::terminate] [position() != last()]">

#define forEveryPrecedingRE() \
<xsl:for-each select="preceding-sibling::*[self::sequence or self::any_sequence or
self::fast_sequence or self::choice or self::condition or self::parallel_sync or
self::parallel_no_sync or self::parallel_part_sync or self::parallel_part_sync_cancel
or self::restricted_parallel_sync or self::wait_all or self::wait_any or self::while_do
or self::task or self::terminate]">

#define forEveryFollowingRE() \
<xsl:for-each select="following-sibling::*[self::sequence or self::any_sequence or
self::fast_sequence or self::choice or self::condition or self::parallel_sync or
self::parallel_no_sync or self::parallel_part_sync or self::parallel_part_sync_cancel
or \
self::restricted_parallel_sync or self::wait_all or self::wait_any or self::while_do or
self::task or self::terminate]">

#define endFor() \
</xsl:for-each>

#define checkSelfTask() \
<xsl:variable name="oneIfTask"><xsl:value-of select="count(self::*[self::task])"/></
xsl:variable>

#define caseSelfTask() \
<xsl:when test="$oneIfTask > 0">

#define checkSuccessorTask() \
<xsl:variable name="oneIfSuccessorTask"><xsl:value-of select="count(following-sib-
ling::*[self::sequence or self::any_sequence or self::fast_sequence or self::choice or
self::condition or self::parallel_sync or self::parallel_no_sync or
self::parallel_part_sync or self::parallel_part_sync_cancel or
self::restricted_parallel_sync or self::wait_all or self::wait_any or self::while_do or
self::task or self::terminate] [1] [self::task])"/></xsl:variable>

#define caseSuccessorTask() \
<xsl:when test="$oneIfSuccessorTask > 0">

#define countSetsAndResets() \
<xsl:variable name="numberOfSets">\
<xsl:value-of select="count(/route//event[@name=$nameOfEvent])"/>\
</xsl:variable>\
<xsl:variable name="numberOfResets">\
<xsl:value-of select="count(/route//event[@name=$nameOfEvent and @type='reset'])"/>\
</xsl:variable>

```

```

#endif

#define ifSets() \
<xsl:if test="$numberOfSets > $numberOfResets">

#define ifResets() \
<xsl:if test="$numberOfResets > 0">

#define endIf() \
</xsl:if>

#define caseReset() \
<xsl:when test="@type = 'reset' ">

```

Appendix E. Extensions, plain XSLT

This appendix shows the plain XSLT specifications of the `restricted_parallel_sync`, the `fast_sequence`, and the event *after* all three extensions have been added.

restricted_parallel_sync

```

<xsl:template match="restricted_parallel_sync">
  <xsl:variable name="numberOfSuspects">
    <xsl:value-of select="count(descendant::*[self::event or self::send or
self::receive or self::wait_all or self::wait_any or self::terminate])"/>
  </xsl:variable>
  <xsl:variable name="numberOfSpawners">
    <xsl:value-of select="count(descendant::*[self::parallel_part_sync or
self::parallel_part_sync_cancel or self::parallel_no_sync])"/>
  </xsl:variable>
  <transition>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/begin</
xsl:attribute>
    <xsl:apply-templates select="attribute::condition"/>
  </transition>
  <transition>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/end</
xsl:attribute>
    <xsl:apply-templates select="attribute::condition"/>
  </transition>
  <place>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/prev</
xsl:attribute>
  </place>
  <place>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/next</
xsl:attribute>
  </place>
  <place>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/exec</
xsl:attribute>
  </place>
  <arc>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/arc1</
xsl:attribute>
    <xsl:attribute name="source"><xsl:apply-templates select="." mode="id"/>/prev</
xsl:attribute>
    <xsl:attribute name="target"><xsl:apply-templates select="." mode="id"/>/begin</

```

