Dependability Analysis of a Class of Probabilistic Petri Nets

HSU-CHUN YEN* AND LIEN-PO YU
Dept. of Electrical Eng., National Taiwan University
Taipei, Taiwan 106, R.O.C.
E-mail: yen@cc.ee.ntu.edu.tw
Fax: +886-2-2363 8247
Phone: + 886-2-2363 5251 ext. 540

Abstract

Verification of various properties associated with concurrent/distributed systems is critical in the process of designing and analyzing dependable systems. While techniques for the automatic verification of finite-state systems *are relatively well studied, one of the main challenges in the domain of verification is concerned with the development of new techniques capable of coping with problems beyond the finite state framework. In this paper, we investigate a number of problems closely related to dependability analysis in the context of probabilistic infinite-state systems modelled by* probabilistic conflict-free Petri nets*. Using a* valuation method*, we are able to demonstrate effective procedures for solving the* termination with probability 1*, the* self-stabilization with probability 1*, and the* controllability with probability 1 *problems in a unified framework.*

Keywords: Controllability, probabilistic Petri net, reachability, self-stabilization, verification.

1. Introduction

As modern hardware and software systems are becoming more complex and at the same time required to be more dependable, there is an ever-increasing need for new evaluation techniques; the advantage of analytical evaluation over experimental one lies in its usefulness in abstracting the essentials of systems (so that various levels of system details can be abstracted out) and analyzing or predicting system behaviors (especially while a system is being designed or implemented), and its being generally far more cost effective than its experiment-based counterpart. With the increasing interest in developing dependable systems, the study of problems regarding *termination*, *controllability*, and *self-stabilization*, among others, has also been gaining increasing popularity in the computer science community due to the following reasons. Dependable systems are often associated with properties like *safety* ('something bad never happens'), *liveness* ('something good eventually happens'), *fault-tolerance* (error detection, recovery and masking), etc. The safety property requires that undesired or failure states be avoided at all times during the course of a computation. The liveness property asserts that a certain desired condition be true eventually. The notion of self-stabilization was introduced by Dijkstra [2] to describe a system having the behavior that regardless of its starting configuration, the computation is guaranteed to return to a *legitimate configuration* eventually. By a legitimate configuration we mean a configuration reachable from its initial configuration. Since a self-stabilizing system has the ability to 'correct' itself even in the presence of certain unpredictable errors leading itself to an illegitimate configuration, one can assert that a self-stabilizing system is, in a sense, fault-tolerant.

In view of the above, it becomes apparent that automatic verification of various properties associated with concurrent/distributed systems is critical in the process of designing and analyzing dependable systems. While techniques for the automatic verification of *finite-state systems* are relatively well studied; see, for example, Clarke, Grumberg and Long [1], one of the main challenges in the domain of verification is concerned with the development of new techniques capable of coping with problems beyond the finite state framework.

The aim of this paper focuses on investigating the following problem. Given an infinite-state system, determine whether the system meets certain criteria (termination, controllability, and self-stabilization) frequently required in dependable computing environments. Taking into consideration that many real-world

[*] Corresponding author

0-7695-2076-6/04 $20.00 © 2004 IEEE

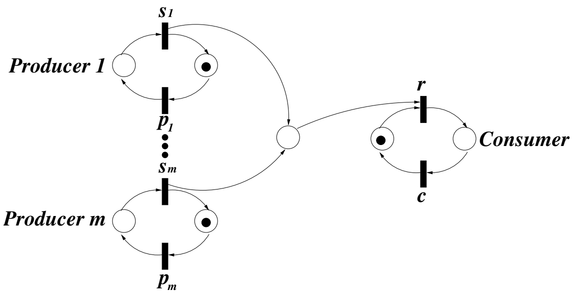

Figure 1. A conflict-free Petri net modelling a system of *m* producers and one consumer.

systems are nondeterministic (or stochastic, to be more precise) in nature, the system model under our investigation is not only infinite-state but also *probabilistic*, allowing us to ask questions such as 'something happens with probability 1', for instance. The systems under investigation are modelled as *Petri nets*, which have been regarded as one of the most successful models for describing the behaviors of systems of concurrent nature [9]. In spite of their popularity, the high expressive power of Petri nets renders most of the nontrivial problems for this model highly intractable or even unsolvable. As a result, it is of interest from the theoretical and practical viewpoints to investigate problems with respect to restricted (either structurally or behaviorally) versions of Petri nets, in hope of making simpler solutions feasible and well as gaining more insights into the factors that make general Petri nets difficult to analyze.

A Petri net is *conflict-free* if every place which is an input of more than one transition is on a self-loop with each such transition [6]; therefore, once a transition becomes enabled, the only way to disable it is to fire the transition itself. Figure 1 is a conflict-free Petri net modelling the m-producer-1-consumer problem.

