

Automated Risk Mitigation in Business Processes (extended version)

R. Conforti¹, A.H.M. ter Hofstede^{1,2}, M. La Rosa¹, and M. Adams¹

¹ Queensland University of Technology, Australia

{raffaele.conforti,a.terhofstede,m.larosa,mj.adams}@qut.edu.au

² Eindhoven University of Technology, The Netherlands

Abstract. This paper proposes a concrete approach for the automatic mitigation of risks that are detected during process enactment. Given a process model affected by risks, e.g. a financial process exposed to the risk of approval fraud, we enact this process and as soon as the likelihood of the associated risk(s) is no longer tolerable, we generate a set of possible mitigation actions to reduce the risks' likelihood, ideally annulling the risks altogether. A mitigation action is a sequence of controlled changes applied to the running process instance, taking into account a snapshot of the process resources and data, and the current status of the system in which the process is executed. These actions are proposed as recommendations to help process administrators mitigate process-related risks as soon as they arise. The approach has been implemented in the YAWL environment and its performance evaluated. The results show that it is possible to mitigate process-related risks within a few minutes.

1 Introduction

Business processes in various sectors such as financial, healthcare and oil&gas, are constantly exposed to a wide range of risks. Take for example the BP oil spill in 2010 which resulted in an environmental disaster, or the fraud at Société Générale in 2008, which led to a €4.9B loss.

A *process-related risk* measures the likelihood and the consequence that something happening will impact on the process objectives [27]. Failing to address process-related risks can result in substantial financial and reputational consequences, potentially threatening an organization's existence, like in the case of Société Générale. There is thus an increasing need to better manage business process risks, as also highlighted by legislative initiatives like Basel II [8] and the Sarbanes-Oxley Act.³ Organizations are attempting to incorporate process-related risks as a distinct view in their operational management, seeking effective ways to *control* such risks. However, whilst conceptually appealing, to date there is little guidance as to how this can be concretely done.

In previous work [11], we presented a mechanism to model risks in executable business process models and detect them as early as possible during process execution. Unfortunately, detecting a risk in time is often not enough to avoid the negative outcome associated. A *prompt* risk mitigation should be taken to restore the process instance to

³ www.gpo.gov/fdsys/pkg/PLAW-107pub1204

a safe state, before the instance progresses any further. Moreover, taking the *right* mitigation at the right time may make the difference between success and failure. In fact, the number of possible ways a process-related risk may be mitigated is potentially very large that is difficult for a process administrator to take the right decision at the right time, without any support. One has to consider all mitigations that are possible, given the current state of the process instance (including a snapshot of the associated data and resources), and the context in which the instance is running, i.e. the state of other running instances, to make such a decision. For example, in order to mitigate the risk of a process instance A to run overtime, a mitigation may entail to reallocate resources from a process instance B (potentially of another process) to A.

In light of the above, in this paper we propose a technique for automatically mitigating process-related risks. Since a process instance may be affected by multiple risks at the same time, we treat this problem as a *multi-objective optimization problem*. A solution to this problem is a variant of the risky process instance obtained by applying a sequence of mitigation actions, in order to reduce the risks' probability down to a tolerable level, or in the best case, to annul the risks altogether. Mitigation actions include control-flow aspects (e.g. skipping a task to be executed), process resources (e.g. reallocating a resource to a different task), and data (e.g. rolling back an executed task to restore its input data). To explore the potentially large solution space, we use dominance-based Multi-Objective Simulated Annealing (MOSA) [26]. At each run, the algorithm generates a small set of solutions similar to the original process instance but with less risks. It stops when either a maximum number of *non-redundant* solutions (i.e. solutions proposing different mitigations) is found or a given timeframe elapses. This approach is not meant to replace human judgement. Instead, it aims to support process administrators in deciding what mitigations to take, by reducing the number of feasible options, and consequently the time needed to take a decision.

We defined the mitigation actions in collaboration with an Australian risk consultant. To prove the feasibility of this approach, we implemented these actions and the MOSA algorithm on top of the YAWL system. We instantiated a set of process models, inspired by an industry standard [28] and a process model from the area of screen business [24], affected by one or more risks, and executed a series of tests to mitigate such risks. The tests show that the technique can find a set of possible solutions within a few minutes of computation, and that in all cases the associated risks are mitigated.

The rest of this paper is organized as follows. Section 2 introduces the required background concepts in the context of an example. Section 3 describes the proposed technique to mitigate process risks which is then evaluated in Section 4. Section 5 covers related work and Section 6 concludes the paper. Appendix A describes the process models used in the experiments; Appendix B provides the definition of the risks associated with these process models; Appendix C provides the solutions to one of the experiments in this paper.

2 Background and running example

Our technique for risk mitigation is part of a holistic approach for managing process-related risks throughout the process lifecycle. Accordingly, the four phases of the traditional BPM lifecycle (Design, Implementation, Enactment and Analysis) [13] are each

extended to incorporate elements of risk, as shown in Fig. 1. First, in a *Risk Identification* phase, the process model to be designed is analyzed for potential risks. Established risk analysis methods such as Fault Tree Analysis [10] or Root Cause Analysis [19] can be employed in this phase. The output of this phase is a set of risks, each expressed as a *risk condition* that describes the set of events that lead to a potential fault occurrence. Then, in the Design phase, these high-level risk conditions are mapped to process model-specific aspects. The result of this phase is a risk-annotated process model. Next, in the Implementation phase, these conditions are linked to workflow-specific aspects, such as the content of data variables and resource allocation states. The model is then executed by a risk-aware process engine in the Enactment phase. To evaluate risk conditions in this phase, we need to consider the current state of all running instances of any process (and not only the instance for which we are computing the risk condition), the resources that are busy and available, and the values of the data variables being created and consumed. Moreover, we need to consider historical data, i.e. the archived execution data of all previous instances of the process. Finally, the Diagnosis phase involves risk monitoring and controlling, which can trigger changes in the current process instance, to *mitigate* the likelihood of a fault occurring, or in the underlying process model, to *prevent* a given risk from occurring in future instances. This risk mitigation phase is the focus of this paper.

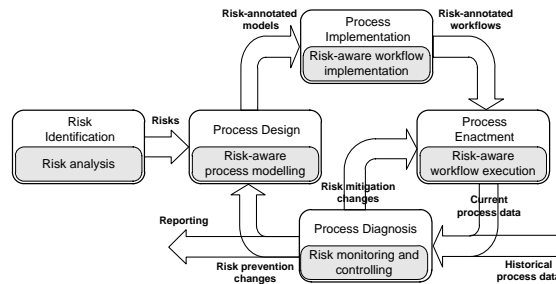


Fig. 1. Risk-aware BPM lifecycle.

Let us now consider an example process for which we have defined several risks, to understand how risk conditions can be formulated in terms of process model elements. These conditions will provide input for the risk mitigation technique presented in the next section. The example process, shown in Figure 2, describes a *Payment* subprocess of an order fulfillment process, inspired by the VICS industry standard for logistics [28]. It begins after freight has been picked up by a carrier and deals with the payment of shipment and freight costs executed in parallel. The freight payment whether required starts with the production of a freight invoice done by a Supply Admin Officer and ends with the processing of the freight payment. During the payment of shipment costs, first a Shipment Invoice is produced for costs related to a specific order. If payment has been made in advance, a Finance Officer simply issues a Shipment Remittance Advice to the customer specifying the amount paid. Otherwise, the Finance Officer issues a Shipment Payment Order, which requires approval by a Senior Finance Officer (a superior of the Finance Officer) who may request amendments be made by the Finance Officer that issued the Order. After the document is finalized and the customer has paid, an Account Manager can process the payment. If the customer underpays, the Account Manager issues a Debit Adjustment, the customer makes a further payment and the payment is

reprocessed. If a customer overpays, the Account Manager issues a Credit Adjustment. In the latter case and in the case of correct payment, the Payment subprocess completes.

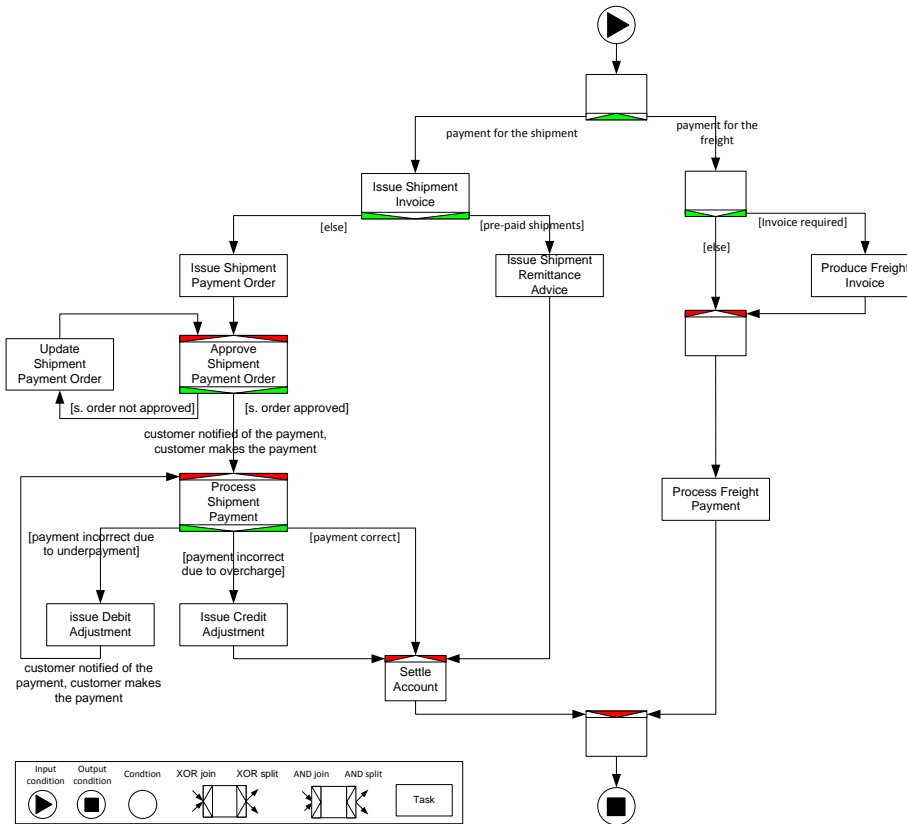


Fig. 2. Order-Fulfillment: Payment subprocess (using YAWL [18] notation).

In this process, we can identify various faults that may occur during execution. For example, a Service Level Agreement (SLA) may establish that the process (or one of its tasks) may not last longer than a Maximum Cycle Time MCT (e.g. 5 days), otherwise a pecuniary penalty may be incurred. To detect the risk of *overtime fault* at run-time, we should check the likelihood that the running instance does not exceed the MCT based on the amount of time T_c expired to that point. Let us consider T_e as the remaining cycle time, i.e. the amount of time estimated to complete the current instance given T_c based on past executions, which can be computed using the approach in [2]. Then the probability of exceeding MCT can be computed as $1 - MCT / (T_e + T_c)$ if $T_e + T_c > MCT$ and is equal to 0 if $T_e + T_c \leq MCT$. If this probability is greater than a tolerance value (e.g. 60%), we notify the risk to the user.

A second fault is related to the resources participating in the process. The Senior Finance Officer who has approved a Shipment Payment Order for a given customer must

have not approved another order by the same customer in the last d days, otherwise there is a potential for *approval fraud*, a violation of a four-eyes principle across different instances of the Payment subprocess. To detect this risk we first have to check that there is an order, say order o of customer c , to be approved. Moreover, we need to check that either of the following conditions holds: i) o has been allocated to a Senior Finance Officer who has already approved another order for customer c in the last d days; or ii) at least one Senior Finance Officer is available who approved an order for customer c in the last d days and all other Senior Finance Officers who did not approve an order for c during the last d days are unavailable.

Finally, a third fault relates to a situation where a process instance executes a given task too many times, typically via a loop. Not only could this lead to a process slow-down, but also to a “livelock” if the task is in a loop whose exit condition is deliberately never met. In general, given a task t , a maximum number of allowable executions of t per process instance, $MAE(t)$, can be fixed as part of the service-level agreement for t . In our example, this fault may occur if task “Update Shipment Payment Order” is re-executed five times within the same process instance. We call this an *order unfulfillment* fault. To detect the risk at run-time, we need to check if: i) the Update task is currently being performed for order o ; and ii) it is likely that the task will be repeated within the same process instance. The probability that the number of times a task will be repeated within the same instance is computed by dividing the number of instances where the MAE for the task has been reached by the number of instances that have executed this task at least as many times as it has been executed by the current instance, and have completed. If the probability to exceed $MAE(t)$ is greater than a tolerance value for t , e.g. 60%, we notify the risk to the user.