```

xsl:attribute>
  </arc>
  <arc>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/arc4</
xsl:attribute>
  <xsl:attribute name="source"><xsl:apply-templates select="." mode="id"/>/end</
xsl:attribute>
  <xsl:attribute name="target"><xsl:apply-templates select="." mode="id"/>/next</
xsl:attribute>
  </arc>
  <xsl:choose>
    <xsl:when test="$reduce = 'No' or $numberOfSuspects > 0">
      <xsl:apply-templates/>
      <xsl:variable name="rememberedNumber">
        <xsl:value-of select="@number"/>
      </xsl:variable>
      <xsl:for-each select="child::*[self::sequence or self::any_sequence or
self::fast_sequence or self::choice or self::condition or self::parallel_sync or
self::parallel_no_sync or self::parallel_part_sync or self::parallel_part_sync_cancel
or self::restricted_parallel_sync or self::wait_all or self::wait_any or self::while_do
or self::task or self::terminate]">
        <xsl:variable name="numberOfSpawnersChild">
          <xsl:value-of select="count(descendant-or-self::*[self::parallel_part_sync
or self::parallel_part_sync_cancel or self::parallel_no_sync])"/>
        </xsl:variable>
        <arc>
          <xsl:attribute name="id"><xsl:apply-templates select=".." mode="id"/>/
arc5_<xsl:number value="position()"/></xsl:attribute>
          <xsl:attribute name="source"><xsl:apply-templates select=".." mode="id"/>/
begin</xsl:attribute>
          <xsl:attribute name="target"><xsl:apply-templates select="." mode="id"/>/
prev</xsl:attribute>
          </arc>
        <arc>
          <xsl:attribute name="id"><xsl:apply-templates select=".." mode="id"/>/
arc6_<xsl:number value="position()"/></xsl:attribute>
          <xsl:attribute name="source"><xsl:apply-templates select=".." mode="id"/>/
exec</xsl:attribute>
          <xsl:attribute name="target"><xsl:apply-templates select="." mode="id"/>/
begin</xsl:attribute>
          </arc>
        <arc>
          <xsl:attribute name="id"><xsl:apply-templates select=".." mode="id"/>/
arc7_<xsl:number value="position()"/></xsl:attribute>
          <xsl:attribute name="source"><xsl:apply-templates select="." mode="id"/>/
next</xsl:attribute>
          <xsl:attribute name="target"><xsl:apply-templates select=".." mode="id"/>/
end</xsl:attribute>
          </arc>
        <arc>
          <xsl:attribute name="id"><xsl:apply-templates select=".." mode="id"/>/
arc8_<xsl:number value="position()"/></xsl:attribute>
          <xsl:attribute name="source"><xsl:apply-templates select="." mode="id"/>/
end</xsl:attribute>
          <xsl:attribute name="target"><xsl:apply-templates select=".." mode="id"/>/
exec</xsl:attribute>
          </arc>
        <xsl:choose>
          <xsl:when test="$reduce = 'No' or $numberOfSpawnersChild > 0">

```

```

        <arc>
            <xsl:attribute name="id"><xsl:apply-templates select=".." mode="id"/>/
arc9_<xsl:number value="position()"/></xsl:attribute>
            <xsl:attribute name="source"><xsl:apply-templates select="." mode="id"/>/
done</xsl:attribute>
            <xsl:attribute name="target"><xsl:apply-templates select=".." mode="id"/
>/term</xsl:attribute>
        </arc>
    </xsl:when>
</xsl:choose>
<xsl:choose>
    <xsl:when test="position() &lt;= $rememberedNumber">
        <arc>
            <xsl:attribute name="id"><xsl:apply-templates select=".." mode="id"/>/
arc2_<xsl:number value="position()"/></xsl:attribute>
            <xsl:attribute name="source"><xsl:apply-templates select=".." mode="id"/
>/begin</xsl:attribute>
            <xsl:attribute name="target"><xsl:apply-templates select=".." mode="id"/
>/exec</xsl:attribute>
        </arc>
        <arc>
            <xsl:attribute name="id"><xsl:apply-templates select=".." mode="id"/>/
arc3_<xsl:number value="position()"/></xsl:attribute>
            <xsl:attribute name="source"><xsl:apply-templates select=".." mode="id"/
>/exec</xsl:attribute>
            <xsl:attribute name="target"><xsl:apply-templates select=".." mode="id"/
>/end</xsl:attribute>
        </arc>
    </xsl:when>
</xsl:choose>
</xsl:for-each>
</xsl:when>
</xsl:choose>
<xsl:choose>
    <xsl:when test="$reduce = 'No' or $numberOfSpawners > 0">
        <transition>
            <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/term</
xsl:attribute>
            <xsl:apply-templates select="attribute::condition"/>
        </transition>
        <place>
            <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/done</
xsl:attribute>
        </place>
        <place>
            <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/sig</
xsl:attribute>
        </place>
        <arc>
            <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/arc10</
xsl:attribute>
            <xsl:attribute name="source"><xsl:apply-templates select="." mode="id"/>/end</
xsl:attribute>
            <xsl:attribute name="target"><xsl:apply-templates select="." mode="id"/>/sig</
xsl:attribute>
        </arc>
        <arc>
            <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/arc11</
xsl:attribute>

```