Probabilistic techniques, capable of modelling unreliable or unpredictable behaviors of systems, are extensively used in the analysis of the performance and dependability of hardware and software systems, see, for example, Marsan, Balbo, etc. [8]. In this paper, we consider a *probabilistic* version of conflict-free Petri nets, in which each marking (i.e., configuration) is associated with a *transition probability function* characterizing the firing of each enabled transition. We investigate through a technique recently developed in [11] (called the *valuation method*) a number of important dependability-related problems, including *termination with probability 1*, *self-stabilization with probability 1*, and *controllability with probability 1*, etc. The idea of the valuation-based approach for Petri nets is to associate a valuation in $\{0, 1, 2, ...\infty\}$ with each marking, and if the set of markings of zero valuation is *forward-closed*, then the valuation along any computation is non-increasing, and in many cases, has the tendency to move towards the ground level (i.e., valuation zero) of which the marking sometimes constitute the set of states of interest, e.g., the termination set.

The main contribution of this paper lies in the development of a unified approach (extending of the work of [11]) for reasoning about various dependability-related problems for probabilistic conflict-free Petri nets. In addition to the results themselves, we feel that the valuation-based approach for the analysis of probabilistic conflict-free Petri net is also interesting in its own right, and may have other applications to the analysis of other probabilistic Petri net models. The remainder of this paper is organized as follows. In Section 2, we define the probabilistic version of Petri nets on which the type of systems under consideration is based. The notations and definitions used throughout this paper as well as the dependability-related problems under investigation are also explained in this section. In Section 3, we develop the basic theory behind the valuation-based approach for probabilistic conflict-free Petri nets to solve those problems defined in Section 2 in a unified framework. Finally, a conclusion and directions for future research are given in Section 4.

2. Preliminaries

2.1. Definitions and notations

Let *N* denote the set of nonnegative integers, and N^k the set of vectors of k nonnegative integers. A *Petri net* (*PN*, for short) is a 3-tuple (P,T,φ), where

- P is a finite set of *places*,
- T is a finite set of *transitions*, and
- φ is a *flow function* φ : (P × T) ∪ (T × P) → {0,1}.

In this paper, k is reserved for |P| (the number of places in P). A *marking* is a mapping μ : P → *N*. (μ assigns *tokens* to each place of the net.) Pictorially, Petri net is a directed, bipartite graph consisting of two kinds of nodes: *places* (represented by circles within which each small black dot denotes a *token*) and *transitions* (represented by bars or boxes), where each arc is either from a place to a transition or vice versa. See Figure 1.

A transition $t \in$ T is *enabled* at a marking μ iff for every $p \in$ P, $\varphi(p, t) \leq \mu(p)$. In a PN (P,T,$\varphi$), a transition t may *fire* at a marking μ if t is enabled at μ; we then write $\mu \overset{t}{\longmapsto} \mu'$, where μ'(p) = μ(p) – φ(p,t) + φ(t,p) for all p ∈ P. (We also write $\mu \rightarrow \mu'$ to denote the reachability of μ' from μ in one step.) A sequence of transitions σ = t$_1$...t$_n$ is a *firing sequence* from μ_0 in a PN iff $\mu_0 \xmapsto{t_1} \mu_1 \xmapsto{t_2} \cdots \xmapsto{t_n} \mu_n$, for some sequence of markings μ_1,...,μ_n; we also write $\mu_0 \xmapsto{\sigma} \mu_n$. In a PN, we write $\mu_0 \xmapsto{\sigma}$ to denote that σ is enabled and can be fired from μ_0, i.e., $\mu_0 \xmapsto{\sigma}$ iff there exists a marking μ such that $\mu_0 \xmapsto{\sigma} \mu$. An infinite sequence σ is a firing sequence from μ, written as $\mu \xmapsto{\sigma}$, iff for every finite prefix σ' of σ, $\mu \xmapsto{\sigma'}$. We write $\mu \xmapsto{*} \mu'$ to denote the existence of a firing sequence σ such that $\mu \xmapsto{\sigma} \mu'$.