In the next section we will show what mitigation actions can be performed that change a process instance on the fly in order to mitigate its risks.

3 Approach

In this paper we deal with the problem of automatically mitigating one or more business process risks for a specific running process instance (*case* for short), without raising other business process risks for the same case. This problem belongs to the family of multi-objective optimization problems, and we propose the use of simulated annealing for finding a Pareto-optimal solution, or a set of such solutions.

The Process Risk Simulated Annealing (PRSA) algorithm is an application of the DBMOSA [26] algorithm where at each iteration a new solution is discovered through the use of one or more random mitigation actions. The algorithm proposes a solution, or mitigation, as a sequence of elementary mitigation actions. A behavioral cost is associated with each action, that measures its impact. The total cost of a solution is the sum of the costs of each mitigation action used. A good solution to the PRSA algorithm is a solution that reduces the likelihood of a risk under its threshold, keeping the total cost as low as possible.

When comparing solutions that have the same cost, a solution that fully mitigates a risk is better than one that mitigates that risk because its risk condition is no longer evaluable. And in turn, this solution is better than one that does not mitigate the risk at all. Finally, if two solutions mitigate the same risk, we privilege the one that yields

the lowest risk probability. Given two solutions a, b we say that a dominates b if it mitigates the same risks mitigated by b with a lower total cost. As result, we define them as mutually non-dominating if neither one dominates the other.

Below we describe the more elementary mitigation actions that can be used to create a solution, and how they affect a process case. Before introducing them, we introduce a number of preliminary concepts and notations.

YAWL Specification. We will not repeat the full definition of a YAWL specification as defined in [18], we will only use selected parts. The set of net identifiers is given by $NetID$ and the process identifier is the net identifier of the root net, $ProcessID \in NetID$. Furthermore, each net has, among others, a set of conditions C , an input condition $i \in C$, an output condition $o \in C$, and a set of tasks T and there is a flow relation $F \subseteq (C \setminus \{o\} \times T) \cup (T \times C \setminus \{i\}) \cup (T \times T)$.

We use the following auxiliary functions from [18]. The pre-set of x is defined as $\bullet x = \{y \in C \cup T \mid (y, x) \in F\}$ and the post-set of x is defined as $x \bullet = \{y \in C \cup T \mid (x, y) \in F\}$. We also introduced other auxiliary functions. The set of tasks that directly or through a place precedes a task ts is referred to as the *task pre-set* of t and is defined as $ot = \{x \in T \mid x \in \bullet t \vee \exists y \in C [y \in \bullet t \wedge x \in \bullet y]\}$. Similarly, the *task post-set* of t is defined as $t\bullet = \{x \in T \mid x \in t \bullet \vee \exists y \in C [y \in t \bullet \wedge x \in y \bullet]\}$. Finally, to detect all the successors of a task, it is defined as $t\bullet^*$ and it is the transitive closure of $t\bullet$.

Following the convention in [18], we write e.g. T_n to access the tasks of net n . Moreover, for a YAWL specification y , T_y is the set of tasks that occur in any of its nets, i.e. $T_y = \cup_{n \in NetID} T_n$, and for a set of YAWL specifications Y , T_Y is the set of tasks that occur in any of the nets of any of the specifications, i.e. $T_Y = \cup_{y \in Y} T_y$.

In our context we have only one Organizational model [18] and what is relevant for us is the set of resources, $UserID$, to whom work items can be assigned. Finally, we defined the set of *skippable* tasks as $\{t \in T_Y \mid \exists r \in UserID [skip \in UserTaskPriv(r, t)]\}$.

The set $StatusType$ contains the various statuses that a work item may go through during its lifecycle. These are: *offered*, *allocated*, *started*, *completed*, *forceCompleted*, *cancelled*, *failed*, *deadlocked* used by the YAWL system and additionally *deoffered*, *deallocated*, *destarted*, *rollback*, *skipped* used for mitigation purposes. Many of these statuses are self-explanatory. The status *rollback* is the status of a work item which was completed but then enabled again though not *offered*. The status *skipped* is the status of a work item that was skipped, which is similar to the status *completed* but the work item was not actually performed. For convenience, we provide certain groupings of event types. In particular, $Rel \triangleq StatusType \setminus \{cancelled, failed, rollback\}$ is the set of event types that identify a work item as subject to mitigation. $Active \triangleq \{offered, allocated, started\}$ is the set of event types that mark a work item as in progress, $Completed \triangleq \{completed, forceCompleted\}$ is the set of event types that mark a work item as completed, and $ActiveC : Active \cup Completed$ is their union.

Given set $ActiveC$ we define a partial order $\subseteq_{\preceq} ActiveC \times ActiveC$ such that it preserves the partial ordering $deoffered < offered < allocated < started < completed = forceCompleted$.

Definition 1 (Log). In the context of a set of YAWL specifications Y , with associated set of tasks T_Y and a set of root nets \mathcal{R} , a log is defined as $L = (\mathcal{E}, \mathcal{W}, \mathcal{C}, Model, WI, Case, Task, EvType, Time, Res, Inp, Outp)$ where:

- \mathcal{E} is a set of events,
- \mathcal{W} is a set of work items,
- \mathcal{C} is a set of case identifiers,
- $Model : \mathcal{C} \rightarrow \mathcal{R}$ is a function relating cases to the root nets of the associated YAWL specification,
- $WI : \mathcal{E} \rightarrow \mathcal{W}$ is a surjective function relating events to work items,
- $Case : \mathcal{E} \rightarrow \mathcal{C}$ is a surjective function relating events to cases,
- $Task : \mathcal{W} \rightarrow T_Y$ is a function relating work items to tasks,
- $EvType : \mathcal{E} \rightarrow StatusType$ is a function relating events to work item statuses,
- $Time : \mathcal{E} \rightarrow \mathcal{T}$ is an injective function relating events to timestamps, hence no two events in the log can have identical timestamps,
- $Res : \mathcal{E} \rightarrow 2^{UserID}$ is a function relating events to sets of resources, as some events may concern multiple resources (e.g. a work item being offered),
- $Inp : \mathcal{E} \times Var \rightarrow \Omega$ is a partial function relating events and variables to (input) values,
- $Outp : \mathcal{E} \times Var \rightarrow \Omega$ is a partial function relating events and variables to (output) values.

Definition 2 (Event Comparison). Let L be a log, given $\mathcal{E}' \subseteq \mathcal{E}$, $\mathcal{E}' \neq \emptyset$, we define the operators $e_1 < e_2$ iff $Time(e_1) < Time(e_2)$ and $e_1 \leq e_2$ iff $Time(e_1) \leq Time(e_2)$, which reflect the temporal ordering on events, and the operators $\min \mathcal{E}' = e_1$ iff $e_1 \in \mathcal{E}'$ and for all $e_2 \in \mathcal{E}'$, $e_1 \leq e_2$, which determines the earliest event of an event set, and $\max \mathcal{E}' = e_1$ iff $e_1 \in \mathcal{E}'$ and for all $e_2 \in \mathcal{E}'$, $e_2 \leq e_1$, which determines the latest event of an event set.

Useful is the possibility of identifying events belonging to the same work item.

Definition 3 (Work Item Event Grouping). Let L be a log, e an event in this log, $e \in \mathcal{E}$, and w a work item in this log, $w \in \mathcal{W}$, we define the set of events that belong to work item w as $\overline{WI}(w) \triangleq \{e \in \mathcal{E} \mid WI(e) = w\}$. Similarly, we define the set of events that belong to the same work item of e as $\overline{WI}(e) \triangleq \overline{WI}(WI(e))$. Finally, the latest event for work item w is defined as $\omega_w \triangleq \max \overline{WI}(w)$.

As for events we are interested in being able to compare work items.

Definition 4 (Work Item Comparison). Let L be a log, with $w_1, w_2 \in \mathcal{W}$, we define $w_1 < w_2$ as $\max \overline{WI}(w_1) < \min \overline{WI}(w_2)$. This operator reflects the partial temporal order between work items, i.e. work item w_1 precedes work item w_2 if its latest event is earlier than the earliest event of w_2 .

An execution graph for a process case provides a view of its execution and is defined on the basis of a log and its corresponding process model.

Definition 5 (Execution Graph). Let L be a process log with case c , Y its YAWL specification, and $UserID$ the set of resources, we define the execution graph of c as $G(c) = (Node, NodeTask, Status, \rightsquigarrow, NodeRes, TimeNode, VarNode)$ where:

- $Node = \{w \in \mathcal{W} \mid EvType(\omega_w) \in Rel \wedge Case(w) = c\}$ is the set of nodes, where each node represents a work item that is not modifiable,
- $NodeTask = Task|_{Node}$ is the restriction of the function $Task$ to the set of nodes,
- $Status = \{(\omega_w, s) \in Node \times Rel \mid s = EvType(\omega_w)\}$ is a function relating a node with its status of execution,

- $\rightsquigarrow = \{(w_1, w_2) \in \text{Node} \times \text{Node} \mid \text{Status}(w_1) \in \{\text{completed}, \text{skip}\} \wedge \text{NodeTask}(w_1) \in \circ \text{NodeTask}(w_2) \wedge \nexists w_3 \in \text{Node}[(\text{NodeTask}(w_1) = \text{NodeTask}(w_3) \vee \text{NodeTask}(w_2) = \text{NodeTask}(w_3)) \wedge w_1 < w_3 \wedge w_3 < w_2]\}$ is the flow relation between work items. Its reflexive transitive closure is defined as \rightsquigarrow^* ,
- $\text{NodeRes} = \{((w, s), r) \in (\text{Node} \times \text{Active}) \times 2^{\text{UserID}} \mid \exists e_1 \in \overline{\text{WI}}(w)[\text{EvType}(e_1) = s \wedge r = \text{Res}(e_1) \wedge \nexists e_2 \in \overline{\text{WI}}(w)[e_1 < e_2 \wedge \text{EvType}(e_2) \preceq s]]\}$ is a function that yields the resources that are involved in the latest changing w to status s ,
- $\text{TimeNode} = \{((w, s), t) \in (\text{Node} \times \text{ActiveC}) \times \mathcal{T} \mid \exists e_1 \in \overline{\text{WI}}(w)[\text{EvType}(e_1) = s \wedge t = \text{Time}(e_1) \wedge \nexists e_2 \in \overline{\text{WI}}(w)[e_1 < e_2 \wedge \text{EvType}(e_2) \preceq s]]\}$ is a partial function that yields the timestamp when w latest moved to status s ,
- $\text{VarNode} = \{((w, x), v) \in (\text{Node} \times \text{Var}) \times \Omega \mid \text{EvType}(\omega_w) \notin \{\text{skip}, \text{deoffered}\} \wedge v = \text{Inp}(\max \{e_2 \in \overline{\text{WI}}(w) \mid \text{EvType}(e_2) = \text{offered}\}, x) \oplus \{((w, x), v) \in (\text{Node} \times \text{Var}) \times \Omega \mid \text{EvType}(\omega_w) \in \text{Completed} \wedge v = \text{Outp}(\omega_w, x)\}\}$ is a partial function relating nodes and variables to values.

As we explore mitigation options the execution graph should evolve along with it, and the initial execution graph becomes a dynamic data structure from which we can modify nodes. We will refer to this modified execution graph as *mitigation graph*.

The concept of *border* identifies work items that can be modified. Such work items are currently in execution, or they are completed work items for which there are no successor work items that are completed or being executed.

Definition 6 (Border). Let G be a mitigation graph. We define the border of G , \circ_G , as $\{n_1 \in \text{Node} \mid \forall n_2 \in \text{Node}[n_1 \rightsquigarrow^* n_2 \Rightarrow \text{Status}(n_2) \in \{\text{deoffered}, \text{skipped}, \text{rollback}\}]\}$.

Discovering which work items are related to a specific resource can be relevant. For this reason we introduce the function \otimes which takes as argument an active event type and a resource, and yields the active border work items associated with that resource.

Definition 7 (Resource Involvement). Let L be a log, r a resource, $r \in \text{UserID}$, and s an active event type, $s \in \text{Active}$, then $\otimes(r, s) \triangleq \{w \in \mathcal{W} \mid \exists e \in \mathcal{E}[e = \omega_w \wedge \text{EvType}(e) = s \wedge r \in \text{Res}(e)]\}$ We will abbreviate active event types to their first letter, e.g. we will write $\otimes(r, o)$ instead of $\otimes(r, \text{offered})$.

To reassign a resource we need to know if the resource is associated with process cases which risks could be eventuate. To obtain an answer to this question we define the concept of *safe resource collection*. Note that in practice we need to know the process model and the resource model in order to obtain this collection.