```

        <xsl:attribute name="source"><xsl:apply-templates select="." mode="id"/>/sig</
xsl:attribute>
        <xsl:attribute name="target"><xsl:apply-templates select="." mode="id"/>/term</
xsl:attribute>
        </arc>
        <arc>
        <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/arc12</
xsl:attribute>
        <xsl:attribute name="source"><xsl:apply-templates select="." mode="id"/>/term</
xsl:attribute>
        <xsl:attribute name="target"><xsl:apply-templates select="." mode="id"/>/done</
xsl:attribute>
        </arc>
        </xsl:when>
        </xsl:choose>
</xsl:template>
<xsl:template match="restricted_parallel_sync" mode="id">
  <xsl:apply-templates select="." mode="id"/>/rps:<xsl:value-of select="1 + count(pre-
ceding-sibling::*[self::sequence or self::any_sequence or self::fast_sequence or
self::choice or self::condition or self::parallel_sync or self::parallel_no_sync or
self::parallel_part_sync or self::parallel_part_sync_cancel or
self::restricted_parallel_sync or self::wait_all or self::wait_any or self::while_do or
self::task or self::terminate] )"/>
</xsl:template>

```

fast_sequence

```

<xsl:template match="fast_sequence">
  <xsl:variable name="numberOfSuspects">
    <xsl:value-of select="count(descendant::*[self::event or self::send or
self::receive or self::wait_all or self::wait_any or self::terminate] )"/>
  </xsl:variable>
  <xsl:variable name="numberOfSpawners">
    <xsl:value-of select="count(descendant::*[self::parallel_part_sync or
self::parallel_part_sync_cancel or self::parallel_no_sync] )"/>
  </xsl:variable>
  <transition>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/begin</
xsl:attribute>
    <xsl:apply-templates select="attribute::condition"/>
  </transition>
  <transition>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/end</
xsl:attribute>
    <xsl:apply-templates select="attribute::condition"/>
  </transition>
  <place>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/prev</
xsl:attribute>
  </place>
  <place>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/next</
xsl:attribute>
  </place>
  <arc>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/arc1</
xsl:attribute>
    <xsl:attribute name="source"><xsl:apply-templates select="." mode="id"/>/prev</

```

