# Polynomial Analysis Algorithms for Free Choice Probabilistic Workflow Nets* 

Javier Esparza ${ }^{1}$, Philipp Hoffmann ${ }^{1}$, and Ratul Saha ${ }^{2}$<br>${ }^{1}$ Technische Universität München<br>${ }^{2}$ National University of Singapore


#### Abstract

We study Probabilistic Workflow Nets (PWNs), a model extending van der Aalst's workflow nets with probabilities. We give a semantics for PWNs in terms of Markov Decision Processes and introduce a reward model. Using a result by Varacca and Nielsen, we show that the expected reward of a complete execution of the PWN is independent of the scheduler. Extending previous work on reduction of non-probabilistic workflow nets, we present reduction rules that preserve the expected reward. The rules lead to a polynomial-time algorithm in the size of the PWN (not of the Markov decision process) for the computation of the expected reward. In contrast, since the Markov decision process of PWN can be exponentially larger than the PWN itself, all algorithms based on constructing the Markov decision process require exponential time. We report on a sample implementation and its performance on a collection of benchmarks.


## 1 Introduction

Workflow Petri Nets are a class of Petri nets for the representation and analysis of business processes $[1,2,5]$. They are a popular formal back-end for different notations like BPMN (Business Process Modeling Notation), EPC (Event-driven Process Chain), or UML Activity Diagrams.

There is recent interest in extending these notations, in particular BPMN, with the concept of cost (see e.g. [15, 18, 19]). The final goal is the development of tool support for computing the worst-case or the average cost of a business process. A sound foundation for the latter requires to extend Petri nets with probabilities and rewards. Since Petri nets can express complex interplay between nondeterminism and concurrency, the extension is a nontrivial semantic problem which has been studied in detail (see e.g. [21,3,4] for untimed probabilistic extensions and [7] for timed extensions).

Fortunately, giving a semantics to probabilistic Petri nets is much simpler for confusion-free Petri nets [21,3], a class that already captures many controlflow constructs of BPMN. In particular, confusion-free Petri nets strictly contain Workflow Graphs, also called free-choice Workflow Nets [1, 11, 12, 9]. In this

[^0]paper we study free choice Workflow Nets extended with rewards and probabilities. Rewards are modeled as real numbers attached to the transitions of the workflow, while, intuitively, probabilities are attached to transitions modeling nondeterministic choices. Our main result is the first polynomial algorithm for computing the expected reward of a workflow.

In order to define expected rewards, we give untimed, probabilistic confusionfree nets a semantics in terms of Markov Decision Processes (MDP), with rewards captured by a reward function. In a nutshell, at each reachable marking the enabled transitions are partitioned into clusters. All transitions of a cluster are in conflict, while transitions of different clusters are concurrent. In the MDP semantics, a scheduler selects one of the clusters, while the transition inside this cluster is chosen probabilistically. We use MDPs instead of probabilistic event structures, as in $[21,3,4]$, because for our purposes the semantics are equivalent, and an MDP semantics allows us to use the well established reward terminology for MDPs [17].

In our first contribution, we prove that the expected reward of a confusionfree workflow net is independent of the scheduler resolving the nondeterministic choices, and so we can properly speak of the expected reward of a freechoice workflow. The proof relies on a result by Varacca and Nielsen [20] on Mazurkiewicz equivalent schedulers.

Since MDP semantics of concurrent systems captures all possible interleavings of transitions, the MDP of a free-choice workflow can grow exponentially in the size of the net, and so MDP-based algorithms for the expected reward have exponential runtime. In our second contribution we provide a polynomial-time reduction algorithm consisting of the repeated application of a set of reduction rules that simplify the workflow while preserving its expected reward. Our rules are an extension to the probabilistic case of a set of rules for free-choice Colored Workflow Nets recently presented in [9]. The rules allow one to merge two alternative tasks, summarize or shortcut two consecutive tasks by one, and replace a loop with a probabilistic guard and an exit by a single task. We prove that the rules preserve the expected reward. The proof makes crucial use of the fact that the expected reward is independent of the scheduler: Given the two workflow nets before and after the reduction, we choose suitable schedulers for both of them, and show that the expected rewards under these two schedulers coincide.

Finally, as a third contribution we report on a prototype implementation, and on experimental results on a benchmark suite of nearly 1500 workflows derived from industrial business processes. We compare our algorithm with the different algorithms based on the construction of the MDP implemented in Prism [14].

## 2 Workflow Nets

We recall the definition of a workflow net, and the properties of soundness and 1 -safeness.

Definition 1 (Workflow Net [1]). A workflow net is a tuple $\mathcal{W}=(P, T, F, i, o)$ where

- $P$ is a finite set of places.
$-T$ is a finite set of transitions ( $P \cap T=\emptyset$ ).
$-F \subseteq(P \times T) \cup(T \times P)$ is a set of arcs.
$-i, o \in P$ are distinguished initial and final places such that $i$ has no incoming arcs and o has no outgoing arcs.
- The graph $(P \cup T, F \cup(o, i))$ is strongly connected.

We write ${ }^{\bullet} p$ and $p^{\bullet}$ to denote the input and output transitions of a place $p$, respectively, and similarly ${ }^{\bullet} t$ and $t^{\bullet}$ for the input and output places of a transition $t$. A marking $M$ is a function from $P$ to the natural numbers that assigns a number of tokens to each place. A transition $t$ is enabled at $M$ if all places of ${ }^{\bullet} t$ contain at least one token in $M$. An enabled transition may fire, removing a token from each place of ${ }^{\bullet} t$ and adding one token to each place of $t^{\bullet}$. We write $M \xrightarrow{t} M^{\prime}$ to denote that $t$ is enabled at $M$ and its firing leads to $M^{\prime}$. The initial marking (final marking) of a workflow net, denoted by $\boldsymbol{i}(\boldsymbol{o})$, puts one token on place $i$ (on place $o$ ), and no tokens elsewhere. A sequence of transitions $\sigma=t_{1} t_{2} \cdots t_{n}$ is an occurrence sequence or firing sequence if there are markings $M_{1}, M_{2}, \ldots, M_{n}$ such that $\boldsymbol{i} \xrightarrow{t_{1}} M_{1} \cdots M_{n-1} \xrightarrow{t_{n}} M_{n}$. Fin $\mathcal{W}_{\mathcal{W}}$ is the set of all firing sequences of $\mathcal{W}$ that end in the final marking. A marking is reachable if some occurrence sequence ends in that marking.


Fig. 1: Three workflow nets

Definition 2 (Soundness and 1-safeness [1]). A workflow net is sound if the final marking is reachable from any reachable marking, and for every transition $t$ there is a reachable marking that enables $t$. A workflow net is 1 -safe if $M(p) \leq 1$ for every reachable marking $M$ and for every place $p$.

Figure 1 shows three sound and 1-safe workflow nets. In this paper we only consider 1-safe workflow nets, and identify a marking with the set of places
that are marked. Markings which only mark a single place are written without brackets and in bold, like the initial marking $\boldsymbol{i}$. In general, deciding if a workflow net is sound and 1-safe is a PSPACE-complete problem. However, for the class of free-choice workflow nets, introduced below, and for which we obtain our main result, there exists a polynomial algorithm [6].

### 2.1 Confusion-Free and Free-Choice Workflow Nets

We recall the notions of independent transitions and transitions in conflict.
Definition 3 (Independent Transitions, Conflict). Two transitions $t_{1}, t_{2}$ of a workflow net are independent if ${ }^{\bullet} t_{1} \cap \bullet t_{2}=\emptyset$. Two transitions are in conflict at a marking $M$ if $M$ enables both of them and they are not independent. The set of transitions in conflict with a transition $t$ at a marking $M$ is called the conflict set of $t$ at $M$.

In Figure 1 transitions $t_{2}$ and $t_{4}$ of the left workflow are independent, while $t_{2}$ and $t_{3}$ are in conflict. The conflict set of $t_{2}$ at the marking $\left\{p_{1}, p_{2}\right\}$ is $\left\{t_{2}, t_{3}\right\}$, but at the marking $\left\{p_{1}, p_{4}\right\}$ it is $\left\{t_{2}\right\}$.