Definition 8 (Safe Resource Collection). Let L be a log, UserID the set of resources, and \mathcal{C} the set of cases, the safe resource collection is defined as $\text{SRC} = (\text{PI}, \text{Risk})$ where:

- $\text{PI} : \text{UserID} \rightarrow 2^{\mathcal{C}}$ is a function relating resources to sets of active cases, in particular given a resource r returns the set of active cases for which there is a task associated to a role, which resource r belongs to,
- $\text{Risk} : \mathcal{C} \rightarrow \{\text{risky}, \text{norisky}\}$ is a function that tells us whether a certain case is considered risky or not.

Definition 9 (Mitigations). Let Y be a set of YAWL specifications, with associated set of tasks T_Y , a set of resources UserID , and a log L . A mitigation is represented as $M = (\mathcal{A}, \text{AcType}, \text{AcTask}, \text{AcRes}, \text{AcCase}, \succ)$ where:

- \mathcal{A} is a set of mitigation actions,

- $AcType : \mathcal{A} \rightarrow \{deoffer, deallocate, destart, offer, allocate, start, rollback, skip\}$, is a function relating actions to types of mitigation,
- $AcTask : \mathcal{A} \rightarrow T_Y$ is a function relating actions to tasks,
- $AcRes : \mathcal{A} \mapsto UserID$ is a partial function relating actions to resources,
- $AcCase : \mathcal{A} \rightarrow \mathcal{C}$ is a function relating actions to cases,
- $\succ \subseteq \mathcal{A} \times \mathcal{A}$ is a total ordering on mitigation actions indicating the order in which they need to be performed. We refer to this total ordering as the mitigation sequence.

The insertion of a new mitigation action $a \notin \mathcal{A}$ into mitigation M , can be expressed as $addMit(M, a, et, t, r, c) \triangleq (\mathcal{A} \cup \{a\}, AcType \cup \{(a, et)\}, AcTask \cup \{(a, t)\}, AcRes \cup \{(a, r)\}, AcCase \cup \{(a, c)\}, \succ \cup \{(x, a) \mid x \in \mathcal{A}\})$.

Now we are in a position to introduce the mitigation actions. For each action we will provide a short description of its behavior; we will quantify its cost and specify the precondition(s) required for its application. All these actions are executed in the context of a mitigation M . As soon as a risk is detected we collect the log L containing all process cases. This log is used to generate an execution graph G' , that we refer to as the original execution graph. It is used as a reference for comparison with the original status of the system. The effects of mitigations actions are explored, though not yet applied, during execution of the mitigation algorithm, and hence they are performed on a clone of the original execution graph which we will refer to as G .

Throughout the remainder of this section G' is the original execution graph, G the mitigation graph in use, and $c \in \mathcal{C}$ is a case. Moreover, whenever a node is modified, we need to store the time this modification occurred. In order to capture the time, we use function $curr()$.

A mitigation is a sequence of mitigation actions. Below we describe the mitigation actions supported by the PRSA algorithm.

Deoffer This action deoffers a task from a resource to whom the task was offered. We can execute $deOff(c, G, M)$ as described in Algorithm 1 if there is a work item $x \in \cup_G$ such that x is an *offered* work item. The cost of this action was set to one and this action serves as a reference for the cost of the other actions.

Algorithm 1: Deoffer Task

```

function deOff(Case c, Mitigation Graph G, Mitigation M);
Output: Mitigation Graph G, Mitigation M
begin
  n ← Any({x ∈ ∪G | Status(x) = offered});
  if n ≠ ⊥ then
    r ← Any(NodeRes(n, offered));
    if |NodeRes(n, offered)| > 1 then
      et ← offered;
      TimeNode ← TimeNode ⊕ {(n, offered), curr()};
      NodeRes ← NodeRes ⊕ {(n, offered), NodeRes(n, offered) \ {r}};
    else
      et ← deoffered;
      TimeNode ← {(n, offered)} ⋖ TimeNode;
      NodeRes ← NodeRes(n, offered) ⋖ NodeRes;
      VarNode ← {(n, v) | VarNode(n, v) ∈ Ω} ⋖ TimeNode;
    Status ← Status ⊕ {(n, et)};
    M ← addMit(M, NewAction(), deoffer, NodeTask(n), r, c);
  return (G, M)

```

Deallocate This action deallocates a task from the resource to whom the task was allocated. If there is a work item $x \in \cup_G$ such that x is an *allocated* work item,

we can execute $deAll(c, G, M)$ as described in Algorithm 2. We set the cost of this action to two, since considering the progress status of a work item, deallocating a work item should be more “expensive” than deoffering it. In the Payment subprocess this action could be used to mitigate the approval fraud risk. The work item of “Approve Shipment Payment Order” can be deallocated from the resource to whom this work item is allocated when the risk is detected, since this resource approved another order for the same customer in the past.

Algorithm 2: Deallocate Task

```

function  $deAll(\text{Case } c, \text{Mitigation Graph } G, \text{Mitigation } M)$ ;
Output: Mitigation Graph  $G$ , Mitigation  $M$ 
begin
   $n \leftarrow \text{Any}(\{x \in \mathcal{O}_G \mid \text{Status}(x) = \text{allocated}\})$ ;
  if  $n \neq \perp$  then
     $r \leftarrow \text{Any}(\text{NodeRes}(n, \text{allocated}))$ ;
     $\text{Status} \leftarrow \text{Status} \oplus \{(n, \text{offered})\}$ ;
     $\text{NodeRes} \leftarrow \{(n, \text{allocated})\} \triangleleft \text{NodeRes}$ ;
     $\text{TimeNode} \leftarrow \{(n, \text{allocated})\} \triangleleft \text{TimeNode}$ ;
     $M \leftarrow \text{addMit}(M, \text{NewAction}(), \text{deallocate}, \text{NodeTask}(n), r, c)$ ;
  return  $(G, M)$ 

```

Destart This action brings an already started work item back to the state *allocated* and allocates it to the resource who started it. We can execute $deSta(c, G, M)$, as described in Algorithm 3, if there is a work item $x \in \mathcal{O}_G$ such that x is a *started* work item. For this action we set the cost to three as destarting a work item requires more effort than deallocating a work item. The *destart* action may be used to mitigate the same risk as the *deallocate* action. This action may be used to destart a resource who never approved an order for the current customer from another work item, reducing this way the probability of allocating the work item to a resource who approved the Shipment Payment Order in the past.

Algorithm 3: Destart Task

```

function  $deSta(\text{Case } c, \text{Mitigation Graph } G, \text{Mitigation } M)$ ;
Output: Mitigation Graph  $G$ , Mitigation  $M$ 
begin
   $n \leftarrow \text{Any}(\{x \in \mathcal{O}_G \mid \text{Status}(x) = \text{started}\})$ ;
  if  $n \neq \perp$  then
     $r \leftarrow \text{Any}(\text{NodeRes}(n, \text{started}))$ ;
     $\text{Status} \leftarrow \text{Status} \oplus \{(n, \text{allocated})\}$ ;
     $\text{NodeRes} \leftarrow \{(n, \text{started})\} \triangleleft \text{NodeRes}$ ;
     $\text{TimeNode} \leftarrow \{(n, \text{started})\} \triangleleft \text{TimeNode}$ ;
     $M \leftarrow \text{addMit}(M, \text{NewAction}(), \text{destart}, \text{NodeTask}(n), r, c)$ ;
  return  $(G, M)$ 

```

Offer This action offers a work item to a resource to whom the task is not currently offered, either because it is not yet part of the set of resources to whom the task is currently offered, or because the task is currently *deoffered*. Given a function D that relates tasks to the set of resources to whom their work items can be offered, we can execute $off(D, c, G, G', M)$, as described in Algorithm 4, if there is a work item $x \in \mathcal{O}_G$ such that x is an *offered* or *deoffered* work item, and this work item is an *offered*, *allocated* or *started* work item in the execution graph G' . This action has a cost of one, the same as the *offer* action.

Algorithm 4: Offer Task

```
function off(Distribution Set  $D$ , Case  $c$ , Mitigation Graph  $G$ , Execution Graph  $G'$ , Mitigation  $M$ );
Output: Mitigation Graph  $G$ , Mitigation  $M$ 
begin
   $n \leftarrow \text{Any}(\{x \in \mathcal{O}_G \mid \text{Status}_G(x) \in \{\text{offered}, \text{deoffered}\} \wedge \text{Status}_{G'}(x) \in \text{Active}\});$ 
   $r \leftarrow \perp;$ 
  if  $n \neq \perp$  then
     $r \leftarrow \text{Any}(D(\text{NodeTask}(n)) \setminus \text{NodeRes}_G(n, \text{offered}) \cup \text{NodeRes}_{G'}(n, \text{offered}));$ 
  if  $r \neq \perp$  then
    if  $\text{Status}_G = \text{deoffered}$  then
       $\text{Status}_G \leftarrow \text{Status}_G \oplus \{(n, \text{offered})\};$ 
       $\text{VarNode}_G \leftarrow \text{TimeNode}_G \cup \{\text{VarNode}_{G'}(n)\};$ 
       $\text{TimeNode}_G \leftarrow \text{TimeNode}_G \cup \{(n, \text{offered}), \text{curr}()\};$ 
       $\text{NodeRes}_G \leftarrow \text{NodeRes}_G \cup \{(n, \text{offered}), \text{NodeRes}_G(n, \text{offered}) \cup \{r\}\};$ 
    else
       $\text{TimeNode}_G \leftarrow \text{TimeNode}_G \oplus \{(n, \text{offered}), \text{curr}()\};$ 
       $\text{NodeRes}_G \leftarrow \text{NodeRes}_G \oplus \{(n, \text{offered}), \text{NodeRes}_G(n, \text{offered}) \cup \{r\}\};$ 
     $M \leftarrow \text{addMit}(M, \text{NewAction}(), \text{offer}, \text{NodeTask}_G(n), r, c);$ 
  return ( $G, M$ )
```

Allocate This action reallocates a work item that was deallocated before (and still has not been allocated) to a resource to whom the task was not allocated when the deallocation took place. We can execute $\text{all}(c, G, G', M)$, as described in Algorithm 5, if there is a work item $x \in \mathcal{O}_G$ such that x is an offered work item, and x is originally an *allocated* or *started* work item. This action has a cost of minus one. The reason behind this is quite simple. Since this action can only be executed if we previously executed a deallocation, these two actions can be seen as a unique action that change the resource involved with the work item. For example this action can be used once we deallocated a resource from a work item of the “Approve Shipment Payment Order” task. Then this action allocates the same work item to another resource who may not have approved an order for the same customer in the past.

Algorithm 5: Allocate Task

```
function all(Case  $c$ , Mitigation Graph  $G$ , Execution Graph  $G'$ , Mitigation  $M$ );
Output: Mitigation Graph  $G$ , Mitigation  $M$ 
begin
   $n \leftarrow \text{Any}(\{x \in \mathcal{O}_G \mid \text{Status}_G(x) = \text{offered} \wedge \text{Status}_{G'}(x) \in \{\text{allocated}, \text{started}\}\});$ 
   $r \leftarrow \perp;$ 
  if  $n \neq \perp$  then
     $r \leftarrow \text{Any}(\text{NodeRes}_G(n, \text{offered}) \setminus \{\text{NodeRes}_{G'}(n, \text{allocated})\});$ 
  if  $r \neq \perp$  then
     $\text{Status}_G \leftarrow \text{Status}_G \oplus \{(n, \text{allocated})\};$ 
     $\text{NodeRes}_G \leftarrow \text{NodeRes}_G \cup \{(n, \text{allocated}), \{r\}\};$ 
     $\text{TimeNode}_G \leftarrow \text{TimeNode}_G \cup \{(n, \text{allocated}), \text{curr}()\};$ 
     $M \leftarrow \text{addMit}(M, \text{NewAction}(), \text{allocate}, \text{NodeTask}_G(n), r, c);$ 
  return ( $G, M$ )
```

Start This action restarts a work item that was previously destarted (and has not yet been restarted) and associates it with a different resource from the one who started the task. We can execute $\text{sta}(c, G, G', M)$, as described in Algorithm 6, if there is a work item $x \in \mathcal{O}_G$ such that x is an *allocated* work item, and x is originally a *started* work item. The cost of this action is minus one, and the reasoning is similar to that used for the *allocate* action.