```

xsl:attribute>
  <xsl:attribute name="target"><xsl:apply-templates select="." mode="id"/>/begin</
xsl:attribute>
  </arc>
  <arc>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/arc2</
xsl:attribute>
  <xsl:attribute name="source"><xsl:apply-templates select="." mode="id"/>/end</
xsl:attribute>
  <xsl:attribute name="target"><xsl:apply-templates select="." mode="id"/>/next</
xsl:attribute>
  </arc>
  <xsl:choose>
    <xsl:when test="$reduce = 'No' or $numberOfSuspects > 0">
      <xsl:apply-templates/>
      <arc>
        <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/arc3</
xsl:attribute>
        <xsl:attribute name="source"><xsl:apply-templates select="." mode="id"/>/
begin</xsl:attribute>
        <xsl:attribute name="target"><xsl:apply-templates
select="child::*[self::sequence or self::any_sequence or self::fast_sequence or
self::choice or self::condition or self::parallel_sync or self::parallel_no_sync or
self::parallel_part_sync or self::parallel_part_sync_cancel or
self::restricted_parallel_sync or self::wait_all or self::wait_any or self::while_do or
self::task or self::terminate] [1]" mode="id"/>/prev</xsl:attribute>
        </arc>
        <arc>
          <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/arc4</
xsl:attribute>
          <xsl:attribute name="source"><xsl:apply-templates
select="child::*[self::sequence or self::any_sequence or self::fast_sequence or
self::choice or self::condition or self::parallel_sync or self::parallel_no_sync or
self::parallel_part_sync or self::parallel_part_sync_cancel or
self::restricted_parallel_sync or self::wait_all or self::wait_any or self::while_do or
self::task or self::terminate] [last()]" mode="id"/>/next</xsl:attribute>
          <xsl:attribute name="target"><xsl:apply-templates select="." mode="id"/>/end</
xsl:attribute>
          </arc>
          <xsl:for-each select="child::*[self::sequence or self::any_sequence or
self::fast_sequence or self::choice or self::condition or self::parallel_sync or
self::parallel_no_sync or self::parallel_part_sync or self::parallel_part_sync_cancel
or self::restricted_parallel_sync or self::wait_all or self::wait_any or self::while_do
or self::task or self::terminate] [position() != last()]">
            <arc>
              <xsl:attribute name="id"><xsl:apply-templates select=".." mode="id"/>/
arc5_<xsl:number value="position()"/></xsl:attribute>
              <xsl:attribute name="source"><xsl:apply-templates select="." mode="id"/>/
begin</xsl:attribute>
              <xsl:attribute name="target"><xsl:apply-templates select="following-sib-
ling::*[self::sequence or self::any_sequence or self::fast_sequence or self::choice or
self::condition or self::parallel_sync or self::parallel_no_sync or
self::parallel_part_sync or self::parallel_part_sync_cancel or
self::restricted_parallel_sync or self::wait_all or self::wait_any or self::while_do or
self::task or self::terminate] [1]" mode="id"/>/prev</xsl:attribute>
              </arc>
              <arc>
                <xsl:attribute name="id"><xsl:apply-templates select=".." mode="id"/>/
arc6_<xsl:number value="position()"/></xsl:attribute>

```

```

        <xsl:attribute name="source"><xsl:apply-templates select="." mode="id"/>/
next</xsl:attribute>
        <xsl:attribute name="target"><xsl:apply-templates select="following-sib-
ling::*[self::sequence or self::any_sequence or self::fast_sequence or self::choice or
self::condition or self::parallel_sync or self::parallel_no_sync or
self::parallel_part_sync or self::parallel_part_sync_cancel or
self::restricted_parallel_sync or self::wait_all or self::wait_any or self::while_do or
self::task or self::terminate] [1]" mode="id"/>/end</xsl:attribute>
    </arc>
    <xsl:choose>
        <xsl:when test="$reduce = 'No' or $numberOfTerminates > 0">
            <xsl:variable name="oneIfTask">
                <xsl:value-of select="count(self::*[self::task])"/>
            </xsl:variable>
            <xsl:choose>
                <xsl:when test="$oneIfTask > 0">
                    <arc>
                        <xsl:attribute name="id"><xsl:apply-templates select=".." mode="id"/
>/arc7_<xsl:number value="position()"/></xsl:attribute>
                        <xsl:attribute name="source"><xsl:apply-templates select="."
mode="id"/>/begin</xsl:attribute>
                            <xsl:attribute name="target"><xsl:apply-templates select="following-
sibling::*[self::sequence or self::any_sequence or self::fast_sequence or self::choice
or self::condition or self::parallel_sync or self::parallel_no_sync or
self::parallel_part_sync or self::parallel_part_sync_cancel or
self::restricted_parallel_sync or self::wait_all or self::wait_any or self::while_do or
self::task or self::terminate] [1]" mode="id"/>/prev</xsl:attribute>
                                </arc>
                            </xsl:when>
                        </xsl:choose>
                        <xsl:variable name="oneIfSuccessorTask">
                            <xsl:value-of select="count(following-sibling::*[self::sequence or
self::any_sequence or self::fast_sequence or self::choice or self::condition or
self::parallel_sync or self::parallel_no_sync or self::parallel_part_sync or
self::parallel_part_sync_cancel or self::restricted_parallel_sync or self::wait_all or
self::wait_any or self::while_do or self::task or self::terminate] [1][self::task])"/>
                        </xsl:variable>
                        <xsl:choose>
                            <xsl:when test="$oneIfSuccessorTask > 0">
                                <arc>
                                    <xsl:attribute name="id"><xsl:apply-templates select=".." mode="id"/
>/arc8_<xsl:number value="position()"/></xsl:attribute>
                                    <xsl:attribute name="source"><xsl:apply-templates select="."
mode="id"/>/next</xsl:attribute>
                                        <xsl:attribute name="target"><xsl:apply-templates select="following-
sibling::*[self::sequence or self::any_sequence or self::fast_sequence or self::choice
or self::condition or self::parallel_sync or self::parallel_no_sync or
self::parallel_part_sync or self::parallel_part_sync_cancel or
self::restricted_parallel_sync or self::wait_all or self::wait_any or self::while_do or
self::task or self::terminate] [1]" mode="id"/>/bend</xsl:attribute>
                                            </arc>
                                        </xsl:when>
                                    </xsl:choose>
                                </xsl:when>
                            </xsl:choose>
                        </xsl:for-each>
                        <xsl:for-each select="child::*[self::sequence or self::any_sequence or
self::fast_sequence or self::choice or self::condition or self::parallel_sync or
self::parallel_no_sync or self::parallel_part_sync or self::parallel_part_sync_cancel

```