It is easy to see that in a 1 -safe workflow net two transitions enabled at a marking are either independent or in conflict. Assume that a 1-safe workflow net satisfies the following property: for every reachable marking $M$, the conflict relation at $M$ is an equivalence relation. Then, at every reachable marking $M$ we can partition the set of enabled transitions into equivalence classes, where transitions in the same class are in conflict and transitions of different classes are independent. For such nets we can introduce the following simple stochastic semantics: at each reachable marking an equivalence class is selected nondeterministically, and then a transition of the class is selected stochastically with probability proportional to a weight attached to the transition. However, not every workflow satisfies this property. For example, the workflow on the left of Figure 1 does not: at the reachable marking marking $\left\{p_{1}, p_{2}\right\}$ transition $t_{3}$ is in conflict with both $t_{2}$ and $t_{4}$, but $t_{2}$ and $t_{4}$ are independent. Confusion-free nets, whose probabilistic semantics is studied in [20], are a class of nets in which this kind of situation cannot occur.

Definition 4 (Confusion-Free Workflow Nets). A marking $M$ of a workflow net is confused if there are two independent transitions $t_{1}, t_{2}$ enabled at $M$ such that $M \xrightarrow{t_{1}} M^{\prime}$ and the conflict sets of $t_{2}$ at $M$ and at $M^{\prime}$ are different. $A$ 1-safe workflow net is confusion-free if no reachable marking is confused.

The workflows in the middle and on the right of Figure 1 are confusion-free.
Lemma 1 ([20]). Let $W$ be a 1-safe, confusion-free workflow net. For every reachable marking of $W$ the conflict relation on the transitions enabled at $M$ is an equivalence relation.

Unfortunately, deciding if a 1-safe workflow net is confusion-free is a PSPACEcomplete problem (this can be proved by an easy reduction from the reachability problem for 1 -safe Petri nets, see [8] for similar proofs). Free-choice workflow nets are a syntactically defined class of confusion-free workflow nets.

Definition 5 (Free-Choice Workflow Nets $[6,1]$ ). A workflow net is freechoice if for every two places $p_{1}, p_{2}$ either $p_{1}^{\bullet} \cap p_{2}^{\bullet}=\emptyset$ or $p_{1}^{\bullet}=p_{2}^{\bullet}$.

The workflow in the middle of Figure 1 is not free-choice, e.g. because of the places $p_{3}$ and $p_{4}$, but the one on the right is.

It is easy to see that free-choice workflow nets are confusion-free, but even more: in free-choice workflow nets, the conflict set of a transition $t$ is the same at all reachable markings that enable $t$. To formulate this, we use the notion of a cluster.

Definition 6 (Transition clusters). Let $\mathcal{W}=(P, T, F, i, o)$ be a free-choice workflow net. The cluster of $t \in T$ is the set of transitions $[t]=\left\{t^{\prime} \in T \mid\right.$ $\left.{ }^{\bullet} t \cap{ }^{\bullet} t^{\prime} \neq \emptyset\right\} .{ }^{3}$

By the free-choice property, if a marking enables a transition of a cluster, then it enables all of them. We say that the marking enables the cluster; we also say that a cluster fires if one of its transitions fires.

Proposition 1. - Let t be a transition of a free-choice workflow net. For every marking that enables $t$, the conflict set of $t$ at $M$ is the cluster $[t]$.

- Free-choice workflow nets are confusion-free.

Proof. The first part follows immediately from the free-choice property. For the second part, let $t_{1}, t_{2}$ be independent transitions enabled at a marking $M$ such that $M \xrightarrow{t_{1}} M^{\prime}$. By the free-choice property, for every $t \in\left[t_{1}\right]$ the transitions $t$ and $t_{2}$ are also independent. So the conflict sets of $t_{1}$ at $M$ and $M^{\prime}$ are both equal to $\left[t_{1}\right]$.

## 3 Probabilistic Workflow Nets

We introduce Probabilistic Workflow Nets, and give them a semantics in terms of Markov Decision Processes. We first recall some basic definitions.

### 3.1 Markov Decision Processes

For a finite set $Q$, let $\operatorname{dist}(Q)$ denote the set of probability distributions over $Q$.
Definition 7 (Markov Decision Process). A Markov Decision Process (MDP) is a tuple $\mathcal{M}=\left(Q, q_{0}\right.$, Steps $)$ where $Q$ is a finite set of states, $q_{0} \in Q$ is the initial state, and Steps: $Q \rightarrow 2^{\text {dist }(Q)}$ is the probability transition function.

[^1]For a state $q$, a probabilistic transition corresponds to first nondeterministically choosing a probability distribution $\mu \in \operatorname{Steps}(q)$ and then choosing the successor state $q^{\prime}$ probabilistically according to $\mu$.

A path is a finite or infinite non-empty sequence $\pi=q_{0} \xrightarrow{\mu_{0}} q_{1} \xrightarrow{\mu_{1}} q_{2} \ldots$ where $\mu_{i} \in \operatorname{Steps}\left(q_{i}\right)$ for every $i \geq 0$. We denote by $\pi(i)$ the $i$-th state along $\pi$ (i.e., the state $q_{i}$ ), and by $\pi^{i}$ the prefix of $\pi$ ending at $\pi(i)$ (if it exists). For a finite path $\pi$, we denote by last $(\pi)$ the last state of $\pi$. A scheduler is a function that maps every finite path $\pi$ of $\mathcal{M}$ to a distribution of $\operatorname{Steps}(\operatorname{last}(\pi))$.

For a given scheduler $S$, let Paths ${ }^{S}$ denote all infinite paths $\pi=q_{0} \xrightarrow{\mu_{0}}$ $q_{1} \xrightarrow{\mu_{1}} q_{2} \ldots$ starting in $s_{0}$ and satisfying $\mu_{i}=S\left(\pi^{i}\right)$ for every $i \geq 0$. We define a probability measure $\operatorname{Prob}^{S}$ on Paths ${ }^{S}$ in the usual way using cylinder sets [13].

We introduce the notion of rewards for an MDP.
Definition 8 (Reward). A reward function for an MDP is a function rew: $S \rightarrow$ $\mathbb{R}_{\geq 0}$. For a path $\pi$ and a set of states $F$, the reward until $F$ is reached is

$$
R(F, \pi):=\sum_{i=0}^{\min \{j \mid \pi(j) \in F\}} \operatorname{rew}(\pi(i))
$$

if the minimum exists, and $\infty$ otherwise. Given a scheduler $S$, the expected reward to reach a set of states $F$ is defined as

$$
E^{S}(F):=\int_{\pi \in \text { Paths }^{S}} R(F, \pi) \mathrm{d} \text { Prob }^{S}
$$

### 3.2 Syntax and Semantics of Probabilistic Workflow Nets

We introduce Probabilistic Workflow Nets with Rewards, just called Probabilistic Workflow Nets or PWNs in the rest of the paper.

Definition 9 (Probabilistic Workflow Net with Rewards). A Probabilistic Workflow Net with Rewards $(P W N)$ is a tuple $(P, T, F, i, o, w, r)$ where $(P, T, F, i, o)$ is a 1-safe confusion-free workflow net, and $w, r: T \rightarrow \mathbb{R}^{+}$are a weight function and a reward function, respectively.

Figure 2a shows a free-choice PWN. All transitions have reward 1, and so only the weights are represented. Unlabeled transitions have weight 1.

The semantics of a PWN is an MDP with a reward function. Intuitively, the states of the MDP are pairs $(M, t)$, where $M$ is a marking, and $t$ is the transition that was fired to reach $M$ (since the same marking can be reached by firing different transitions, the MDP can have states $\left(M, t_{1}\right),\left(M, t_{2}\right)$ for $\left.t_{1} \neq t_{2}\right)$. Additionally there is a distinguished initial and final states $I, O$. The transition relation Steps is independent of the transition $t$, i.e., $\operatorname{Steps}\left(\left(M, t_{1}\right)\right)=\operatorname{Steps}\left(\left(M, t_{2}\right)\right)$ for any two transitions $t_{1}, t_{2}$, and the reward of a state $(M, t)$ is the reward of the transition $t$. Figure 2b shows the MDP of the PWN of Figure 2a, representing only the states reachable from the initial state.