Algorithm 6: Start Task

```
function sta(Case  $c$ , Mitigation Graph  $G$ , Execution Graph  $G'$ , Mitigation  $M$ );
Output: Mitigation Graph  $G$ , Mitigation  $M$ 
begin
   $n \leftarrow \text{Any}(\{x \in \mathcal{O}_G \mid \text{Status}_G(x) = \text{allocated} \wedge \text{Status}_{G'}(x) = \text{started}\});$ 
   $r \leftarrow \perp$ ;
  if  $n \neq \perp$  then
     $r \leftarrow \{NodeRes_G(n, \text{allocated})\} \setminus \{NodeRes_{G'}(n, \text{started})\};$ 
  if  $r \neq \perp$  then
     $Status_G \leftarrow Status_G \oplus \{(n, \text{started})\};$ 
     $NodeRes_G \leftarrow NodeRes_G \cup \{(n, \text{started}), \{r\}\};$ 
     $TimeNode_G \leftarrow TimeNode_G \cup \{(n, \text{started}), \text{curr}()\};$ 
     $M \leftarrow \text{addMit}(M, \text{NewAction}(), \text{start}, \text{NodeTask}_G(n), r, c);$ 
  return ( $G, M$ )
```

Rollback This action returns a completed work item to the status of unoffered. We can execute $\text{rollbackTask}(c, G)$ if there is a work item $x \in \mathcal{O}_G$ such that x is a *completed* work item. Its operationalization is described in Algorithm 7. The rollback action restores the case to a consistent status where the execution of a given work item never happened. A compensation routine can be associated with a task, so that it is triggered when the task is rolled back. The idea of this compensation routine is to deal with elements outside the control of the workflow engine (e.g. returning the money to a client after their payment has been rolled back). The rollback action is our most powerful action and has a cost of nine, obtained by adding the absolute values of all the actions introduced until now. In our Payment subprocess, we could use this action when we execute a large number of updates on the same Payment Order.

Algorithm 7: Rollback Task

```
function rollbackTask(Case  $c$ , Execution Graph  $G$ , Mitigation  $M$ );
Output: Execution Graph  $G$ , Mitigation  $M$ 
begin
   $n_1 \leftarrow \text{Any}(\{x \in \mathcal{O}_G \mid \text{Status}_G(x) \in \text{Completed}\});$ 
  if  $n_1 \neq \perp$  then
    foreach  $n_2 \in Node_G$  do
      if  $n_1 \rightsquigarrow_G^* n_2$  then  $Status_G \leftarrow Status_G \oplus \{(n_2, \text{rollback})\};$ 
     $Status_G \leftarrow Status_G \oplus \{(n_1, \text{rollback})\};$ 
     $M \leftarrow \text{addMit}(M, \text{NewAction}(), \text{rollback}, \text{NodeTask}_G(n_1), \perp, c);$ 
  return ( $G, M$ )
```

Skip This action marks an unoffered and skippable task as ‘to be skipped’. If there exists a task $t \in \text{skippable}$ which does not have any work item active or completed, and there not exists a mitigation action $a \in \mathcal{A}$ which skipped task t for case c , then we define $\text{skipTask}(c, G)$ as described in Algorithm 8. To limit the use of this action, since this action may produce inconsistency in the data, we decided to assign a cost of nine. The utility of this action can be seen in two situations when we consider our running example. The first situation is the order unfulfillment. In this case, to prevent the reiterated execution of an update, we may decide to skip the “Update Shipment Payment Order” task. The second situation is the overtime process risk. In this case we may decide to skip some tasks in order to complete the process in time.

Algorithm 8: Skip Task

```
function skip(Case c, Mitigation Graph G, Mitigation M, YAWL Specification Y);
Output: Mitigation Graph G, Mitigation M
begin
   $D \leftarrow \{t \in T_{Model(c)} \mid \exists x \in \circ_G [Status(x) = deoffered \wedge NodeTask(x) = t]\};$ 
   $R \leftarrow \{t \in T_{Model(c)} \mid \exists x \in \circ_G [Status(x) = rollback \wedge NodeTask(x) = t]\};$ 
   $U \leftarrow \{t \in T_{Model(c)} \mid \exists x \in \circ_G [t \in NodeTask(x) \circ^*]\};$ 
   $t \leftarrow Any((D \cup R) \cap U);$ 
   $a \leftarrow \exists a \in \mathcal{A} \mid AcType(a) = skip \wedge AcTask(a) = t \wedge AcCase(a) = c;$ 
  if  $t \neq \perp \wedge skippable(t) = yes \wedge \neg a$  then
     $M \leftarrow addMit(M, NewAction(), skip, t, \perp, c);$ 
  return (G, M)
```

Relocate Resource This action looks for a resource that is only involved in the execution of a work item belonging to a case for which no risk was defined. If once such a resource is found, it deallocates (and destarts if necessary) the work item associated with this resource and allocates the resource to a work item of the process case that we want to mitigate. The cost of this action is seven since this action performs a (partial) sequence of destart and deallocate on two work items, and another allocate and a start action on one work item. Let x be an active border work item $x \in \circ_G$ in case c , r be a resource involved only in case c_2 , and c_2 be process which is not risky. If resource r only started or allocated one work item (of any active border events), then we can execute $relRes(c, G, M, SRC)$ as described in Algorithm 9.

Algorithm 9: Relocate Resource

```
function relRes(Case c, Mitigation Graph G, Mitigation M, Safe Resource Collection SRC);
Output: Mitigation Graph G, Mitigation M
begin
   $n \leftarrow Any(\{x \in \circ_G \mid Status(x) = started\});$ 
   $r_1 \leftarrow NodeRes(n, started);$ 
   $t \leftarrow \perp;$ 
  foreach  $r_2 \in NodeRes(n, offered)$  do
    if  $t = \perp \wedge |PI(r)| = 1 \wedge Risk(Any(PI(r))) = norisky \wedge (|\otimes(r, s)| + |\otimes(r, a)|) = 1$  then
      if  $|\otimes(r, s)| = 1$  then
         $c_2 \leftarrow Case(Any(\otimes(r_2, s)));$ 
         $t \leftarrow Task(Any(\otimes(r_2, s)));$ 
         $M \leftarrow addMit(M, NewAction(), destart, t, r_2, c_2);$ 
      else
         $c_2 \leftarrow Case(Any(\otimes(r_2, a)));$ 
         $t \leftarrow Task(Any(\otimes(r_2, a)));$ 
       $M \leftarrow addMit(M, NewAction(), deallocate, t, r_2, c_2);$ 
       $NodeRes \leftarrow NodeRes \oplus \{(n, allocated), \{r\}\};$ 
       $NodeRes \leftarrow NodeRes \oplus \{(n, started), \{r\}\};$ 
       $TimeNode \leftarrow TimeNode \oplus \{(n, allocated), curr()\};$ 
       $TimeNode \leftarrow TimeNode \oplus \{(n, started), curr()\};$ 
       $M \leftarrow addMit(M, NewAction(), destart, NodeTask(n), r_1, c);$ 
       $M \leftarrow addMit(M, NewAction(), deallocate, NodeTask(n), r_1, c);$ 
       $M \leftarrow addMit(M, NewAction(), allocate, NodeTask(n), r, c);$ 
       $M \leftarrow addMit(M, NewAction(), start, NodeTask(n), r, c);$ 
  return (G, M)
```

4 Evaluation

We implemented the PRSA algorithm as a custom service in the YAWL system.⁴ We extended the YAWL system as it is built on a service-oriented architecture, which facilitates the addition of new services; it is open-source, which facilitates its distribution among academics and practitioners; and as the underlying YAWL language provides comprehensive supports for the workflow patterns [18].

The risk mitigation service interacts with the *risk detection service*⁵ that we developed previously [11], for the sake of identifying risks and computing their probabilities. It uses as input a reference to the process instance whose risks need to be mitigated, the complete YAWL specification for this instance, a log of the process (as extracted from the YAWL system), and a copy of the risk sensors associated with the process instance, as provided by the risk detection service. Modifications that a mitigation may introduce are communicated to the risk detection service, which recomputes the risk probabilities. The final solutions are returned to the user as recommendations. The one chosen by the user is then applied to the process instance under exam using the APIs provided by the YAWL engine. We implemented compensation actions associated with rolled back work items via the YAWL Worklet mechanism [18]. Accordingly, we equipped the YAWL Editor with an interface to allow users to associate a Worklet containing a compensation action to a task. When an instance of this task is rolled back, the associated Worklet is run as a separate process instance in the YAWL engine, so that from an engine perspective, the Worklet and its invoking processes are two distinct cases.

To prove the feasibility of our approach, we ran three experiments. First, we tested the required time to mitigate the same set of risks on different process models. Second, we checked the dependency of the mitigation time on different variables. Third, we checked the quality of the mitigations proposed on a specific process model.

For the first experiment, we selected four process models from real-life (see appendix A), with a minimum of five tasks and a maximum of twenty tasks. The first model (Process A) describes a film production process, carried out on a daily basis. This process is taken from a case study we conducted in collaboration with the Australian Film, Television and Radio School [24]. The other three models are subprocesses of an order fulfillment process inspired by the VICS industry standard for logistics [28]. The first of these processes (Process B) deals with the ordering, the second (Process C) deals with the payment for the goods and is the process we showed in Section 2, the third process (Process D) deals with the delivery of the goods.

Next, we defined seven generic risk conditions on these processes (see appendix B). These conditions represent possible undesirable situations that may arise in a process, and relate to different process aspects such as data, resources and control-flow elements. They are domain-independent so that we could define them on all four process models. The first condition detects a situation where two concurrent work items may not complete in a desired order. The second one is used to detect a violation of the four-eyes principle between parallel work items. The third one detects whether a time limit is exceeded when executing a loop. The fourth condition detects a possible delay

⁴ <http://www.yawlfoundation.org>

⁵ <http://www.yawlfoundation.org/prsa>

with the execution of a work item. The fifth one detects the possibility that two concurrent work items that should be executed by the same resource are actually allocated to two different resources (a situation that is not possible to enforce with many workflow management systems). The sixth one detects a delay with the execution of a portion of the process while the seventh one detects a data error, specifically if the data values produced by two concurrent work items are not the same.

For each process model we generated a variant with a specific combination of the above risk conditions. This led to a total of 180 process models (not every risk identified could be applied to every process model), as follows: 19 models with one risk condition, 40 models with two risk conditions, 50 with three risk conditions, 41 models with four risk conditions, 22 process models with five risk conditions, seven process models with six risk conditions, and one process model with all seven risk conditions.

Process	Size	Variants	Risks avg/max	Mitigation time [sec]			Candidates		
				min	max	avg	min	max	avg
Process A	20	127	3.53 / 7	0.003	178.891	26.415	2	20,181	3,456
Process B	5	7	1.71 / 3	0.001	0.033	0.015	3	54	32
Process C	15	31	3.05 / 5	0.001	0.117	0.030	2	256	60.93
Process D	5	15	2.13 / 4	0.004	0.929	0.170	2	553	78.2
Total	45	180	3.18 / 7	0.001	178.891	18.657	2	20,181	2,457

For each process model we ran ten tests and averaged the results. Each test was executed on the first state of a process instance where all the risk conditions evaluated to true. For each group of tests on the same process model we measured the time required to obtain the first solution that mitigates all risks, and the number of candidate solutions generated by the algorithm in order to obtain this solution. We performed the tests on an Intel Core I5 M560 2.67GHz processor with 4GB RAM, running Linux Ubuntu v11.10.

Table 1. Time and number of candidate solutions explored to find the first solution.

Table 1 shows the results of this experiment. The second, third and fourth columns show the size (as number of tasks), the number of variants and the number of risk conditions for each of the four process models. The fifth and sixth columns show the mitigation time required to find the first solution, and the number of candidate solutions explored to find such a solution. From this table we can observe that the algorithm takes at most 3 mins (179 secs) to mitigate multiple risks in a variant of Process A (this timing refers to a combination of 5 risks for this process), though the average time is much lower (19 secs across all models). It seems reasonable to assume that in most business scenarios mitigation times in the order of a few minutes are acceptable, compared to the average time required to perform a task, and thus the average duration of a process instance. For example, let us assume an average duration of 24 hours for the Payment subprocess, with a new task being executed every 30 mins. Let us also assume that we sample the risk conditions every 5 mins. This means we have up to 6 mins to mitigate all identified risks before a new task is executed which may change the risk conditions.

Table 1 also shows that the algorithm needs to explore a very large number of candidate solutions to find the first solution (2,456 solutions on average across all models). While it is not fair to compare the computation power of a machine to that of humans, this result highlights the complexity of finding a solution. It is reasonable to think that

many of these candidate solutions explored by the algorithm would also need be evaluated by a human in order to find the right solution.

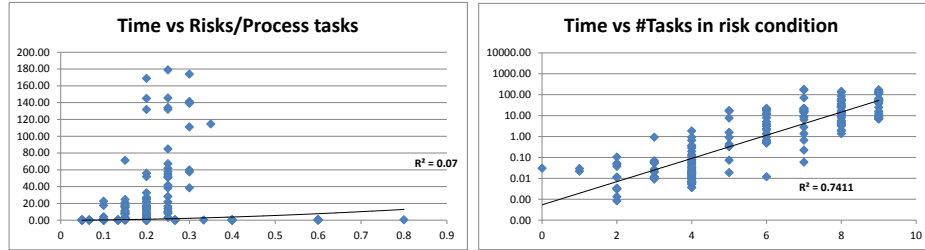


Fig. 3. Correlation between time and a) risks/tasks ratio, b) tasks in risk conditions.