```

or self::restricted_parallel_sync or self::wait_all or self::wait_any or self::while_do
or self::task or self::terminate]">
  <xsl:variable name="numberOfSpawnersChild">
    <xsl:value-of select="count(descendant-or-self::*[self::parallel_part_sync
or self::parallel_part_sync_cancel or self::parallel_no_sync])"/>
  </xsl:variable>
  <xsl:choose>
    <xsl:when test="$reduce = 'No' or $numberOfSpawnersChild > 0">
      <arc>
        <xsl:attribute name="id"><xsl:apply-templates select=".." mode="id"/>/
arc12_<xsl:number value="position()"/></xsl:attribute>
        <xsl:attribute name="source"><xsl:apply-templates select="." mode="id"/>/
done</xsl:attribute>
        <xsl:attribute name="target"><xsl:apply-templates select=".." mode="id"/
>/term</xsl:attribute>
      </arc>
    </xsl:when>
  </xsl:choose>
</xsl:for-each>
</xsl:when>
<xsl:otherwise>
  <place>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/exec</
xsl:attribute>
    </place>
  <arc>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/arc10</
xsl:attribute>
    <xsl:attribute name="source"><xsl:apply-templates select="." mode="id"/>/
begin</xsl:attribute>
    <xsl:attribute name="target"><xsl:apply-templates select="." mode="id"/>/exec</
xsl:attribute>
    </arc>
  <arc>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/arc11</
xsl:attribute>
    <xsl:attribute name="source"><xsl:apply-templates select="." mode="id"/>/exec</
xsl:attribute>
    <xsl:attribute name="target"><xsl:apply-templates select="." mode="id"/>/end</
xsl:attribute>
    </arc>
  </xsl:otherwise>
</xsl:choose>
<xsl:choose>
  <xsl:when test="$reduce = 'No' or $numberOfSpawners > 0">
    <transition>
      <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/term</
xsl:attribute>
      <xsl:apply-templates select="attribute::condition"/>
    </transition>
    <place>
      <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/done</
xsl:attribute>
    </place>
    <place>
      <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/sig</
xsl:attribute>
    </place>
  <arc>

```

```

        <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/arc7</
xsl:attribute>
        <xsl:attribute name="source"><xsl:apply-templates select="." mode="id"/>/end</
xsl:attribute>
        <xsl:attribute name="target"><xsl:apply-templates select="." mode="id"/>/sig</
xsl:attribute>
        </arc>
        <arc>
            <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/arc8</
xsl:attribute>
            <xsl:attribute name="source"><xsl:apply-templates select="." mode="id"/>/sig</
xsl:attribute>
            <xsl:attribute name="target"><xsl:apply-templates select="." mode="id"/>/term</
xsl:attribute>
            </arc>
        <arc>
            <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/arc9</
xsl:attribute>
            <xsl:attribute name="source"><xsl:apply-templates select="." mode="id"/>/term</
xsl:attribute>
            <xsl:attribute name="target"><xsl:apply-templates select="." mode="id"/>/done</
xsl:attribute>
            </arc>
        </xsl:when>
    </xsl:choose>
</xsl:template>
<xsl:template match="fast_sequence" mode="id">
    <xsl:apply-templates select=".." mode="id"/>/fs[<xsl:value-of select="1 + count(pre-
ceding-sibling::*[self::sequence or self::any_sequence or self::fast_sequence or
self::choice or self::condition or self::parallel_sync or self::parallel_no_sync or
self::parallel_part_sync or self::parallel_part_sync_cancel or
self::restricted_parallel_sync or self::wait_all or self::wait_any or self::while_do or
self::task or self::terminate])"/>]
</xsl:template>

```

event