Fig. 2: Running example

Definition 10 (Probability distribution). Let $\mathcal{W}=(P, T, F, i, o, w, r)$ be a $P W N$, let $M$ be a 1-safe marking of $\mathcal{W}$ enabling at least one transition, and let $C$ be a conflict set enabled at $M$. The probability distribution $P_{M, C}$ over $T$ is obtained by normalizing the weights of the transitions in $C$, and assigning probability 0 to all other transitions.

Definition 11 (MDP and reward function of a PWN). Let $\mathcal{W}=(P, T, F$, $i, o, w, r)$ be a $P W N$. The $M D P M_{\mathcal{W}}=\left(Q, q_{0}\right.$, Steps $)$ of $\mathcal{W}$ is defined as follows:
$-Q=(\mathcal{M} \times T) \cup\{I, O\}$ where $\mathcal{M}$ are the 1 -safe markings of $\mathcal{W}$, and $q_{0}=I$.

- For every transition t:
- Steps $((\boldsymbol{o}, t))$ contains exactly one distribution, which assigns probability 1 to state o, and probability 0 to all other states.
- For every marking $M \neq \boldsymbol{o}$ enabling no transitions, Steps $((M, t))$ contains exactly one distribution, which assigns probability 1 to $(M, t)$, and probability 0 to all other states.
- For every marking $M$ enabling at least one transition, Steps $((M, t))$ contains a distribution $\mu_{C}$ for each conflict set $C$ of transitions enabled at $M$. The distribution $\mu_{C}$ is defined as follows. For the states $I, O: \mu_{C}(I)=0=\mu_{C}(O)$. For each state $\left(M^{\prime}, t^{\prime}\right)$ such that $t^{\prime} \in C$ and $M \xrightarrow{t^{\prime}} M^{\prime}: \mu_{C}\left(\left(M^{\prime}, t^{\prime}\right)\right)=P_{M, C}\left(t^{\prime}\right)$. For all other states $\left(M^{\prime}, t^{\prime}\right)$ : $\mu_{C}\left(\left(M^{\prime}, t^{\prime}\right)\right)=0$.
- $\operatorname{Steps}(I)=\operatorname{Steps}((\boldsymbol{i}, t))$ for any transition $t$.
- Steps $(O)=\operatorname{Steps}((\boldsymbol{o}, t))$ for any transition $t$.

The reward function $\operatorname{rew}_{\mathcal{W}}$ of $\mathcal{W}$ is defined by: $\operatorname{rew}_{\mathcal{W}}(I)=0=\operatorname{rew}_{\mathcal{W}}(O)$, and $r e w_{\mathcal{W}}((M, t))=r(t)$.

In Figure 2a, Steps $(i)$ is a singleton set that contains the probability distribution which assigns probability $\frac{2}{5}$ to the state $\left(\boldsymbol{p}_{\mathbf{1}}, t_{1}\right)$ and probability $\frac{3}{5}$ to the
state $\left(\left\{p_{2}, p_{3}\right\}, t_{2}\right)$. Steps $\left(\left(\left\{p_{2}, p_{3}\right\}, t_{2}\right)\right)$ contains two probability distributions, one that assigns probability 1 to $\left(\left\{p_{5}, p_{3}\right\}, t_{4}\right)$ and one that assigns probability 1 to $\left(\left\{p_{2}, p_{6}\right\}, t_{4}\right)$.

We establish a correspondence between firing sequences and paths of the MDP.

Definition 12. Let $\mathcal{W}$ be a $P W N$, and let $M_{\mathcal{W}}$ be its associated MDP. Let $\sigma=t_{1} t_{2} \ldots t_{n}$ be a firing sequence of $\mathcal{W}$. The path $\Pi(\sigma)$ of $M_{\mathcal{W}}$ corresponding to $\sigma$ is $\pi_{\sigma}=I \xrightarrow{\mu_{0}}\left(M_{1}, t_{1}\right) \xrightarrow{\mu_{1}}\left(M_{2}, t_{2}\right) \xrightarrow{\mu_{2}} \ldots$, where $M_{0}=\boldsymbol{i}$ and for every $1 \leq k$ :

- $M_{k}$ is the marking reached by firing $t_{1} \ldots t_{k}$ from $\boldsymbol{i}$, and
- $\mu_{k}$ is the unique distribution of Steps $\left(M_{k-1}, t_{k-1}\right)$ such that $\mu\left(t_{k}\right)>0$.

Let $\pi=I \xrightarrow{\mu_{0}}\left(M_{1}, t_{1}\right) \cdots\left(M_{n}, t_{n}\right)$ be a path of $M_{\mathcal{W}}$. The sequence $\Sigma(\pi)$ corresponding to $\pi$ is $\sigma_{\pi}=t_{1} \ldots t_{n}$.

It follows immediately from the definition of $M_{\mathcal{W}}$ that the functions $\Pi$ and $\Sigma$ are inverses of each other. For a path $\pi$ of the MDP that ends in state last $(\pi)$, the distributions in Steps (last $(\pi)$ ) are obtained from the conflict sets enabled after $\Sigma(\pi)$ has fired, if any. If no conflict set is enabled the choice is always trivial by construction. Therefore, a scheduler of the MDP $\mathcal{M}_{W}$ can be equivalently defined as a function that assigns to each firing sequence $\sigma \in T^{*}$ one of the conflict sets enabled after $\sigma$ has fired. In our example, after $t_{2}$ fires, the conflict sets $\left\{t_{3}\right\}$ and $\left\{t_{4}\right\}$ are concurrently enabled. A scheduler chooses either $\left\{t_{3}\right\}$ or $\left\{t_{4}\right\}$. A possible scheduler always chooses $\left\{t_{3}\right\}$ every time the marking $\left\{p_{2}, p_{3}\right\}$ is reached, and produces sequences in which $t_{3}$ always occurs before $t_{4}$, while others may behave differently.
Convention: In the rest of the paper we define schedulers as functions from firing sequences to conflict sets.

In particular, this definition allows us to define the probabilistic language of a scheduler as the function that assigns to each finite firing sequence $\sigma$ the probability of the cylinder of all paths that "follow" $\sigma$. Formally:
Definition 13 (Probabilistic language of a scheduler [20]). The probabilistic language $\nu_{S}$ of a scheduler $S$ is the function $\nu_{S}: T^{*} \rightarrow \mathbb{R}^{+}$defined by $\nu_{S}(\sigma)=\operatorname{Prob}^{S}\left(\right.$ cyl $\left.{ }^{S}(\Pi(\sigma))\right)$. A transition sequence $\sigma$ is produced by $S$ if $\nu_{S}(\sigma)>0$.

The reward function $r$ extends to transition sequences in the natural way by taking the sum of all rewards. When we draw a PWN, the labels of transitions have the form $(w, c)$ where $w$ is the weight and $c$ is the reward of the transition. See for example Figure 4a.

We now introduce the expected reward of a PWN under a scheduler.
Definition 14 (Expected reward of a PWN under a scheduler). Let $\mathcal{W}$ be a PWN, and let $S$ be a scheduler of its $M D P M_{\mathcal{W}}$. The expected reward $V^{S}(\mathcal{W})$ of $\mathcal{W}$ under $S$ is the expected reward $E^{S}(O)$ to reach the final state $O$ of $M_{\mathcal{W}}$.

Given a firing sequence $\sigma$, we have $r(\sigma)=R(O, \Pi(\sigma))$ by the definition of the reward function and the fact that $O$ can only occur at the very end of $\pi_{\sigma}$.