In the second experiment, we investigated the factors affecting the performance of the algorithm. One would think that the mitigation time is proportional to the number of risks defined in a process model, and to the model size itself. The larger the number of risks and/or the model size, the longer it should take to mitigate such risks. However the data we extrapolated from Table 1 does not confirm this hypothesis. For example, the 21 variants of Process A with 5 risks have mitigation times ranging from 3.3 to 179 secs, despite their sizes and number of risks being the same. To verify that the mitigation time is not sensitive to the number of risks, nor to the process size, we plotted the correlation between the mitigation time and the ratio risks/process size in Figure 3a (the solid line is the linear regression of the points). The low value of the coefficient of determination R^2 (0.07) confirms this intuition. We then checked the correlation between the mitigation time and the number of tasks used in risk conditions. The intuition is that the more work items of these tasks are pending in a given state of the process instance, the larger the number of possible mitigation actions. The corresponding scatter plot is shown in Figure 3b, which indeed confirms this intuition ($R^2 = 0.74$).

Finally, we checked the feasibility of the solutions proposed by the algorithm, when mitigating the domain-specific risks associated with the Payment subprocess (cf. Section 2). We recall that two of these risks (overtime process and order unfulfillment) are detected when the associated probability, obtained by analyzing historical data, exceeds a tolerance threshold, whereas the third risk (approval fraud) involves a complex risk condition. We considered the first state of an instance of the Payment subprocess when all three risks are active. This occurs after executing “Update shipment payment order” for the third time, once task “Approve shipment payment order” has been allocated to a resource who has already executed this task in the past.

To obtain a small number of solutions, we stopped the algorithm after one min of execution. In this time-frame, five solutions were retrieved. For each solution, Table 2 reports whether the solution mitigates each of the three risks, and the cost of the solution in terms of mitigation actions performed on the initial process instance. In particular, a “-” indicates a risk not mit-

Solutions [at 1 min]	1	2	3	4	5
Overtime Process	+	+	+	+	+
Approval Fraud	+	+	+	+	+
Order Unfulfillment	+	+	±	±	-
Cost	50	50	40	40	19

Table 2. Payment subprocess mitigation.

igated, a “+” indicates a risk mitigated (with risk probability lower than the specific threshold if the condition depends on the risk probability), and a “±” indicates a risk mitigated whose condition cannot be computed for lack of information, i.e. some of the variables used in the risk condition are null. We recall that the algorithm prioritizes a solution whose risk is mitigated by computing the risk condition, than a solution whose risk is mitigated because the respective condition cannot be computed.

The five solutions identified are pairwise mutually non-dominating. Solutions 1 and 2 are dominated by solutions 3, 4 and 5 cost-wise, but dominate these solutions w.r.t. the mitigation of the order unfulfillment risk. Solution 5 dominates solutions 3 and 4 cost-wise but is dominated by these two solutions w.r.t. the mitigation of the order unfulfillment risk.

Let us briefly examine the mitigations performed by the five solutions (see appendix C). The first four solutions mitigate the approval fraud by deallocating the resource that was allocated “Approve shipment payment order”, while solution 5 additionally allocates the work item to a resource who did not execute this task for the same customer in the past. All these mitigations are feasible, though the one provided by solution 5 is more robust, since there is no risk that the task gets allocated to a resource who has already executed it. The order unfulfillment risk is mitigated by solutions 1 and 2 through rolling back the work item of task “Update shipment payment order” (which leads to a deoffer of the work item of task “Approve shipment payment order” that comes afterwards). Solutions 3 and 4 do this too but also mark this task ‘to be skipped’ preventing a possible re-execution of it. This action sets to null the risk variables associated with this task that retrieve the number of executions and its estimated remaining time making the risk mitigated but not computable. Thus, while all four solutions are feasible, we would prioritise the first two since these ensure that the risk probability has actually dropped below the threshold. Finally, all solutions differ in the way they mitigate the overtime process risk. Each of them skips a different task among those not yet executed (for simplicity, all of them have the same estimated duration). Despite the fact that all these solutions are feasible, only the mitigation proposed by solution 3 is interesting since it proposes to skip tasks “Update Shipment Payment Order” and “Approve Shipment Payment Order” avoiding this way that the loop is taken again. In other words, it prevents the order to undergo further updates, and subsequent approvals.

5 Related Work

Risk mitigation is an essential step in the risk management process [27]. Several risk analysis methods such as OCTAVE [4], CRAMM [7] and CORAS [21] describe risk mitigation guidelines. Although helpful, these guidelines are too generic and no support is offered on how they could be operationalized. Similarly, the academic literature recognizes the importance of mitigating process-related risks, though it focuses more on risk-aware BPM methodologies than on concrete algorithms for automating risk mitigation. For a comprehensive comparison of these methodologies, we refer to [11]. A well known example is the ROPE (Risk-Oriented Process Evaluation) methodology [16]. ROPE is concerned with threats to the resources required for process executions. If a required resource becomes unavailable, pre-planned countermeasures and recov-

ery procedures are manually enacted to handle the fault. These procedures are defined and validated via a simulator at design-time; enactment at runtime is designated to a ‘responsible person’, that is, the mitigation and recovery operations are not automated.

Various frameworks have been proposed for the dynamic adaptation of process instances. For example, ADEPT [12] supports adding, deleting and changing the sequence of tasks at both the model and instance levels, however such changes must be achieved via manual intervention by an administrator. AgentWork [23] provides the ability to modify process instances by dropping and adding individual tasks based on events and rules. CBRFlow [29] uses case-based reasoning to support runtime adaptation by allowing users to annotate rules during process execution. CEVICHE [17] is a service-based framework that uses the AO4BPEL (Aspect-Oriented for BPEL) language [9] to provide an option for skipping or reallocating tasks to other services in an ad-hoc manner. While these approaches could be used for risk mitigation purposes, they do not provide any help for the identification of which particular mitigation actions should be used. The YAWL Worklet Service [3] provides each task of a process instance with the ability to be associated with an extensible repertoire of actions (‘drop-in’ processes), one of which is contextually and dynamically bound to the task at runtime. It also supports capabilities for dynamically detecting and handling runtime exceptions, however the approach is generic and not specifically designed for risk detection and mitigation. Also a new situation cannot automatically be dealt with but requires a workflow administrator to intervene.

Our work is also related to *operational support* in process mining [1]. Operational support deals with the analysis of current and historical execution data, with the aim to predict future states of a running process instance, and provide recommendations to guide the user in selecting the next activity to execute based on certain objectives. For example, the approach for cycle time prediction in [2] could be, with the opportune modifications, adapted for risk prediction. Using this approach it would be possible to estimate the probability of an overtime risk and suggest the next steps the current instance should take in order to keep this risk under control. The application of this approach unfortunately requires that the process model captures all the possible mitigation actions as normal activities, i.e. as control-flow alternatives. For instance, if a task can be skipped, there should be a path without that task that leads to the end node of the process model. This may drastically increase the complexity of the process model. Moreover, this approach would not be applicable to capture mitigation operations on resources (i.e. deallocating a resource) or on task states (e.g. suspending a task). That said, more in general, our approach can be seen as a possible provider for operation support, and could thus be integrated in process mining environments like ProM.⁶

Our work provides recommendations to users as to which mitigation actions can be applied to the specific context at hand. As such, it shares commonalities with recommendation and decision support systems (DSS). Alter [5] states that the focus of such systems should be towards improving decision making within work systems, rather than externalizing support. This view is shared by our technique, which provides an extension for existing process-aware information systems, rather than a separate standalone tool. As such, it may be considered a member of the domain known as *Group Decision*

⁶ <http://processmining.org>

Support Systems, which facilitate task support in group environments. The conjunction of our work and DSS can be demonstrated via Alter's twenty-four work system principles [6], which may be used to analyze the capabilities of systems, using commonly understood terminology. In particular, our work meets principles #5 (encourage appropriate use of judgement), #6 (control problems at their source), #9 (match work practices with participants), #13 (provide information where it will affect action), #14 (protect information from inappropriate use), #22 (minimize unnecessary risks) and #24 (maintain the ability to adapt, change and grow). In addition, our work provides all four features associated with a DSS's ability to improve practice, as identified in a survey of seventy studies within the healthcare domain: (a) decision support provided automatically as part of workflow; (b) decision support delivered at the time and location of decision making; (c) actionable recommendations provided; and (d) computer based [20].

The mitigation operations we perform on resources share commonalities with task rescheduling systems, which become particularly important in those domains where appointment-based task executions are critical (e.g. healthcare). For example, a system is presented in [22] that supports the integration of unscheduled and scheduled tasks within a process instance by interfacing a workflow engine with individual user calendars maintained by MS Exchange Server. A similar, though more generic, approach is encapsulated in the YAWL Scheduling Service [25]. Another approach, by Eder et al. [14], focuses on the creation of personal schedules for each user so that a workflow engine, when making decisions about work assignment, can take into account information about a user's time constraints and availability to perform particular tasks.

In previous work [15] we explored the use of dominance-based MOSA for automatically fixing behavioral errors in process models, at design-time. Our work on risk mitigation can thus be seen as an adaptation of that idea to run-time aspects, since we aim to improve running process instances. Besides their distinct aims, the main difference between the two approaches is that for correcting behavioral errors we defined three objective functions capturing the structural and behavioral similarity of a solution to the incorrect model, whereas in risk mitigation the number and type of objective functions depends on which risks are active in a given state of a process instance.

6 Conclusion

This paper contributes a concrete technique for the automatic mitigation of process-related risks at run-time. The technique requires as input an executable process model and a set of associated risk conditions. At run-time, when one or more risk conditions evaluate to true, a process administrator can launch our technique to mitigate the identified risks and bring the process instance back to a safe state. This is achieved by generating a set of possible mitigations that change the current instance in order to bring the likelihood of the identified risks below a tolerance level. These mitigation actions are not performed directly on the instance under consideration. Rather, their effects are simulated and those solutions that mitigate the most risks in a given timeframe, are proposed as recommendations to the process administrator.

The mitigation actions are determined via a dominance-based MOSA algorithm. This choice allows us to explore the solution space as widely as possible, avoiding local optima. In essence, each risk is treated as an objective function whose likelihood needs to be minimized. The objective is reached as soon as the likelihood goes below the tolerance value for that particular risk. Mitigation actions affect various aspects of a process, such as task execution and resources utilization. To the best of our knowledge, this is the first time that process-related risks can be mitigated automatically.

The technique was implemented in the YAWL system and its performance evaluated with real-life process models. The tests show that on the analyzed process models a set of possible solutions can be found in a matter of seconds, or within a few minutes in the worst case, and that in all cases the associated risks are mitigated. We expect this technique to reduce the effort and time required by process administrators to understand what mitigation actions are feasible based on a particular state of the system. That said, we still need to validate the feasibility and appropriateness of the proposed mitigation actions with domain experts. We plan to do so by comparing the solutions obtained with our algorithm with those proposed by them. We also plan to improve the exploration of the solution space by prioritizing the mitigation of those risks that have the highest impact on the process objectives. In fact, currently all risks are treated alike whereas in reality this might not be the case. Finally, the algorithm could also be extended to prioritize certain mitigation actions based on how these have been ranked by the users in previously mitigated instances.

Acknowledgments We thank Peter Hughes for his valuable comments on the mitigation operations. This research is funded by the ARC Discovery Project “Risk-aware Business Process Management” (DP110100091).

References

1. W.M.P. van der Aalst. *Process Mining - Discovery, Conformance and Enhancement of Business Processes*. Springer, 2011.
2. W.M.P. van der Aalst, M.H. Schonenberg, and M. Song. Time prediction based on process mining. *Information Systems*, 36(2):450–475, 2011.
3. M. Adams, A.H.M. ter Hofstede, W.M.P. van der Aalst, and D. Edmond. Dynamic, extensible and context-aware exception handling for workflows. In *CoopIS*, volume 4803 of *LNC3*, pages 95–112. Springer, 2007.
4. C.J. Alberts and A.J. Dorofee. OCTAVE criteria, version 2.0. Technical Report CMU/SEI-2001-TR-016, Carnegie Mellon University, 2001.
5. S. Alter. A work system view of DSS in its fourth decade. *DSS*, 38, December 2004.
6. S. Alter and R. Wright. Validating work system principles for use in systems analysis and design. In *Proceedings of the 30th International Conference on Information System (ICIS 2010)*. AIS eLibrary, 2010.
7. B. Barber and J. Davey. The use of the CCTA Risk Analysis and Management Methodology CRAMM in health information systems. In *MEDINFO*. North Holland Publishing, 1992.
8. Basel Committee on Bankin Supervision. *Basel II - International Convergence of Capital Measurement and Capital Standards*, 2006.
9. A. Charfi and M. Mezini. AO4BPEL: An aspect-oriented extension to BPEL. *World Wide Web*, 10(3):309–344, 2007.