```

<xsl:template match="event">
    <xsl:variable name="nameOfEvent">
        <xsl:value-of select="@name"/>
    </xsl:variable>
    <xsl:variable name="rankOfEvent">
        <xsl:value-of select="count(preceding::event[@name=$nameOfEvent]|ances-
tor::event[@name=$nameOfEvent])"/>
    </xsl:variable>
    <xsl:choose>
        <xsl:when test="@type = 'reset'">
            <arc>
                <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/arc1</
xsl:attribute>
                <xsl:attribute name="source"><xsl:apply-templates select=".." mode="id"/>/
events</xsl:attribute>
                <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/
resetevent</xsl:attribute>
            </arc>
        <arc>
            <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/arc2</
xsl:attribute>
            <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/even-

```

```

treset</xsl:attribute>
  <xsl:attribute name="target"><xsl:apply-templates select=".." mode="id"/>/end</
xsl:attribute>
  </arc>
</xsl:when>
<xsl:otherwise>
  <arc>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/arc1</
xsl:attribute>
    <xsl:attribute name="source"><xsl:apply-templates select=".." mode="id"/>/
events</xsl:attribute>
    <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/set-
event</xsl:attribute>
    </arc>
  <arc>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="id"/>/arc2</
xsl:attribute>
    <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/
eventset</xsl:attribute>
    <xsl:attribute name="target"><xsl:apply-templates select=".." mode="id"/>/end</
xsl:attribute>
    </arc>
  </xsl:otherwise>
</xsl:choose>
<xsl:choose>
  <xsl:when test="$rankOfEvent = 0">
    <xsl:variable name="numberOfSets">
      <xsl:value-of select="count (/route//event [@name=$nameOfEvent])" />
    </xsl:variable>
    <xsl:variable name="numberOfResets">
      <xsl:value-of select="count (/route//event [@name=$nameOfEvent and
@type='reset'])" />
    </xsl:variable>
    <transition>
      <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/clear</
xsl:attribute>
    </transition>
  <place>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/event</
xsl:attribute>
  </place>
  <place>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/non-
event</xsl:attribute>
  </place>
  <arc>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc1</
xsl:attribute>
    <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/
event</xsl:attribute>
    <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/
clear</xsl:attribute>
    </arc>
  <arc>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc2</
xsl:attribute>
    <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/
clear</xsl:attribute>
    <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/non-

```

```

event</xsl:attribute>
  </arc>
  <arc>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc3</
xsl:attribute>
    <xsl:attribute name="source"><xsl:apply-templates select="/" mode="id"/>/
almost</xsl:attribute>
    <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/
clear</xsl:attribute>
  </arc>
  <arc>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc4</
xsl:attribute>
    <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/
clear</xsl:attribute>
    <xsl:attribute name="target"><xsl:apply-templates select="/" mode="id"/>/
almost</xsl:attribute>
  </arc>
  <arc>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc5</
xsl:attribute>
    <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/non-
event</xsl:attribute>
    <xsl:attribute name="target"><xsl:apply-templates select="/" mode="id"/>/done</
xsl:attribute>
  </arc>
  <arc>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc6</
xsl:attribute>
    <xsl:attribute name="source"><xsl:apply-templates select="/" mode="id"/>/
begin</xsl:attribute>
    <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/non-
event</xsl:attribute>
  </arc>
  <xsl:if test="$numberOfSets > $numberOfResets">
    <transition>
      <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/set</
xsl:attribute>
    </transition>
    <transition>
      <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/isset</
xsl:attribute>
    </transition>
    <place>
      <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/set-
event</xsl:attribute>
    </place>
    <place>
      <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/event-
set</xsl:attribute>
    </place>
  </arc>
  <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc7</
xsl:attribute>
    <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/
setevent</xsl:attribute>
    <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/
set</xsl:attribute>
  </arc>

```