Lemma 2. Let $\mathcal{W}$ be a sound $P W N$, and let $S$ be a scheduler. Then $V^{S}(\mathcal{W})$ is finite and $V^{S}(\mathcal{W})=\sum_{\pi \in \Pi} R(O, \pi) \cdot \operatorname{Prob}^{S}\left(\operatorname{cyl}^{S}(\pi)\right)=\sum_{\sigma \in \text { Fin }_{\mathcal{W}}} r(\sigma) \cdot \nu_{S}(\sigma)$, where $\Pi_{O}$ are the paths of the MDP $M_{\mathcal{W}}$ leading from the initial state $I$ to the state $O$ (without looping in $O$ ).
Proof. By definition, $V^{S}(\mathcal{W})=E^{S}(O)=\int_{\pi \in \text { Paths }^{S}} R(O, \pi) \mathrm{d}$ Prob $^{S}$. Since $\mathcal{W}$ is sound, the final marking is reachable from every marking. Furthermore, since the weights are all positive, and the marking graph is finite, the probability to reach the final marking from any given marking can be bounded away from zero. Therefore the probability to eventually reach the final marking is equal to one, and so $O$ is the only absorbing state of the Markov chain induced by the scheduler $S$. It thus holds that

$$
\int_{\pi \in \text { Paths }^{S}} R(O, \pi) \mathrm{d} \operatorname{Prob}^{S}=\int_{\pi \in c y l^{S}\left(\Pi_{o}\right)} R(O, \pi) \mathrm{d} \operatorname{Prob}^{S} .
$$

Furthermore, for a path $\pi \in \Pi_{O}$, it holds that $R(O, \pi)=R\left(O, \pi^{\prime}\right)$ for all $\pi^{\prime} \in$ $\operatorname{cyl}^{S}(\pi)$ because last $(\pi)=O$. We obtain

$$
\int_{\pi \in c y l^{S}\left(\Pi_{O}\right)} R(O, \pi) \mathrm{d} \operatorname{Prob}^{S}=\sum_{\pi \in \Pi_{O}} R(O, \pi) \cdot \operatorname{Prob}^{S}\left(c y l^{S}(\pi)\right)
$$

and therefore the first equality. Together with $r(\sigma)=R(O, \Pi(\sigma))$, the fact that $\Pi$ is a bijection between $\Pi_{O}$ and $F i n_{\mathcal{W}}$, and the definition of $\nu_{S}$, the second equality follows.

### 3.3 Expected Reward of a PWN

Using a result by Varacca and Nielsen [20], we prove that the expected reward of a PWN is the same for all schedulers, which allows us to speak of "the" expected reward of a PWN. We first define partial schedulers.

Definition 15 (Partial schedulers). A partial scheduler of length $n$ is the restriction of a scheduler to firing sequences of length less than n. Given two partial schedulers $S_{1}, S_{2}$ of lengths $n_{S_{1}}, n_{S_{2}}$, we say that $S_{1}$ extends $S_{2}$ if $n_{S_{1}} \geq$ $n_{S_{2}}$ and $S_{2}$ is the restriction of $S_{1}$ to firing sequences of length less than $n_{S_{2}}$. The probabilistic language $\nu_{S}$ of a partial scheduler $S$ of length $n$ is the function $\nu_{S}: T^{\leq n} \rightarrow \mathbb{R}^{+}$defined by $\nu_{S}(\sigma)=\operatorname{Prob}^{S}\left(c y l^{S}(\Pi(\sigma))\right)$. A transition sequence $\sigma$ is produced by $S$ if $\nu_{S}(\sigma)>0$.

Observe that if $\sigma$ is not a firing sequence, then $\nu_{S}(\sigma)=0$ for every scheduler $S$. In our running example there are exactly two partial schedulers $S_{1}, S_{2}$ of length 2 ; after $t_{2}$ they choose $t_{3}$ or $t_{4}$, respectively:

$$
\begin{array}{lll}
S_{1}: \epsilon \mapsto\left\{t_{1}, t_{2}\right\} & t_{1} \mapsto\left\{t_{6}\right\} & t_{2} \mapsto\left\{t_{3}\right\} \\
S_{2}: \epsilon \mapsto\left\{t_{1}, t_{2}\right\} & t_{1} \mapsto\left\{t_{6}\right\} & t_{2} \mapsto\left\{t_{4}\right\}
\end{array}
$$

For example we have $\nu_{S_{1}}\left(t_{2} t_{3}\right)=3 / 5$, and $\nu_{S_{2}}\left(t_{2} t_{3}\right)=0$.
For finite transition sequences, Mazurkiewicz equivalence, denoted by $\equiv$, is the smallest congruence such that $\sigma t_{1} t_{2} \sigma^{\prime} \equiv \sigma t_{2} t_{1} \sigma^{\prime}$ for every $\sigma, \sigma^{\prime} \in T^{*}$ and for any two independent transitions $t_{1}, t_{2}$ [16]. We extend Mazurkiewicz equivalence to partial schedulers.

Definition 16 (Mazurkiewicz equivalence of partial schedulers). Given a partial scheduler $S$ of length $n$, we denote by $F_{S}$ the set of firing sequences $\sigma$ of $\mathcal{W}$ produced by $S$ such that either $|\sigma|=n$ or $\sigma$ leads to a marking that enables no transitions.

Two partial schedulers $S_{1}, S_{2}$ with probabilistic languages $\nu_{S_{1}}$ and $\nu_{S_{2}}$ are Mazurkiewicz equivalent, denoted $S_{1} \equiv S_{2}$, if they have the same length and there is a bijection $\phi: F_{S_{1}} \rightarrow F_{S_{2}}$ such that $\sigma \equiv \phi(\sigma)$ and $\nu_{S_{1}}(\sigma)=\nu_{S_{2}}(\phi(\sigma))$ for every $\sigma \in F_{n}$.

The two partial schedulers of our running example are not Mazurkiewicz equivalent. Indeed, we have $F_{S_{1}}=\left\{t_{1} t_{6}, t_{2} t_{3}\right\}$ and $F_{S_{2}}=\left\{t_{1} t_{6}, t_{2} t_{4}\right\}$, and no bijection satisfies $\sigma \equiv \phi(\sigma)$ for every $\sigma \in F_{S_{1}}$.

We can now present the main result of [20], in our terminology and for PWNs. ${ }^{4}$

Theorem 1 (Equivalent extension of schedulers [20] ${ }^{5}$ ). Let $S_{1}, S_{2}$ be two partial schedulers. There exist two partial schedulers $S_{1}^{\prime}, S_{2}^{\prime}$ such that $S_{1}^{\prime}$ extends $S_{1}, S_{2}^{\prime}$ extends $S_{2}$ and $S_{1}^{\prime} \equiv S_{2}^{\prime}$.

In our example, $S_{1}$ can be extended to $S_{1}^{\prime}$ by adding $t_{1} t_{6} \mapsto \emptyset$ and $t_{2} t_{3} \mapsto t_{4}$, and $S_{2}$ to $S_{2}^{\prime}$ by adding $t_{1} t_{6} \mapsto \emptyset$ and $t_{2} t_{4} \mapsto t_{3}$. Now we have $F_{S_{1}^{\prime}}=\left\{t_{1} t_{6}, t_{2} t_{3} t_{4}\right\}$ and $F_{S_{2}^{\prime}}=\left\{t_{1} t_{6}, t_{2} t_{4} t_{3}\right\}$. The obvious bijection shows $S_{1}^{\prime} \equiv S_{2}^{\prime}$, because we have $t_{2} t_{3} t_{4} \equiv t_{2} t_{4} t_{3}$ and $\nu_{S_{1}^{\prime}}\left(t_{2} t_{3} t_{4}\right)=3 / 5=\nu_{S_{2}}\left(t_{2} t_{4} t_{3}\right)$.

We now prove that the expected reward of a PWN is independent of the scheduler. We need a preliminary proposition, which follows immediately from the definition of Mazurkiewicz equivalence and the commutativity of addition.

Proposition 2. Let $\mathcal{W}$ be a $P W N$. Then for any two firing sequences $\sigma$ and $\tau$ that are Mazurkiewicz equivalent, it holds that $r(\sigma)=r(\tau)$.

Theorem 2. Let $\mathcal{W}$ be a $P W N$. There exists a value $v$ such that for every scheduler $S$ of $M_{\mathcal{W}}$, the expected reward $V^{S}(\mathcal{W})$ is equal to $v$.

Proof. Pick any two schedulers $R, S$. We show that there is a bijection between Mazurkiewicz equivalent firing sequences that end in the final marking and that are produced by those schedulers.

[^2]By Theorem 1, any two partial schedulers can be extended to two equivalent partial schedulers, in particular the partial schedulers $R^{k}, S^{k}$ that are the restrictions of $R$ and $S$ to firing sequences of length less than $k$.