10. International Electrotechnical Commission. *IEC 61025 Fault Tree Analysis (FTA)*, 1990.
11. R. Conforti, G. Fortino, M. La Rosa, and A.H.M. ter Hofstede. History-aware, real-time risk detection in business processes. In *CoopIS*, volume 7044 of *LNCS*. Springer, 2011.
12. P. Dadam and M. Reichert. The ADEPT project: a decade of research and development for robust and flexible process support. *CSRD*, 23:81–97, 2009.
13. M. Dumas, W.M.P. van der Aalst, and A.H.M. ter Hofstede. *Process-Aware Information Systems: Bridging People and Software through Process Technology*. Wiley & Sons, 2005.
14. Johann Eder, Horst Pichler, Wolfgang Gruber, and Michael Ninaus. Personal schedules for workflow systems. In Arthur ter Hofstede and Mathias Weske, editors, *Business Process Management*, volume 2678 of *Lecture Notes in Computer Science*, pages 1018–1018. Springer Berlin / Heidelberg, 2003.
15. M. Gambini, M. La Rosa, S. Migliorini, and A.H.M. ter Hofstede. Automated error correction of business process models. In *BPM*, volume 6896 of *LNCS*. Springer, 2011.
16. G. Goluch, S. Tjoa, S. Jakoubi, and G. Quirchmayr. Deriving resource requirements applying risk-aware business process modeling and simulation. In *ECIS*. AISEL, 2008.
17. G. Hermosillo, L. Seinturier, and L. Duchien. Using complex event processing for dynamic business process adaptation. In *SCC*, pages 466–473. IEEE, 2010.
18. A.H.M. ter Hofstede, W.M.P. van der Aalst, M. Adams, and N. Russell, editors. *Modern Business Process Automation: YAWL and its Support Environment*. Springer, 2010.
19. W.G. Johnson. *MORT - The Management Oversight and Risk Tree*. U.S. Atomic Energy Commission, 1973.
20. Kensaku Kawamoto, Caitlin A Houlihan, E Andrew Balas, and David F Lobach. Improving clinical practice using clinical decision support systems: a systematic review of trials to identify features critical to success. *British Medical Journal*, 330(7494), April 2005.
21. M.S. Lund, B. Solhaug, and K. Stølen. *Model-Driven Risk Analysis - The CORAS Approach*. Springer, 2011.
22. Ronny Mans, Nick Russell, Wil van der Aalst, Arnold Moleman, and Piet Bakker. Schedule-aware workflow management systems. In Kurt Jensen, Susanna Donatelli, and Maciej Koutny, editors, *Transactions on Petri Nets and Other Models of Concurrency IV*, volume 6550 of *Lecture Notes in Computer Science*, pages 121–143. Springer Berlin / Heidelberg, 2010.
23. R. Muller, U. Greiner, and E. Rahm. AgentWork: a workflow system supporting rule-based workflow adaptation. *Data & Knowledge Engineering*, 51(2):223–256, 2004.
24. C. Ouyang, M. La Rosa, A.H.M. ter Hofstede, M. Dumas, and K. Shortland. Toward web-scale workflows for film production. *IEEE, Internet Computing*, 12(5):53–61, 2008.
25. C. Ouyang, M.T. Wynn, J-C. Kuhr, M. Adams, T. Becker, A.H.M. ter Hofstede, and C. Fidge. Workflow support for scheduling in surgical care processes. In *ECIS*.
26. K.I. Smith, R.M. Everson, J.E. Fieldsend, C. Murphy, and R. Misra. Dominance-based multiobjective simulated annealing. *IEEE TEC*, 12(3):323–342, 2008.
27. Standards Australia and Standards New Zealand. *Standard AS/NZS ISO 31000*, 2009.
28. Voluntary Interindustry Commerce Solutions Association. *Voluntary Inter-industry Commerce Standard (VICS)*. <http://www.vics.org>. Accessed: June 2011.
29. B. Weber, W. Wild, and R. Breu. CBRFlow: Enabling adaptive workflow management through conversational case-based reasoning. In *ECCBR*, volume 3155 of *LNCS*. Springer, 2004.

A Process Models

In this section we describe the business process models used for testing our approach.

A.1 Process A: Film Production Process

The process model in Figure 4 show the Film Production process [24]. This process starts with the collections of documents produced during the pre-production phase. These documents are collected during the execution of *Input Cast List*, *Input Crew List*, *Input Location Notes* and *Input Shooting Schedule*. Once all documents are collected the process is carried out on a daily basis, following two parallel paths. The first focuses on producing a “call sheet”. A call sheet is a schedule for a particular day. It is produced by the *Production Office* and delivered to all cast and crew one day in advance. The second path focuses on the technical part of the production, producing log of activities and technical notes. In particular the tasks *Fill Out Continuity Report* and *Fill Out Continuity Daily Report* are executed by the *Continuity person*, task *Fill Out Sound Sheets* is executed by *Sound Recordist*, task *Fill Out Camera Sheets* is executed by *Camera Assistant*, and task *Fill Out AD Report* is executed by *2nd Assistant Director*. Once these five activities are completed a daily progress report is generated in *Create DPR* and distributed to the *Producer* and *Executive Producer* in *Distribute DPR*.

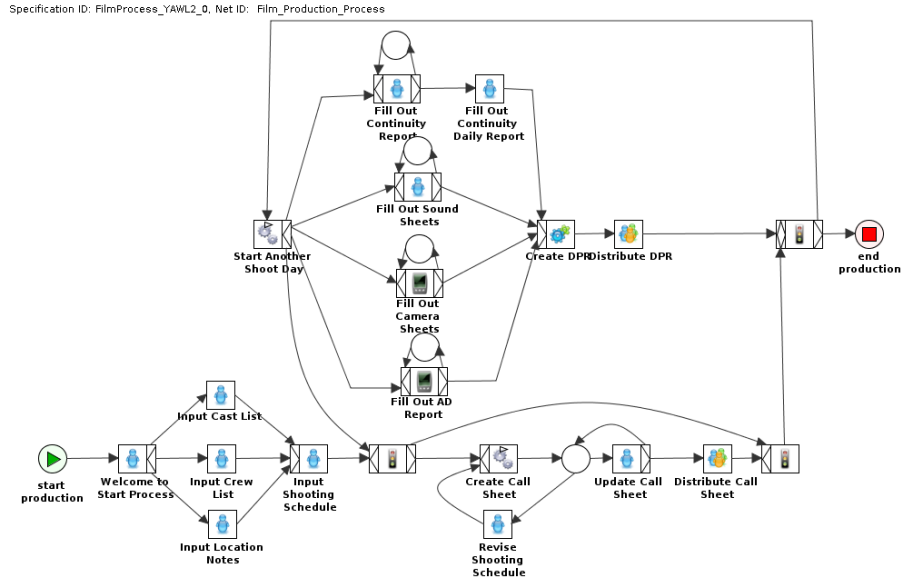


Fig. 4. Film Production Process (using YAWL [18] notation).

A.2 Process B: Ordering Process

The process model in Figure 5 show the Ordering subprocess of the Order Fulfilment process. This process deals with the placing of an purchase order. The process start with the creation of a purchase order, carried out by a *Process Order Manager* (*Create Purchase Order* task), and need to be approved by a *Senior Supply Officer* (*Approve Purchase Order* task). Once an order is approved it can be modified and in this case need to be approved again, or confirmed by the same *Process Order Manager* who created it. If an order is not approved within three days it is automatically cancelled.

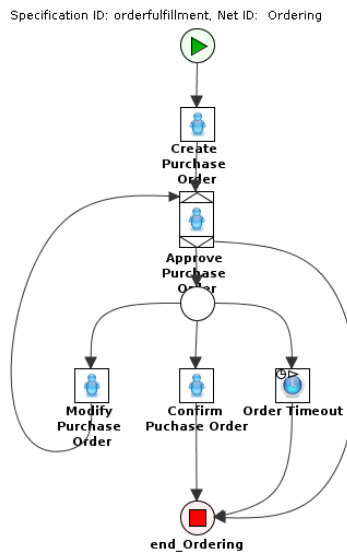


Fig. 5. Order Fulfilment Process: Ordering Subprocess (using YAWL [18] notation).

A.3 Process C: Payment Process

The process model in Figure 6 show the Payment subprocess of the Order Fulfilment process. This process starts after the freight has been picked up by a carrier and deals with the shipment and freight payment executed in parallel. The freight payment whether required starts with the production of a freight invoice done by a Supply Admin Officer and ends with the processing of the freight payment. During the payment of shipment costs, first a Shipment Invoice is produced for costs related to a specific order. If payment has been made in advance, a Finance Officer simply issues a Shipment Remittance Advice to the customer specifying the amount paid. Otherwise, the Finance Officer issues a Shipment Payment Order, which requires approval by a Senior Finance Officer (a superior of the Finance Officer) who may request amendments be made by the Finance Officer that issued the Order. After the document is finalized and the customer

has paid, an Account Manager can process the payment. If the customer underpays, the Account Manager issues a Debit Adjustment, the customer makes a further payment and the payment is reprocessed. If a customer overpays, the Account Manager issues a Credit Adjustment. In the latter case and in the case of correct payment, the Payment subprocess completes.

Specification ID: orderfulfillment, Net ID: Payment

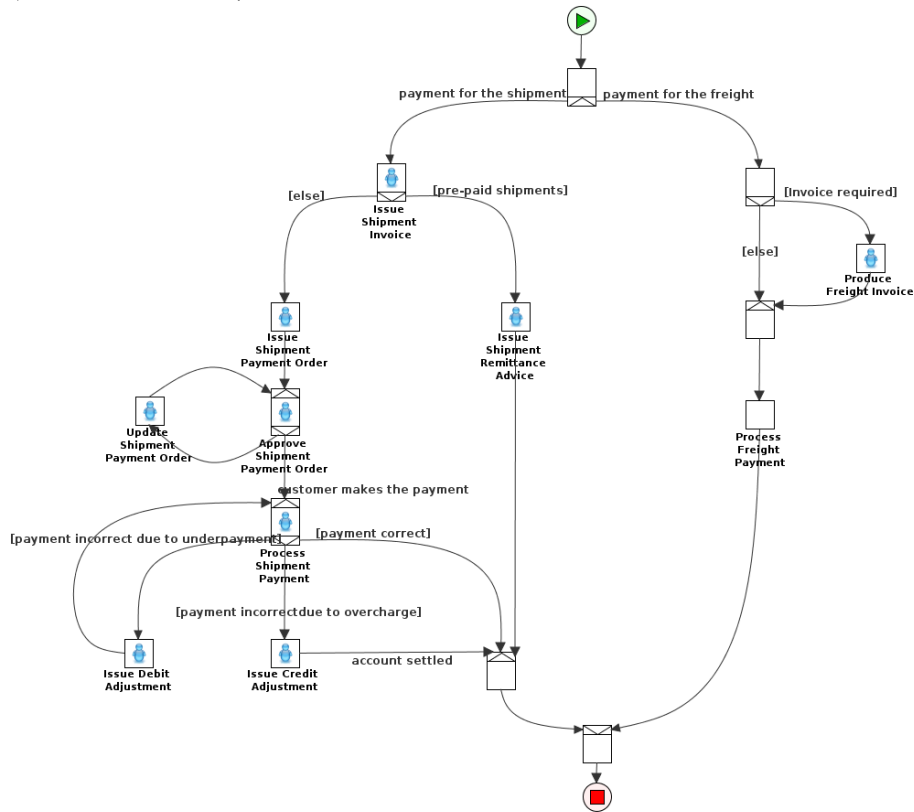


Fig. 6. Order Fulfillment Process: Payment Subprocess (using YAWL [18] notation).

A.4 Process D: Freight in Transit Process

The process model in Figure 7 show the Freight in Transit subprocess of the Order Fulfillment process. This process starts after the freight has been picked up by a carrier and deals with its transit until it is delivered. During its delivery of a freight a *courier* may issue one or more trackpoint notices, that are then used by the *Carrier Admin Officer* to generate a report. At the same time a *Client Liaison* can initiate an inquiry about the shipment status of the freight. Once the freight is physically delivered and

a trackpoint report is generated, the process end with a *Process Order Manager* who create an acceptance certificate.