```

        <arc>
            <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc8</
xsl:attribute>
            <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/
setevent</xsl:attribute>
            <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/
isset</xsl:attribute>
        </arc>
        <arc>
            <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc9</
xsl:attribute>
            <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/
nonevent</xsl:attribute>
            <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/
set</xsl:attribute>
        </arc>
        <arc>
            <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc10</
xsl:attribute>
            <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/
set</xsl:attribute>
            <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/
event</xsl:attribute>
        </arc>
        <arc>
            <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc11</
xsl:attribute>
            <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/
event</xsl:attribute>
            <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/
isset</xsl:attribute>
        </arc>
        <arc>
            <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc12</
xsl:attribute>
            <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/
isset</xsl:attribute>
            <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/
event</xsl:attribute>
        </arc>
        <arc>
            <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc13</
xsl:attribute>
            <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/
set</xsl:attribute>
            <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/
setevent</xsl:attribute>
        </arc>
        <arc>
            <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc14</
xsl:attribute>
            <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/
isset</xsl:attribute>
            <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/
eventset</xsl:attribute>
        </arc>
    </xsl:if>
    <xsl:if test="$numberOfResets > 0">
        <transition>

```

```

        <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/reset</
xsl:attribute>
    </transition>
    <transition>
        <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/isre-
set</xsl:attribute>
    </transition>
    <place>
        <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/
resetevent</xsl:attribute>
    </place>
    <place>
        <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/even-
treset</xsl:attribute>
    </place>
    <arc>
        <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc17</
xsl:attribute>
        <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/
resetevent</xsl:attribute>
        <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/
reset</xsl:attribute>
    </arc>
    <arc>
        <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc18</
xsl:attribute>
        <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/
resetevent</xsl:attribute>
        <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/
isreset</xsl:attribute>
    </arc>
    <arc>
        <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc19</
xsl:attribute>
        <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/
nonevent</xsl:attribute>
        <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/
reset</xsl:attribute>
    </arc>
    <arc>
        <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc20</
xsl:attribute>
        <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/
reset</xsl:attribute>
        <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/
event</xsl:attribute>
    </arc>
    <arc>
        <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc21</
xsl:attribute>
        <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/
event</xsl:attribute>
        <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/
isreset</xsl:attribute>
    </arc>
    <arc>
        <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc22</
xsl:attribute>
        <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/

```