Let $R^{\prime}$ be a partial scheduler extending $R^{k}, S^{\prime}$ a partial scheduler extending $S^{k}$ such that $R^{\prime} \equiv S^{\prime}$. Let $\sigma$ be a firing sequence of length $k$ produced by $R$ that ends in the final marking. By the definition of equivalence, there is a firing sequence $\tau$ such that $\sigma \equiv \tau$ and $\nu_{R^{\prime}}(\sigma)=\nu_{S^{\prime}}(\tau)$. Since $\sigma$ and $\tau$ are Mazurkiewicz equivalent, $\tau$ also ends in the final marking and also has length $k$. Since $\sigma$ is of length $k$, it was already produced by $R^{k}$ and thus by $R$, and $\tau$ was already produced by $S$.

Repeating this for every $k$, we can construct a bijection $\phi$ that maps every firing sequence $\sigma$ produced by $R$ that ends in the final marking to a Mazurkiewicz equivalent firing sequence $\phi(\sigma)$ of the same length produced by $S$ that ends in the final marking.

Using Proposition 2, we know that $r(\sigma)=r(\phi(\sigma))$. Now we apply Lemma 2 and get:
$V^{R}(\mathcal{W})=\sum_{\sigma \in \Sigma} r(\sigma) \cdot \nu_{R}(\sigma)=\sum_{\sigma \in \Sigma} r(\phi(\sigma)) \cdot \nu_{S}(\phi(\sigma))=\sum_{\sigma \in \Sigma} r(\sigma) \cdot \nu_{S}(\sigma)=V^{S}(\mathcal{W})$
where the third equality is just a reordering of the sum.

### 3.4 Free-choice PWNs

By Proposition 1, in free-choice PWNs the conflict set of a given transition is exactly its cluster, and so its probability is always the same at any reachable marking that enables it. So we can label a transition directly with this probability.

Convention: From now on we assume that the weights are normalized for each cluster, i.e. the weights are already a probability distribution.

In the next section we present a reduction algorithm that decides if a given free-choice PWN is sound or not, and if sound computes its expected reward. If the PWN is unsound, then we just apply the following lemma:

Lemma 3. The expected reward of an unsound free-choice PWN is infinite.
Proof. Let $\mathcal{W}=(P, T, F, i, o)$ be an unsound free-choice PWN. Since, by the definition of a workflow net, the graph $(P \cup T, F \cup(o, i))$ is strongly connected, if we add a transition to $\mathcal{W}$ with $o$ as input and $i$ as output transition, we obtain a strongly connected and 1 -safe free-choice net $N$. Since $\mathcal{W}$ is unsound, by Theorem 1 of [1] the net $N$ with the marking $M_{0}$ that puts one token in place $i$ is either non-live or non-bounded, and so, since $\mathcal{W}$ is 1 -safe, it must be non-live. By Theorem 4.31 of [6], the net $N$ with $M_{0}$ as initial marking has a deadlock $M$, which clearly is also a deadlock of $\mathcal{W}$. Let $i \xrightarrow{\sigma} M$ be an occurrence sequence leading to $M$. Choose a scheduler $S$ such that $\nu_{S}(\sigma)>0$. We show
that the expected reward $V^{S}(\mathcal{W})$ is infinite which, by Theorem 2, implies that the expected reward is also infinite.

The cylinder of paths of $M_{\mathcal{W}}$ that extend the path $\pi_{\sigma}$ has positive probability and infinite reward (by the definition of $\mathcal{M}_{\mathcal{W}}$ this is the cylinder of paths that extend $\pi_{\sigma}$ by staying in the state $(M, t)$ forever, where $t$ is the last transition of $\sigma$ (or in state $i$ forever, if $\sigma=\epsilon$ ). So the expected reward $V^{S}(\mathcal{W})$ is also infinite.


Fig. 3: An unsound confusion free PWN

Notice that the above lemma is not true for confusion-free workflow nets, as can be seen in the example net in Figure 3. The transition $t_{3}$ can never be enabled and thus the net is unsound. However the net contains no deadlock and indeed the only maximal transition sequence is $t_{1} t_{2}$. Thus the value of the net is finite.

## 4 Reduction rules

We transform the reduction rules of [9] for non-probabilistic (colored) workflow nets into rules for probabilistic workflow nets.

Definition 17 (Rules, correctness, and completeness). $A$ rule $R$ is a binary relation on the set of $P W N s$. We write $\mathcal{W}_{1} \xrightarrow{R} \mathcal{W}_{2}$ for $\left(\mathcal{W}_{1}, \mathcal{W}_{2}\right) \in R$.

A rule $R$ is correct if $\mathcal{W}_{1} \xrightarrow{R} \mathcal{W}_{2}$ implies that $\mathcal{W}_{1}$ and $\mathcal{W}_{2}$ are either both sound or both unsound, and have the same expected reward.
$A$ set $\mathcal{R}$ of rules is complete for a class of PWNs if for every sound $P W N$ $\mathcal{W}$ in that class there exists a sequence $\mathcal{W} \xrightarrow{R_{1}} \mathcal{W}_{1} \cdots \mathcal{W}_{n-1} \xrightarrow{R_{n}} \mathcal{W}^{\prime}$ such that $\mathcal{W}^{\prime}$ is a $P W N$ consisting of a single transition $t$ between the two only places $i$ and $o$.

Observe that if $\mathcal{W}$ is reduced to a $\mathcal{W}^{\prime}$ as above, then the expected reward of $\mathcal{W}$ is equal to the reward of $t$ in $\mathcal{W}^{\prime}$.

As in [9], we describe rules as pairs of a guard and an action. $\mathcal{W}_{1} \xrightarrow{R} \mathcal{W}_{2}$ holds if $\mathcal{W}_{1}$ satisfies the guard, and $\mathcal{W}_{2}$ is a possible result of applying the action to $\mathcal{W}_{1}$.

Merge rule. The merge rule merges two transitions with the same input and output places into one single transition. The weight of the new transition is the sum of the old weights, and the reward is the weighted average of the reward of the two merged transitions.

Definition 18. Merge rule
Guard: $\mathcal{W}$ contains two distinct transitions $t_{1}, t_{2} \in T$ such that ${ }^{\bullet} t_{1}={ }^{\bullet} t_{2}$ and $t_{1}^{\bullet}=t_{2}^{\bullet}$.
Action: (1) $T:=\left(T \backslash\left\{t_{1}, t_{2}\right\}\right) \cup\left\{t_{m}\right\}$, where $t_{m}$ is a fresh name.
(2) $t_{m}^{\bullet}:=t_{1}^{\bullet}$ and ${ }^{\bullet} t_{m}:={ }^{\bullet} t_{1}$.
(3) $r\left(t_{m}\right):=w\left(t_{1}\right) \cdot r\left(t_{1}\right)+w\left(t_{2}\right) \cdot r\left(t_{2}\right)$.
(4) $w\left(t_{m}\right)=w\left(t_{1}\right)+w\left(t_{2}\right)$.

Iteration rule. Loosely speaking, the iteration rule removes arbitrary iterations of a transition by adjusting the weights of the possible successor transitions. The probabilities are normalized again and the reward of each successor transition increases by a geometric series dependent on the reward and weight of the removed transition.

Definition 19. Iteration rule
Guard: $\mathcal{W}$ contains a cluster $c$ with a transition $t \in c$ such that $t^{\bullet}={ }^{\bullet} t$.
Action: (1) $T:=(T \backslash\{t\})$.
(2) For all $t^{\prime} \in c \backslash\{t\}: r\left(t^{\prime}\right):=\frac{w(t)}{1-w(t)} \cdot r(t)+r\left(t^{\prime}\right)$
(3) For all $t^{\prime} \in c \backslash\{t\}: w\left(t^{\prime}\right):=\frac{w\left(t^{\prime}\right)}{1-w(t)}$

Observe that $\frac{w(t)}{1-w(t)} \cdot r(t)=(1-w(t)) \cdot \sum_{i=0}^{\infty} w(t)^{i} \cdot i \cdot r(t)$ captures the fact that $t$ can be executed arbitrarily often, each execution yields the reward $r(t)$, and eventually some other transition occurs.