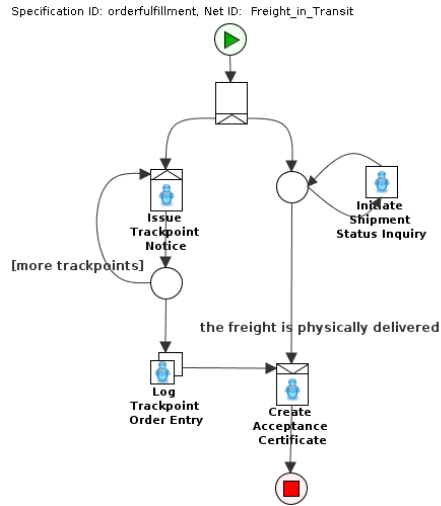


Fig. 7. Order Fulfilment Process: Freight in Transit Subprocess (using YAWL [18] notation).

B Process Risks

In this section we show the definition of the risks used for each the process models described in Appendix A.

B.1 First Experiment: Risks for Process A (Film Production)

Here we list the seven risks associated with Process A.

Risk1:

```
<sensor name="risk1">
  <vars>
    <var name="a11" mapping="case(current).Fill Out Continuity Daily Report(StartTimeInMillis)" type="" />
    <var name="a12" mapping="case(current).Fill Out Continuity Daily Report(TimeEstimationInMillis)" type="" />
    <var name="b11" mapping="case(current).Fill Out Sound Sheets(StartTimeInMillis)" type="" />
    <var name="b12" mapping="case(current).Fill Out Sound Sheets(TimeEstimationInMillis)" type="" />
    <var name="c11" mapping="case(current).Fill Out Camera Sheets(isStarted)" type="" />
    <var name="c12" mapping="case(current).Fill Out Camera Sheets(isCompleted)" type="" />
    <var name="d11" mapping="case(current).Fill Out AD Report(isStarted)" type="" />
    <var name="d12" mapping="case(current).Fill Out AD Report(isCompleted)" type="" />
  </vars>
  <riskCondition>((a11+a12)&lt;(b11+b12))&amp;((c11&amp;!c12)|(d11&amp;!d12))</riskCondition>
  <riskProbability />
  <riskThreshold />
  <riskMessage />
</sensor>
```

```

<faultCondition />
<faultProbability />
<faultThreshold />
<faultMessage />
</sensor>

```

Risk2:

```

<sensor name="risk2">
<vars>
<var name="a21" mapping="case(current).Fill Out Continuity Daily Report(allocateResource)" type="" />
<var name="b21" mapping="case(current).Fill Out Sound Sheets(allocateResource)" type="" />
</vars>
<riskCondition>a21==b21</riskCondition>
<riskProbability />
<riskThreshold />
<riskMessage />
<faultCondition />
<faultProbability />
<faultThreshold />
<faultMessage />
</sensor>

```

Risk3:

```

<sensor name="risk3">
<vars>
<var name="a31" mapping="case(current).Revise Shooting Schedule(Count)" type="" />
<var name="b31" mapping="case(current).Input Shooting Schedule(CompleteTimeInMillis)" type="" />
<var name="a32" mapping="case(current).Revise Shooting Schedule(TimeEstimationInMillis)" type="" />
<var name="a33" mapping="case(current).Revise Shooting Schedule(isCompleted)" type="" />
<var name="a34" mapping="case(current).Revise Shooting Schedule(isStarted)" type="" />
<var name="a35" mapping="case(current).Revise Shooting Schedule(StartTimeInMillis)" type="" />
</vars>
<riskCondition>(a31&gt;5) &amp; (!a33&amp;a34) &amp; (a35-b31&gt;a32)</riskCondition>
<riskProbability />
<riskThreshold />
<riskMessage />
<faultCondition />
<faultProbability />
<faultThreshold />
<faultMessage />
</sensor>

```

Risk4:

```

<sensor name="risk4">
<vars>
<var name="a41" mapping="case(current).Revise Shooting Schedule(PassTimeInMillis)" type="" />
<var name="a52" mapping="case(current).Revise Shooting Schedule(TimeEstimationInMillis)" type="" />
</vars>
<riskCondition>a41+a52&gt;5</riskCondition>
<riskProbability />
<riskThreshold />
<riskMessage />
<faultCondition />
<faultProbability />
<faultThreshold />
<faultMessage />
</sensor>

```

Risk5:

```

<sensor name="risk5">
<vars>
<var name="a51" mapping="case(current).Input Cast List.callSheet" type="" />
<var name="b51" mapping="case(current).Input Crew List.timeSheetInfo" type="" />
<var name="c51" mapping="case(current).Fill Out Camera Sheets(allocateResource)" type="" />
<var name="d51" mapping="case(current).Fill Out Continuity Daily Report(allocateResource)" type="" />

```

```

</vars>
<riskCondition>!(a51==b51)&amp;! (c51==d51)</riskCondition>
<riskProbability />
<riskThreshold />
<riskMessage />
<faultCondition />
<faultProbability />
<faultThreshold />
<faultMessage />
</sensor>

```

Risk6:

```

<sensor name="risk6">
<vars>
  <var name="a61" mapping="case(current).Input Crew List.production" type="" />
  <var name="b61" mapping="case(current).Input Cast List(PassTimeInMillis)" type="" />
  <var name="c61" mapping="case(current).Input Shooting Schedule(PassTimeInMillis)" type="" />
</vars>
<riskCondition>(a61==5)&amp;((b61+c61)&lt;20)</riskCondition>
<riskProbability />
<riskThreshold />
<riskMessage />
<faultCondition />
<faultProbability />
<faultThreshold />
<faultMessage />
</sensor>

```

Risk7:

```

<sensor name="risk7">
<vars>
  <var name="a71" mapping="case(current).Fill Out Continuity Report.producer" type="" />
  <var name="b71" mapping="case(current).Fill Out Sound Sheets.production" type="" />
  <var name="c71" mapping="case(current).Fill Out Camera Sheets.production" type="" />
  <var name="d71" mapping="case(current).Fill Out AD Report.production" type="" />
</vars>
<riskCondition>!(c71==b71)&amp;(! (b71==a71) || (b71==d71))</riskCondition>
<riskProbability />
<riskThreshold />
<riskMessage />
<faultCondition />
<faultProbability />
<faultThreshold />
<faultMessage />
</sensor>

```

B.2 First Experiment: Risks for Process B (Ordering)

Here we list the four risks associated with Process B.

Risk2:

```

<sensor name="risk2">
<vars>
  <var name="a21" mapping="case(current).Initiate Shipment Status Inquiry(allocateResource)" type="" />
  <var name="b21" mapping="case(current).Issue Trackpoint Notice(allocateResource)" type="" />
</vars>
<riskCondition>a21==b21</riskCondition>
<riskProbability />
<riskThreshold />
<riskMessage />
<faultCondition />
<faultProbability />
<faultThreshold />
<faultMessage />
</sensor>

```

Risk4:

```
<sensor name="risk4">
  <vars>
    <var name="a41" mapping="case(current).Issue Trackpoint Notice(PassTimeInMillis)" type="" />
    <var name="a42" mapping="case(current).Issue Trackpoint Notice(TimeEstimationInMillis)" type="" />
  </vars>
  <riskCondition>a41+a42>5</riskCondition>
  <riskProbability />
  <riskThreshold />
  <riskMessage />
  <faultCondition />
  <faultProbability />
  <faultThreshold />
  <faultMessage />
</sensor>
```

Risk5:

```
<sensor name="risk5">
  <vars>
    <var name="a51" mapping="case(current).Issue Trackpoint Notice.AcceptanceCertificate" type="" />
    <var name="b51" mapping="case(current).Initiate Shipment Status Inquiry.TrackpointNotice" type="" />
    <var name="c51" mapping="case(current).Log Trackpoint Order Entry(allocateResource)" type="" />
    <var name="d51" mapping="case(current).Create Acceptance Certificate(allocateResource)" type="" />
  </vars>
  <riskCondition>! (a51==b51) & ! (c51==d51)</riskCondition>
  <riskProbability />
  <riskThreshold />
  <riskMessage />
  <faultCondition />
  <faultProbability />
  <faultThreshold />
  <faultMessage />
</sensor>
```

Risk6:

```
<sensor name="risk6">
  <vars>
    <var name="a61" mapping="case(current).Initiate Shipment Status Inquiry.Report" type="" />
    <var name="b61" mapping="case(current).Issue Trackpoint Notice(PassTimeInMillis)" type="" />
    <var name="c61" mapping="case(current).Log Trackpoint Order Entry(PassTimeInMillis)" type="" />
  </vars>
  <riskCondition>(a61==5) & ((b61+c61) < 20)</riskCondition>
  <riskProbability />
  <riskThreshold />
  <riskMessage />
  <faultCondition />
  <faultProbability />
  <faultThreshold />
  <faultMessage />
</sensor>
```

B.3 First experiment: Risks for Process C (Payment)

Here we list the three risks associated with Process C.

Risk3:

```
<sensor name="risk3">
  <vars>
    <var name="a31" mapping="case(current).Approve Purchase Order(Count)" type="" />
    <var name="b31" mapping="case(current).Create Purchase Order(CompleteTime)" type="" />
    <var name="a32" mapping="case(current).Approve Purchase Order(TimeEstimationInMillis)" type="" />
    <var name="a33" mapping="case(current).Approve Purchase Order(isCompleted)" type="" />
  </vars>
```

```

    <var name="a34" mapping="case(current).Approve Purchase Order(isStarted)" type="" />
    <var name="a35" mapping="case(current).Approve Purchase Order(StartTimeInMillis)" type="" />
  </vars>
  <riskCondition>(a31>5) &amp; (!a33&amp;a34) &amp; (a35+b31>a32)</riskCondition>
  <riskProbability />
  <riskThreshold />
  <riskMessage />
  <faultCondition />
  <faultProbability />
  <faultThreshold />
  <faultMessage />
</sensor>

```

Risk4:

```

<sensor name="risk4">
  <vars>
    <var name="a41" mapping="case(current).Modify Purchase Order(PassTimeInMillis)" type="" />
    <var name="a42" mapping="case(current).Modify Purchase Order(TimeEstimationInMillis)" type="" />
  </vars>
  <riskCondition>a41+a42>5</riskCondition>
  <riskProbability />
  <riskThreshold />
  <riskMessage />
  <faultCondition />
  <faultProbability />
  <faultThreshold />
  <faultMessage />
</sensor>

```

Risk6:

```

<sensor name="risk6">
  <vars>
    <var name="a61" mapping="case(current).Modify Purchase Order.POApproval" type="" />
    <var name="b61" mapping="case(current).Approve Purchase Order(PassTimeInMillis)" type="" />
    <var name="c61" mapping="case(current).Confirm Purchase Order(PassTimeInMillis)" type="" />
  </vars>
  <riskCondition>(a61==5) &amp; ((b61+c61) &lt;20)</riskCondition>
  <riskProbability />
  <riskThreshold />
  <riskMessage />
  <faultCondition />
  <faultProbability />
  <faultThreshold />
  <faultMessage />
</sensor>

```

B.4 First experiment: Risks for Process D (Freight in Transit)

Here we list the five risks associated with Process D.