```

isreset</xsl:attribute>
  <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/
event</xsl:attribute>
  </arc>
  <arc>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc23</
xsl:attribute>
    <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/
reset</xsl:attribute>
    <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/
resetevent</xsl:attribute>
    </arc>
  </arc>
    <xsl:attribute name="id"><xsl:apply-templates select="." mode="eid"/>/arc24</
xsl:attribute>
    <xsl:attribute name="source"><xsl:apply-templates select="." mode="eid"/>/
isreset</xsl:attribute>
    <xsl:attribute name="target"><xsl:apply-templates select="." mode="eid"/>/
eventreset</xsl:attribute>
    </arc>
  </xsl:if>
</xsl:when>
</xsl:choose>
</xsl:template>
<xsl:template match="event" mode="id">
  <xsl:apply-templates select=".." mode="id"/>/e<xsl:number/>
</xsl:template>
<xsl:template match="event" mode="eid">
  <xsl:apply-templates select="/" mode="id"/>/<xsl:value-of select="@name"/>
</xsl:template>

```

Appendix F. Coverage of routing elements

For sake of simplicity, we assume that the invariants of the child routing elements are recalibrated such that the places *next* and *done* have weight 1. As a result, some weights might be fractions instead of natural numbers, but this is irrelevant: If need be, we can always multiply all weights with the product of all denominators.

Any_sequence:

$$\begin{aligned}
P_{as} &= (N + 1) \times (\text{prev} + \text{next}) + \left(\sum_{i=1}^N P_{RE_i} + RE_{i_{exec}} \right) + \text{exec} \\
Q_{as} &= (2N + 1) \times (\text{prev} + \text{done}) + \left(\sum_{i=1}^N P_{RE_i} + Q_{RE_i} + RE_{i_{exec}} \right) + \text{exec} + (N + 1) \times \text{sig}
\end{aligned}$$

Choice:

$$P_{ce} = N \times (\text{prev} + \text{join} + \text{next}) + \left(\sum_{i=1}^N N \times P_{RE_i} - (N - 1) \times RE_{i_{prev}} \right)$$

$$Q_{ce} = 2N \times (\text{prev} + \text{done}) + \left(\sum_{i=1}^N N \times (P_{RE_i} + Q_{RE_i}) - (N-1) \times RE_{i_{\text{prev}}} \right) + N \times (\text{join} + \text{next} + \text{sig})$$

Condition, with T true's and F false's:

$$P_{cn} = (T + F + 1) \times (\text{prev} + \text{split} + \text{join} + \text{next}) + \left(\sum_{i=1}^T P_{\text{true}_{RE_i}} \right) + (F + 1) \times \text{texec} +$$

$$\left(\sum_{i=1}^F P_{\text{false}_{RE_i}} \right) + (T + 1) \times \text{fexec}$$

$$Q_{cn} = (2T + 2F + 2) \times (\text{prev} + \text{split} + \text{done}) + \left(\sum_{i=1}^T P_{\text{true}_{RE_i}} + Q_{\text{true}_{RE_i}} \right) +$$

$$(2F + 2) \times \text{texec} + (F + 1) \times \text{tsig} + \left(\sum_{i=1}^F P_{\text{false}_{RE_i}} + Q_{\text{false}_{RE_i}} \right) +$$

$$(2T + 2) \times \text{fexec} + (T + 1) \times \text{fsig} + (T + F + 1) \times (\text{join} + \text{sig} + \text{joindone})$$

Parallel_no_sync:

$$P_{pns} = \text{prev} + \text{exec} + \text{next}$$

$$Q_{pns} = (2N + 1) \times (\text{prev} + \text{done}) + \left(\sum_{i=1}^N P_{RE_i} + Q_{RE_i} \right) + \text{exec} + \text{sig}$$

Parallel_part_sync or parallel_part_sync_cancel:

$$P_{pps} = P_{ppsc} = \text{prev} + \text{exec} + \text{next}$$

$$Q_{pps} = Q_{ppsc} = 2N \times (\text{prev} + \text{done}) + \text{free} + \left(\sum_{i=1}^N P_{RE_i} + Q_{RE_i} \right) + \text{exec} + \text{count} + (K + 1) \times \text{sig}$$

Parallel_sync:

$$P_{ps} = N \times (\text{prev} + \text{next}) + \left(\sum_{i=1}^N P_{RE_i} \right)$$

$$Q_{ps} = 2N \times (\text{prev} + \text{done}) + \left(\sum_{i=1}^N P_{RE_i} + Q_{RE_i} \right) + N \times \text{sig}$$

Sequence:

$$P_s = \text{prev} + \left(\sum_{i=1}^N P_{RE_i} \right) + \text{next}$$

$$Q_s = (N+1) \times (\text{prev} + \text{done}) + \left(\sum_{i=1}^N (N+1-i) \times P_{RE_i} + Q_{RE_i} \right) + \text{sig}$$

Task:

$$P_{tk} = \text{prev} + \text{exec} + \text{wait} + \text{bypass} + \text{next}$$

$$Q_{tk} = \text{prev} + \text{exec} + \text{wait} + \text{bypass} + \text{sig} + \text{done}$$

Terminate:

$$P_{tt} = \text{prev} + \text{exec} + \text{next}$$

$$Q_{tt} = \text{prev} + \text{exec} + \text{sig} + \text{done}$$

Wait_all or wait_any, with N timeouts:

$$P_{wl} = P_{wy} = \text{prev} + \text{wprev} + \text{wnext} + \text{next}$$

$$Q_{wl} = Q_{wy} = 2 \times (\text{prev} + \text{wprev} + \text{done}) + \left(\sum_{i=1}^N P_{\text{timeout}_{RE_i}} + Q_{\text{timeout}_{RE_i}} \right) + \text{wnext} + \text{wdone} + \text{sig}$$

While_do:

$$P_{wd} = \text{prev} + P_{RE} + \text{next}$$

$$Q_{wd} = (X+1) \times (\text{prev} + \text{done}) + P_{RE} + Q_{RE} + \text{sig}$$