For an example of an application of the iteration rule, consult Figure 4b and Figure 4c. Transition $t_{9}$ has been removed and as a result the label of transition $t_{7}$ changed.

Shortcut rule. The shortcut rule merges transitions of two clusters into one single transition with the same effect. The reward of the new transition is the sum of the rewards of the old transitions, and its weight the product of the old weights.

A transition $t$ unconditionally enables a cluster $c$ if ${ }^{\bullet} t^{\prime} \subseteq t^{\bullet}$ for some transition $t^{\prime} \in c$. Observe that if $t$ unconditionally enables $c$ then any marking reached by firing $t$ enables every transition in $c$.

## Definition 20. Shortcut rule

Guard: $\mathcal{W}$ contains a transition $t$ and a cluster $c \neq[t]$ such that $t$ unconditionally enables $c$.
Action: (1) $T:=(T \backslash\{t\}) \cup\left\{t_{s}^{\prime} \mid t^{\prime} \in c\right\}$, where $t_{s}^{\prime}$ are fresh names.
(2) For all $t^{\prime} \in c: \bullet t_{s}^{\prime}:={ }^{\bullet} t$ and $t_{s}^{\prime}:=\left(t^{\bullet} \backslash \bullet t^{\prime}\right) \cup t^{\bullet}$.
(3) For all $t^{\prime} \in c: r\left(t_{s}^{\prime}\right):=r(t)+r\left(t^{\prime}\right)$.
(4) For all $t^{\prime} \in c: w\left(t_{s}^{\prime}\right)=w(t) \cdot w\left(t^{\prime}\right)$.
(5) If ${ }^{\bullet} p=\emptyset$ for all $p \in c$, then remove $c$ from $\mathcal{W}$.

For an example shortcut rule application, compare the example of Figure 2a with the net in Figure 4a. The transition $t_{1}$ which unconditionally enabled the cluster $\left[t_{6}\right]$ has been shortcut, a new transition $t_{8}$ has been created, and $t_{1}, p_{1}$ and $t_{6}$ have been removed.

Theorem 3. The merge, shortcut and iteration rules are correct for PWNs.
Proof. It was already shown in [9] that the rules preserve soundness for freechoice workflow nets. We thus only have to show that the rules preserve the expected reward of the net. In the unsound case this is easy: Since there is a reachable marking from which the final marking is unreachable, there is a cylinder which occurs with positive probability and never reaches the final marking. For such a cylinder, the reward is infinite by Definition 8, thus the expected reward is infinite. As the rules preserve unsoundness, they also preserve the expected reward in that case.

By Theorem 2 the expected reward of the net does not depend on the scheduler. We use this fact in the following way: For each rule, we pick two schedulers, one for the net before the rule application and one for the net after the rule was applied. These schedulers will be such that it is easy to show that their expected rewards are equal. We begin with the shortcut rule.

Shortcut rule. Let $\mathcal{W}_{1}, \mathcal{W}_{2}$ be such that $\mathcal{W}_{1} \xrightarrow{\text { shortcut }} \mathcal{W}_{2}$. Let $c, t$ be as in Definition 20. Let $S_{1}$ be a scheduler for $W_{1}$ such that $S_{1}\left(\sigma_{1}\right)=c$ if $\sigma_{1}$ ends with $t$. Since $t$ unconditionally enables $c$, this is a valid scheduler.

We define a mapping $\phi$ that maps firing sequences in $\mathcal{W}_{2}$ to firing sequences in $\mathcal{W}_{1}$ by replacing every occurrence of $t_{s}^{\prime}$ by $t t^{\prime}$. Next we define a scheduler $S_{2}$ for $\mathcal{W}_{2}$ by $S_{2}\left(\sigma_{2}\right)=S_{1}\left(\phi\left(\sigma_{2}\right)\right)$.

Observe that $\phi$ is a bijection between sequences produced by $S_{1}$ that do not end with $t$ and sequences produced by $S_{2}$. In particular $\phi$ is a bijection between sequences produced by $S_{1}$ and $S_{2}$ that end with the final marking.

Let now $\sigma_{2}$ be a firing sequence in $\mathcal{W}_{2}$ and let $\sigma_{1}=\phi\left(\sigma_{2}\right)$. We claim that $\sigma_{1}$ and $\sigma_{2}$ have the same reward and also $\nu_{S_{1}}\left(\sigma_{1}\right)=\nu_{S_{2}}\left(\sigma_{2}\right)$. Indeed, since the only difference is that every occurrence of $t_{s}^{\prime}$ is replaced by $t t^{\prime}$ and $r\left(t_{S}^{\prime}\right)=r(t)+r\left(t^{\prime}\right)$ and $w\left(t_{s}^{\prime}\right)=w(t) w\left(t^{\prime}\right)$ by the definition of the shortcut rule, the reward must be equal and $\nu_{S_{1}}\left(\sigma_{1}\right)=\nu_{S_{2}}\left(\sigma_{2}\right)$.

We now use these equalities, the fact that there is a bijection between firing sequences that end with the final marking, and Lemma 2:

$$
\begin{aligned}
V\left(\mathcal{W}_{2}\right) & =\sum_{\sigma_{2} \in \text { Fin }_{\mathcal{W}_{2}}} r\left(\sigma_{2}\right) \cdot \nu_{S_{2}}(\sigma)=\sum_{\sigma_{2} \in \text { Fin }_{\mathcal{W}_{2}}} r\left(\phi\left(\sigma_{2}\right)\right) \cdot \nu_{S_{1}}\left(\phi\left(\sigma_{2}\right)\right) \\
& =\sum_{\sigma_{1} \in \text { Fin }_{\mathcal{W}_{1}}} r\left(\sigma_{1}\right) \cdot \nu_{S_{1}}\left(\sigma_{1}\right)=V\left(\mathcal{W}_{1}\right)
\end{aligned}
$$

Iteration rule. Let $\mathcal{W}_{1}, \mathcal{W}_{2}$ be such that $\mathcal{W}_{1} \xrightarrow{\text { iteration }} \mathcal{W}_{2}$. Let $c, t$ be as in Definition 19. Let $S_{2}$ be a scheduler for $W_{2}$ such that $S_{2}\left(\sigma_{2}\right)=c$ if $c$ is enabled after $\sigma_{2}$.

We define a mapping $\phi$ that maps firing sequences in $\mathcal{W}_{1}$ to firing sequences in $\mathcal{W}_{2}$ by removing all occurrences of $t$. Next we define a scheduler $S_{1}$ for $\mathcal{W}_{1}$ by $S_{1}\left(\sigma_{1}\right)=S_{2}\left(\phi\left(\sigma_{1}\right)\right)$. Note that $\phi$ is not a bijection but it is surjective.

Let $r_{1}$ and $r_{2}$ be the reward functions of $\mathcal{W}_{1}$ and $\mathcal{W}_{2}$. For a sequence $\sigma_{2}$ in $\mathcal{W}_{2}$, we claim:

$$
r_{2}\left(\sigma_{2}\right) \cdot \nu_{S_{2}}\left(\sigma_{2}\right)=\sum_{\sigma_{1} \in \phi^{-1}\left(\sigma_{2}\right)} r_{1}\left(\sigma_{1}\right) \cdot \nu_{S_{1}}\left(\sigma_{1}\right) .
$$

Let $k$ be the number of times $c$ is enabled during $\sigma_{2}$. We only consider the case $k=1$, the general case being similar. We observe that $\sigma_{2}$ is also a sequence in $\mathcal{W}_{1}$. We have

$$
\begin{align*}
\nu_{S_{1}}\left(\sigma_{2}\right) & =\nu_{S_{2}}\left(\sigma_{2}\right) \cdot(1-w(t))  \tag{1}\\
r_{1}\left(\sigma_{2}\right) & =r_{2}\left(\sigma_{2}\right)-\frac{w(t)}{1-w(t)} \cdot c(t) \tag{2}
\end{align*}
$$

because the probabilistic choice must pick something other than $t$, and because the iteration rule adds $\frac{w(t)}{1-w(t)} \cdot c(t)$ to the reward of every transition in $c$ in $\mathcal{W}_{2}$.