Risk2:

```

<sensor name="risk2">
  <vars>
    <var name="a21" mapping="case(current).Produce Freight Invoice(allocateResource)" type="" />
    <var name="b21" mapping="case(current).Issue Shipment Payment Order(allocateResource)" type="" />
  </vars>
  <riskCondition>a21==b21</riskCondition>
  <riskProbability />
  <riskThreshold />
  <riskMessage />
  <faultCondition />
  <faultProbability />
  <faultThreshold />
  <faultMessage />
</sensor>

```

Risk3:

```
<sensor name="risk3">
  <vars>
    <var name="a31" mapping="case(current).Approve Shipment Payment Order(Count)" type="" />
    <var name="b31" mapping="case(current).Issue Shipment Payment Order(CompleteTime)" type="" />
    <var name="a32" mapping="case(current).Approve Shipment Payment Order(TimeEstimationInMillis)" type="" />
    <var name="a33" mapping="case(current).Approve Shipment Payment Order(isCompleted)" type="" />
    <var name="a34" mapping="case(current).Approve Shipment Payment Order(isStarted)" type="" />
    <var name="a35" mapping="case(current).Approve Shipment Payment Order(StartTimeInMillis)" type="" />
  </vars>
  <riskCondition>(a31>5) & (!a33&a34) & (a35+b31>a32) </riskCondition>
  <riskProbability />
  <riskThreshold />
  <riskMessage />
  <faultCondition />
  <faultProbability />
  <faultThreshold />
  <faultMessage />
</sensor>
```

Risk4:

```
<sensor name="risk4">
  <vars>
    <var name="a41" mapping="case(current).Approve Shipment Payment Order(PassTimeInMillis)" type="" />
    <var name="a42" mapping="case(current).Approve Shipment Payment Order(TimeEstimationInMillis)" type="" />
  </vars>
  <riskCondition>a41+a42>5</riskCondition>
  <riskProbability />
  <riskThreshold />
  <riskMessage />
  <faultCondition />
  <faultProbability />
  <faultThreshold />
  <faultMessage />
</sensor>
```

Risk5:

```
<sensor name="risk5">
  <vars>
    <var name="a51" mapping="case(current).Produce Freight Invoice.ShipmentPaymentOrder" type="" />
    <var name="b51" mapping="case(current).Approve Shipment Payment Order.ShipmentPaymentOrder" type="" />
    <var name="c51" mapping="case(current).Issue Shipment Invoice(allocateResource)" type="" />
    <var name="d51" mapping="case(current).Issue Shipment Payment Order(allocateResource)" type="" />
  </vars>
  <riskCondition>!(a51==b51) & !(c51==d51) </riskCondition>
  <riskProbability />
  <riskThreshold />
  <riskMessage />
  <faultCondition />
  <faultProbability />
  <faultThreshold />
  <faultMessage />
</sensor>
```

Risk6:

```
<sensor name="risk6">
  <vars>
    <var name="a61" mapping="case(current).Produce Freight Invoice.ShipmentPaymentOrder" type="" />
    <var name="b61" mapping="case(current).Issue Shipment Payment Order(PassTimeInMillis)" type="" />
    <var name="c61" mapping="case(current).Approve Shipment Payment Order(PassTimeInMillis)" type="" />
  </vars>
  <riskCondition>(a61==5) & ((b61+c61)<20) </riskCondition>
  <riskProbability />
  <riskThreshold />
</sensor>
```

```

<riskMessage />
<faultCondition />
<faultProbability />
<faultThreshold />
<faultMessage />
</sensor>

```

B.5 Second experiment: Risks Process C (Payment Process)

Here we list the three risks used in the second experiment. These risks are different from the ones used in the first experiment since they use historical information to compute the risk probability.

Overtime Process risk:

```

<sensor name="Overtime Process">
<vars>
  <var name="d" mapping="5" type="" />
  <var name="Tc" mapping="case(current).Payment(PassTimeInMillis)" type="" />
  <var name="Te" mapping="case(current).(TimeEstimationInMillis)" type="" />
</vars>
<riskCondition>(1-(d*24*60*60*1000)/(Te+Tc))>0.6</riskCondition>
<riskProbability />
<riskThreshold />
<riskMessage />
<faultCondition />
<faultProbability />
<faultThreshold />
<faultMessage />
</sensor>

```

Approval Fraud risk:

```

<sensor name="Approval Fraud">
<vars>
  <var name="sfoA" mapping="case(current).Approve Shipment Payment Order(isAllocated)" type="" />
  <var name="sfo1" mapping="case(current).Approve Shipment Payment Order(allocateResource)" type="" />
  <var name="c" mapping="case(current).Issue Shipment Invoice.ShipmentInvoice.Company" type="" />
  <var name="ASPOa" mapping="case(current).Approve Shipment Payment Order(OfferTimeInMillis)" type="" />
  <var name="ASPOn" mapping="case(Approve Shipment Payment Order(completeResource)=sfo1 AND
  Issue Shipment Invoice.ShipmentInvoice.Company=c AND
  Approve Shipment Payment Order(CompleteTimeInMillis)>(ASPOa-(5*24*60*60*1000)) AND
  (ID)!=[IDCurr]).Approve Shipment Payment Order(CountElements)" type="" />
  <var name="sfo2" mapping="case(Issue Shipment Invoice.ShipmentInvoice.Company=c AND
  Approve Shipment Payment Order(isCompleted)="&quot;true&quot;
  AND Approve Shipment Payment Order(CompleteTimeInMillis)>(ASPOa-(5*24*60*60*1000))
  AND (ID)!=[IDCurr]).Approve Shipment Payment Order(completeResource)" type="" />
  <var name="sfo" mapping="case(current).Approve Shipment Payment Order(offerDistribution)" type="" />
</vars>
<riskCondition>
  (ASPOn&gt;0) | (!sfoA&amp;sfo2.startedMinNumber==0) &amp; (sfo.startedMinNumberExcept.sfo2&gt;=1)
</riskCondition>
<riskProbability />
<riskThreshold />
<riskMessage />
<faultCondition />
<faultProbability />
<faultThreshold />
<faultMessage />
</sensor>

```

Order Unfulfillment risk:

```

<sensor name="Order Unfulfillment">
<vars>

```

```

<var name="USPOuc" mapping="case(current).Update Shipment Payment Order(Count)" type="" />
<var name="USPOu5" mapping="case(Update Shipment Payment Order(Count)>=5).
  Update Shipment Payment Order(CountElements)" type="" />
<var name="USPOus" mapping="case(Update Shipment Payment Order(Count)>=USPOuc AND
  Process Shipment Payment(isOffered)=&quot;true&quot;);.
  Update Shipment Payment Order(CountElements)" type="" />
</vars>
<riskCondition>(USPOu5/USPOus)>0.7</riskCondition>
<riskProbability />
<riskThreshold />
<riskMessage />
<faultCondition />
<faultProbability />
<faultThreshold />
<faultMessage />
</sensor>

```

C Solutions second experiment

Here we list the solutions proposed for the second experiment.

Model: - Time: 60009 -

Solution1

```

risk1 0.0, prob1 0.0, risk2 0.0, prob2 0.0, risk3 0.0, prob3 0.0, cost 50.0
RemoveAllocateRes: Approve Shipment Payment Order <participant id="PA-223a865e-9b16-4c0b-a496-4e445eaadlec">
RemoveOfferRes: Approve Shipment Payment Order <participant id="PA-0b7f1d4c-f3b2-4fd4-8e77-c33f65bb2821">
RemoveOfferRes: Approve Shipment Payment Order <participant id="PA-223a865e-9b16-4c0b-a496-4e445eaadlec">
RemoveOfferRes: Approve Shipment Payment Order <participant id="PA-3dad7d51-ffbd-4e02-aeba-462deac95ef8">
Rollback: Update Shipment Payment Order
Rollback: Approve Shipment Payment Order
Tauize: Issue Shipment Remittance Advice
Rollback: Update Shipment Payment Order
Tauize: Approve Shipment Payment Order

```

Solution2

```

risk1 0.0, prob1 0.0, risk2 0.0, prob2 0.0, risk3 0.0, prob3 0.0, cost 50.0
RemoveAllocateRes: Approve Shipment Payment Order <participant id="PA-223a865e-9b16-4c0b-a496-4e445eaadlec">
RemoveOfferRes: Approve Shipment Payment Order <participant id="PA-0b7f1d4c-f3b2-4fd4-8e77-c33f65bb2821">
RemoveOfferRes: Approve Shipment Payment Order <participant id="PA-223a865e-9b16-4c0b-a496-4e445eaadlec">
RemoveOfferRes: Approve Shipment Payment Order <participant id="PA-3dad7d51-ffbd-4e02-aeba-462deac95ef8">
Rollback: Update Shipment Payment Order
Rollback: Approve Shipment Payment Order
Tauize: Issue Shipment Remittance Advice
Rollback: Update Shipment Payment Order
Tauize: Issue Debit Adjustment

```

Solution3

```

risk1 0.0, prob1 0.0, risk2 0.0, prob2 0.0, risk3 25.0, prob3 0.0, cost 41.0
RemoveAllocateRes: Approve Shipment Payment Order <participant id="PA-223a865e-9b16-4c0b-a496-4e445eaadlec">
RemoveOfferRes: Approve Shipment Payment Order <participant id="PA-0b7f1d4c-f3b2-4fd4-8e77-c33f65bb2821">
RemoveOfferRes: Approve Shipment Payment Order <participant id="PA-223a865e-9b16-4c0b-a496-4e445eaadlec">
RemoveOfferRes: Approve Shipment Payment Order <participant id="PA-3dad7d51-ffbd-4e02-aeba-462deac95ef8">
Rollback: Update Shipment Payment Order
Rollback: Approve Shipment Payment Order
Tauize: Update Shipment Payment Order
Tauize: Process Shipment Payment

```

Solution4

```

risk1 0.0, prob1 0.0, risk2 0.0, prob2 0.0, risk3 25.0, prob3 0.0, cost 41.0
RemoveAllocateRes: Approve Shipment Payment Order <participant id="PA-223a865e-9b16-4c0b-a496-4e445eaadlec">
RemoveOfferRes: Approve Shipment Payment Order <participant id="PA-0b7f1d4c-f3b2-4fd4-8e77-c33f65bb2821">
RemoveOfferRes: Approve Shipment Payment Order <participant id="PA-223a865e-9b16-4c0b-a496-4e445eaadlec">
RemoveOfferRes: Approve Shipment Payment Order <participant id="PA-3dad7d51-ffbd-4e02-aeba-462deac95ef8">
Rollback: Update Shipment Payment Order

```


Rollback: Approve Shipment Payment Order
Tauize: Update Shipment Payment Order
Tauize: Approve Shipment Payment Order

Solution5

risk1 0.0, prob1 0.0, risk2 0.0, prob2 0.0, risk3 75.0, prob3 0.0, cost 19.0
RemoveAllocateRes: Approve Shipment Payment Order <participant id="PA-223a865e-9b16-4c0b-a496-4e445eadlec">
Tauize: Issue Shipment Remittance Advice
Tauize: Process Shipment Payment
RestoreAllocateRes: Approve Shipment Payment Order <participant id="PA-3dad7d51-ffbd-4e02-aeba-462deac95ef8">

Solution6

risk1 0.0, prob1 0.0, risk2 0.0, prob2 0.0, risk3 75.0, prob3 0.0, cost 19.0
RemoveAllocateRes: Approve Shipment Payment Order <participant id="PA-223a865e-9b16-4c0b-a496-4e445eadlec">
RestoreAllocateRes: Approve Shipment Payment Order <participant id="PA-0b7f1d4c-f3b2-4fd4-8e77-c33f65bb2821">
Tauize: Issue Debit Adjustment
Tauize: Issue Shipment Remittance Advice

Solution7

risk1 75.0, prob1 0.0, risk2 0.0, prob2 0.0, risk3 00.0, prob3 0.0, cost 41.0
RemoveAllocateRes: Approve Shipment Payment Order <participant id="PA-223a865e-9b16-4c0b-a496-4e445eadlec">
RemoveOfferRes: Approve Shipment Payment Order <participant id="PA-0b7f1d4c-f3b2-4fd4-8e77-c33f65bb2821">
RemoveOfferRes: Approve Shipment Payment Order <participant id="PA-223a865e-9b16-4c0b-a496-4e445eadlec">
RemoveOfferRes: Approve Shipment Payment Order <participant id="PA-3dad7d51-ffbd-4e02-aeba-462deac95ef8">
Rollback: Update Shipment Payment Order
Rollback: Approve Shipment Payment Order
Tauize: Issue Shipment Remittance Advice
Rollback: Update Shipment Payment Order

Solution8

risk1 75.0, prob1 0.0, risk2 0.0, prob2 0.0, risk3 25.0, prob3 0.0, cost 32.0
RemoveAllocateRes: Approve Shipment Payment Order <participant id="PA-223a865e-9b16-4c0b-a496-4e445eadlec">
RemoveOfferRes: Approve Shipment Payment Order <participant id="PA-0b7f1d4c-f3b2-4fd4-8e77-c33f65bb2821">
RemoveOfferRes: Approve Shipment Payment Order <participant id="PA-223a865e-9b16-4c0b-a496-4e445eadlec">
RemoveOfferRes: Approve Shipment Payment Order <participant id="PA-3dad7d51-ffbd-4e02-aeba-462deac95ef8">
Rollback: Update Shipment Payment Order
Rollback: Approve Shipment Payment Order
Tauize: Update Shipment Payment Order

Solution9

risk1 75.0, prob1 0.0, risk2 0.0, prob2 0.0, risk3 75.0, prob3 0.0, cost 1.0
RemoveAllocateRes: Approve Shipment Payment Order <participant id="PA-223a865e-9b16-4c0b-a496-4e445eadlec">
RestoreAllocateRes: Approve Shipment Payment Order <participant id="PA-0b7f1d4c-f3b2-4fd4-8e77-c33f65bb2821">

Solution10

risk1 0.0, prob1 0.0, risk2 75.0, prob2 0.0, risk3 75.0, prob3 0.0, cost 18.0
Tauize: Process Shipment Payment
Tauize: Issue Shipment Remittance Advice

total model generated: 1015