We now insert $l$ occurrences of $t$ in $\sigma_{2}$, at the position at which $c$ is enabled, and call the new sequence $\tau_{l}$. We have $\phi^{-1}\left(\sigma_{2}\right)=\left\{\tau_{l} \mid l \geq 0\right\}$. Further $r_{1}\left(\tau_{l}\right)=$ $r_{1}\left(\sigma_{2}\right)+l \cdot c(t)$ and $\nu_{S_{1}}\left(\tau_{l}\right)=\nu_{S_{1}}(\tau) \cdot w(t)^{l}$, and so summing over all $l$ we get:

$$
\begin{align*}
\sum_{\sigma_{1} \in \phi^{-1}\left(\sigma_{2}\right)} r_{1}\left(\sigma_{1}\right) \cdot \nu_{S_{1}}\left(\sigma_{1}\right) & =\sum_{l=0}^{\infty} r_{1}\left(\tau_{l}\right) \cdot \nu_{S_{1}}\left(\tau_{l}\right) \\
& =\nu_{S_{1}}\left(\sigma_{2}\right) \cdot \sum_{l=0}^{\infty}\left(r_{1}\left(\sigma_{2}\right)+l \cdot c(t)\right) \cdot w(t)^{l} \\
& =\nu_{S_{1}}\left(\sigma_{2}\right) \cdot\left(\frac{r_{1}\left(\sigma_{2}\right)}{1-w(t)}+\frac{c(t) \cdot w(t)}{(1-w(t))^{2}}\right) \\
& =\nu_{S_{2}}\left(\sigma_{2}\right) \cdot\left(r_{1}\left(\sigma_{2}\right)+\frac{c(t) \cdot w(t)}{1-w(t)}\right)  \tag{by1}\\
& =\nu_{S_{2}}\left(\sigma_{2}\right) \cdot r_{2}\left(\sigma_{2}\right) \tag{by2}
\end{align*}
$$

and the claim is proved.

Now, using the claim we obtain:

$$
\begin{aligned}
V\left(\mathcal{W}_{2}\right) & =\sum_{\sigma_{2} \in \text { Fin }_{W_{2}}} r_{2}\left(\sigma_{2}\right) \cdot \nu_{S_{2}}(\sigma)=\sum_{\sigma_{2} \in \text { Fin }_{W_{2}}} \sum_{\sigma_{1} \in \phi^{-1}\left(\sigma_{2}\right)} r_{1}\left(\sigma_{1}\right) \cdot \nu_{S_{1}}\left(\sigma_{1}\right) \\
& =\sum_{\sigma_{1} \in \text { Finn }_{w_{1}}} r_{1}\left(\sigma_{1}\right) \cdot \nu_{S_{1}}\left(\sigma_{1}\right)=V\left(\mathcal{W}_{1}\right)
\end{aligned}
$$

where the third equality follows from the fact that $\phi$ is defined on all sequences of $\mathcal{W}_{1}$ and thus $\phi^{-1}$ hits every sequence in $\mathcal{W}_{1}$ exactly once.

Merge rule. Let $\mathcal{W}_{1}, \mathcal{W}_{2}$ be such that $\mathcal{W}_{1} \xrightarrow{\text { merge }} \mathcal{W}_{2}$. Let $t_{1}, t_{2}$ be as in Definition 18. Let $S_{2}$ be a scheduler for $\mathcal{W}_{2}$.

We define a mapping $\phi$ that maps firing sequences in $\mathcal{W}_{1}$ to firing sequences in $\mathcal{W}_{2}$ by replacing all occurrences of $t_{1}$ and $t_{2}$ by $t_{m}$. We define a scheduler $S_{1}$ for $\mathcal{W}_{1}$ by $S_{1}\left(\sigma_{1}\right)=S_{2}\left(\phi\left(\sigma_{1}\right)\right)$.

Once again, $\phi$ is a surjective function. For a sequence $\sigma_{2}$ in $\mathcal{W}_{2}$, we claim $r\left(\sigma_{2}\right) \cdot \nu_{S_{2}}\left(\sigma_{2}\right)=\sum_{\sigma_{1} \in \phi^{-1}\left(\sigma_{2}\right)} r\left(\sigma_{1}\right) \cdot \nu_{S_{1}}\left(\sigma_{1}\right)$. Indeed, every sequence $\sigma_{1}$ the set $\phi^{-1}\left(\sigma_{2}\right)$ can be obtained by replacing $t_{m}$ by either $t_{1}$ or $t_{2}$. So, by Definition 18, the sums are equal.

As for the iteration rule, this equality and the fact that $\phi$ is defined for every sequence in $\mathcal{W}_{1}$ imply that the expected rewards of $\mathcal{W}_{1}$ and $\mathcal{W}_{2}$ are equal.

In [9] we provide a reduction algorithm for non-probabilistic free-choice workflow, and prove the following result.

Theorem 4 (Completeness[9]). The reduction algorithm summarizes every sound free choice workflow net in at most $\mathcal{O}\left(|C|^{4} \cdot|T|\right)$ applications of the shortcut rule and $\mathcal{O}\left(|C|^{4}+|C|^{2} \cdot|T|\right)$ applications of the merge and iteration rules, where $C$ is the set of clusters of the net. Any unsound free-choice workflow nets can be recognized as unsound in the same number of rule applications.

We illustrate a complete reduction by reducing the example of Figure 2a. We set the reward for each transition to 1 , so the expected reward of the net is the expected number of transition firings until the final marking is reached. Initially, $t_{1}$ unconditionally enables $\left[t_{6}\right]$ and we apply the shortcut rule. Since $\left[t_{6}\right]=\left\{t_{6}\right\}$, exactly one new transition $t_{8}$ is created. Furthermore $t_{1}, p_{1}$ and $t_{6}$ are removed (Figure 4a). Now, $t_{5}$ unconditionally enables $\left[t_{3}\right]$ and $\left[t_{4}\right]$. We apply the shortcut rule twice and call the result $t_{9}$ (Figure 4b). Transition $t_{9}$ now satisfies the guard of the iteration rule and can be removed, changing the label of $t_{7}$ (Figure 4c). Since $t_{2}$ unconditionally enables $\left[t_{3}\right]$ and $\left[t_{4}\right]$, we apply the shortcut rule twice and call the result $t_{10}$ (Figure 4d). After short-cutting $t_{10}$, we apply the merge rule to the two remaining transitions, which yields a net with one single transition labeled by $(1,5)$ (Figure 4e). So the net terminates with probability 1 after firing 5 transitions in average.


Fig. 4: Example of reduction

Fixing a scheduler. Since the expected reward of a PWN $\mathcal{W}$ is independent of the scheduler, we can fix a scheduler $S$ and compute the expected reward $V^{S}(\mathcal{W})$. This requires to compute only the Markov chain induced by $S$, which can be much smaller than the MDP. However, it is easy to see that this idea does not lead to a polynomial algorithm. Consider the free-choice PWN of Figure 5, and the scheduler that always chooses the largest enabled cluster according to the order

$$
\left\{t_{11}, t_{12}\right\}>\cdots>\left\{t_{n 1}, t_{n 2}\right\}>\left\{u_{11}\right\}>\left\{u_{12}\right\}>\cdots>\left\{u_{n 1}\right\}>\left\{u_{n 2}\right\}
$$

Then for every subset $K \subset\{1, \ldots, n\}$ the Markov chain contains a state enabling $\left\{u_{i 1} \mid i \in K\right\} \cup\left\{u_{i 2} \mid i \notin K\right\}$, and has therefore exponential size. There might be a procedure to find a suitable scheduler for a given PWN such that the Markov chain has polynomial size, but we do not know of such a procedure.


Fig. 5: Example

## 5 Experimental evaluation

We have implemented our reduction algorithm as an extension of the algorithm described in [9]. In this section we report on its performance and on a comparison with Prism[14]. The results confirm what could be expected: our polynomial algorithm for free-choice workflows outperforms Prism's exponential, but more generally applicable algorithm. More interestingly, they provide quantitative information on the speed-up achieved by our algorithm.

Industrial benchmarks. The benchmark suite consists of 1385 free-choice workflow nets, previously studied in [10], of which 470 nets are sound. The workflows correspond to business models designed at IBM. Since they do not contain probabilistic information, we assigned to each transition $t$ the probability $\frac{1}{\mid t t] \mid}$ (i.e., the probability is distributed uniformly among the transitions of a cluster). We study the following questions, which can be answered by both our algorithm and Prism: Is the probability to reach the final marking equal to one (equivalent to "is the net sound?"). And if so, how many transitions must be fired in average to reach the final marking? (This corresponds to a reward function assigning reward 1 to each transition.)

All experiments were carried out on an i7-3820 CPU using 1 GB of memory.
Prism has three different analysis engines able to compute expected rewards: explicit, sparse and symbolic (bdd). In a preliminary experiment with a timeout of 30 seconds, we observed that the explicit engine clearly outperforms the other two: It solved 1309 cases, while the bdd and sparse engines only solved 636 and 638 cases, respectively. Moreover, 418 and 423 of the unsolved cases were due to memory overflow, so even with a larger timeout the explicit engine is still leading. For this reason, in the comparison we only used the explicit engine.

After increasing the timeout to 10 minutes, the explicit engine did not solve any further case, leaving 76 cases unsolved. This was due to the large state space of the nets: 69 out of the 76 have over $10^{6}$ reachable states.

The 1309 cases were solved by the explicit engine in 353 seconds, with about 10 seconds for the larger nets. Our implementation solved all 1385 cases in 5 seconds combined. It never needs more than 20 ms for a single net, even for those with more than $10^{7}$ states (for these nets we do not know the exact number of reachable states).


Fig. 6: Academic benchmark

In the unsound case, our implementation still reduces the reachable state space by a lot, which makes it easier to apply state exploration tools for other problems than the expected reward, like the distribution of the rewards. After reduction, the 69 nets with at least $10^{6}$ states had an average of 5950 states, with the largest at 313443 reachable states.

An academic benchmark. Many workflows in our suite have a large state space because of fragments modeling the following situation. Multiple processes do a computation step in parallel, after which they synchronize. Process $i$ may execute its step normally with probability $p_{i}$, or a failure may occur with probability $1-p_{i}$, which requires to take a recovery action and therefore has a higher cost. Such a scenario is modeled by the free-choice PWNs net of Figure 6a, where the probabilities and costs are chosen at random. The scenario can also be easily modeled in Prism. Figure 6b shows the time needed by the three Prism engines and by our implementation for computing the expected reward using a time limit of 10 minutes. The number of reachable states grows exponentially in the number processes, and the explicit engine runs out of memory for 15 processes. Since the failure probabilities vary between the processes, there is little structure that the symbolic engine can exploit, and it times out for 13 processes. The sparse engine reaches the time limit at 20 processes. However, since the rule-based approach does not need to construct the state space, we can easily solve the problem with up to 500 processes.

## 6 Conclusion

We have presented a set of reduction rules for probabilistic workflow nets with rewards that preserve soundness and the expected reward of the net, and are
complete for free-choice nets. While the semantics and the expected reward are defined via an associated Markov Decision Process, our rules work directly on the workflow net. The rules lead to the first polynomial-time algorithm to compute the expected reward.

In future work we want to generalize our algorithm in several ways. First, we think that the cost model can be extended to any semiring satisfying some mild conditions. A particular instance of this result should lead to an algorithm for computing the probability on non-termination and the conditional expected reward under termination, which is of interest in the unsound case. Second, we plan to extend our approach to GSPNs with the semantics introduced in [7]. Third, we think that the expected time to termination of a free-choice workflow can also be computed by means of a reduction algorithm.

Acknowledgments. We thank the anonymous referees for their comments, and especially the one who helped us correct a mistake in Lemmas 2 and 3.

## References

1. W. v. d. Aalst. The application of Petri nets to workflow management. Journal of Circuits, Systems, and Computers, 8(1):21-66, 1998.
2. W. v. d. Aalst and K. M. v. Hee. Workflow management: models, methods, and systems. MIT press, 2004.
3. S. Abbes and A. Benveniste. True-concurrency probabilistic models: Markov nets and a law of large numbers. Theor. Comput. Sci., 390(2-3):129-170, 2008.
4. S. Abbes and A. Benveniste. Concurrency, $\sigma$-algebras, and probabilistic fairness. In Proceedings of FOSSACS 2009, LNCS, vol. 5504, pages 380-394, 2009.
5. J. Desel and T. Erwin. Modeling, simulation and analysis of business processes. In Business Process Management, LNCS, vol. 1806, pages 129-141. Springer, 2000.
6. J. Desel and J. Esparza. Free choice Petri nets, volume 40. Cambridge university press, 2005.
7. C. Eisentraut, H. Hermanns, J.-P. Katoen, and L. Zhang. A semantics for every GSPN. In Application and Theory of Petri Nets and Concurrency, pages 90-109. Springer, 2013.
8. J. Esparza. Decidability and complexity of Petri net problems - An introduction. In Lectures on Petri Nets I: Basic Models, Advances in Petri Nets, LNCS, vol. 1491, pages 374-428, 1996.
9. J. Esparza and P. Hoffmann. Reduction rules for colored workflow nets. In Proceedings of FASE 2016, LNCS, vol. 9633, pages 342-358, 2016.
10. D. Fahland, C. Favre, B. Jobstmann, J. Koehler, N. Lohmann, H. Völzer, and K. Wolf. Instantaneous soundness checking of industrial business process models. In Business Process Management, LNCS, vol. 5701, pages 278-293. Springer, 2009.
11. C. Favre, D. Fahland, and H. Völzer. The relationship between workflow graphs and free-choice workflow nets. Inf. Syst., 47:197-219, 2015.
12. C. Favre, H. Völzer, and P. Müller. Diagnostic information for control-flow analysis of Workflow Graphs (a.k.a. Free-Choice Workflow nets). In Proceedings of TACAS 2016, LNCS, vol. 9636, pages 463-479, 2016.
13. J. G. Kemeny, J. L. Snell, and A. W. Knapp. Denumerable Markov chains: with a chapter of Markov random fields by David Griffeath, volume 40. Springer Science \& Business Media, 2012.
14. M. Kwiatkowska, G. Norman, and D. Parker. PRISM 4.0: Verification of probabilistic real-time systems. In Proceedings of CAV 2011, LNCS, vol. 6806, pages 585-591, 2011.
15. M. Magnani and D. Montesi. BPMN: How much does it cost? An incremental approach. In Business Process Management, LNCS, vol. 4714, pages 80-87. Springer, 2007.
16. A. Mazurkiewicz. Trace theory. In Petri nets: applications and relationships to other models of concurrency, pages 278-324. Springer, 1986.
17. M. L. Puterman. Markov decision processes: discrete stochastic dynamic programming. John Wiley \& Sons, 2014.
18. K. Saeedi, L. Zhao, and P. R. F. Sampaio. Extending BPMN for supporting customer-facing service quality requirements. In ICWS 2010, pages 616-623. IEEE Computer Society, 2010.
19. P. Sampath and M. Wirsing. Evaluation of cost based best practices in business processes. In Enterprise, Business-Process and Information Systems Modeling, Lecture Notes in Business Information Processing, vol. 81, pages 61-74. Springer, 2011.
20. D. Varacca and M. Nielsen. Probabilistic Petri nets and Mazurkiewicz equivalence. 2003. Unpublished Manuscript. Available online at http://www.lacl.fr/ ~dvaracca/works.html. Last retrieved on May 27, 2016.
21. D. Varacca, H. Völzer, and G. Winskel. Probabilistic event structures and domains. Theor. Comput. Sci., 358(2-3):173-199, 2006.

[^0]:    * This work was partially funded by the DFG Project 5090812 (Negotiations: Ein Modell für nebenläufige Systeme mit niedriger Komplexität).

[^1]:    ${ }^{3}$ In [6] clusters are defined in a slightly different way.

[^2]:    ${ }^{4}$ In [20], enabled conflict sets are called actions, and markings are called cases.
    ${ }^{5}$ Stated as Theorem 2, the original paper gives this theorem with $S_{1}^{\prime}$ and $S_{2}^{\prime}$ being (non-partial) schedulers. However, in the paper equivalence is only defined for partial schedulers and the schedulers constructed in the proof are also partial.