The *reachability set* of a PN $\mathcal{P}$ with respect to initial marking μ_0 is the set $R(\mathcal{P}, \mu_0)$ = $\{\mu \mid \mu_0 \xmapsto{\sigma} \mu$ for some $\sigma \in T^*\}$. Given a set of markings S, the *successor* (resp., *predecessor*) of S, written as $succ(S)$ (resp., $pred(S)$), is the set $\{\mu \mid \exists t \in T, \mu' \in S, \mu' \xmapsto{t} \mu\}$ (resp., $\{\mu \mid \exists t \in T, \mu' \in S, \mu \xmapsto{t} \mu'\}$). Let $pred^*$ (resp., $succ^*$) be the reflexive and transitive closure of $pred$ (resp., $succ$). (That is, $succ^*(S) = \{\mu \mid \exists \mu' \in S, \mu' \xmapsto{*} \mu\}$, and $pred^*(S) = \{\mu \mid \exists \mu' \in S, \mu \xmapsto{*} \mu'\}$.) The sets $succ^*(S)$ and $pred^*(S)$ will be referred to as the *forward reachability set* and the *backward reachability set* of S, respectively. Notice that $R(\mathcal{P}, \mu_0) = succ^*(\{\mu_0\})$. A set of markings S is said to be *forward-closed* if $\forall \mu \in S, \forall t \in T, \mu \xmapsto{t} \mu'$ implies $\mu' \in S$. An infinite computation $\mu_1 \xmapsto{t_1} \mu_2 \xmapsto{t_2} \cdots \mu_i \xmapsto{t_i} \mu_{i+1} \cdots$ is *fair* if for every transition t, if t is enabled at infinitely many μ_{i_l} $(l \geq 1)$, then there exist infinitely many j_l $(l \geq 1)$ such that $t_{j_l} = t$. (In words, if a transition is enabled infinitely many times, then the transition must occur infinitely often as well.) See, e.g., [9, 10] for more about Petri nets and their related problems.

Our analytical model is based on the model of *probabilistic Petri nets*, which is defined as a 4-tuple (P,T,φ,p), where P, T, and φ are the same as those defined earlier, and p : M × T → [0,1] is the *transition probability function* such that $\forall \mu \in$ M, $\sum_{t \in T} p_\mu(t) = 1$. (Here M denotes the set of all markings.) In words, a probabilistic Petri net associates each marking $\mu \in$ M with an individual transition probability function.

Of note is that $p_\mu(t)$ is indeed a conditional probability of firing a transition t given the system being at marking μ (sometimes denoted by $Pr(t \mid \mu)$ as often seen in probability textbooks); hence it may differ from $p_{\mu'}(t)$ once $\mu' \neq \mu$. A *path* of a probabilistic Petri net $\mathcal{P}$=(P, T, φ, p) is a nonempty (finite or infinite) sequence $\mu_1 \xmapsto{t_1} \mu_2 \xmapsto{t_2} \cdots \mu_i \xmapsto{t_i} \mu_{i+1} \cdots$ of alternative markings and transitions, such that $\mu_i \in M, t_i \in T$ and $p_{\mu_i}(t_i) > 0$ for all $i \geq 0$. For each $\mu \in$ M, let Π_μ denote the set of all infinite paths starting from μ, and $\mathcal{B}_\mu \subseteq 2^{\Pi_\mu}$ the smallest σ-algebra of measurable subsets that contains all the *cylindrical sets*

$$\Pi_\mu(\sigma_1) \equiv \{\sigma \in \Pi_\mu \mid \sigma_1 \text{ is a prefix of } \sigma\},$$

where σ_1 ranging over the finite paths starting from μ. The probability measure π on $\mathcal{B}_\mu$ is defined so that for each cylindrical set containing prefix σ_1, say $\mu_1 \xmapsto{t_1} \mu_2 \xmapsto{t_2} \cdots \mu_{n-1} \xmapsto{t_{n-1}} \mu_n$, we have

$$\pi(\Pi_\mu(\sigma_1)) = Pr(\sigma_1), \text{ where } Pr(\sigma_1) = \prod_{i=1}^{n-1} p_{\mu_i}(t_i).$$

Those probabilities for paths following a prefix give rise to a unique probability measure on B_μ.

For ease of expression, the following notations will be used throughout the rest of this paper. Let σ, σ' be transition sequences, and t be a transition.

- $\#_\sigma(t)$ represents the number of occurrences of t in σ.
- $Tr(\sigma) = \{t | t \in T, \#_\sigma(t) > 0\}$, denoting the set of transitions used in σ.
- $\sigma \dot{-} \sigma'$ is defined inductively as follows. Suppose σ' = t$_1$...t$_n$. Let σ_0 be σ. If t$_i$ is in σ_{i-1}, let σ_i be σ_{i-1} with the leftmost occurrence of t$_i$ deleted; otherwise, let $\sigma_i = \sigma_{i-1}$. Finally, let $\sigma \dot{-} \sigma' = \sigma_n$. For instance, if $\sigma = t_1 t_2 t_3 t_4 t_5$ and $\sigma' = t_1 t_3 t_4$, then $\sigma \dot{-} \sigma' = t_2 t_5$. Intuitively, $\sigma \dot{-} \sigma'$ represents the transition sequence resulting from removing each transition of σ' from the leftmost occurrence of such a transition in σ (if the transition exists).

Given a computation $\mu_0 \xmapsto{\sigma} \mu$, a sequence σ' is said to be a *rearrangement* of σ if $\#_\sigma(t) = \#_{\sigma'}(t), \forall t \in T$, and $\mu_0 \xmapsto{\sigma'} \mu$.

2.2. Dependability-related problems

In this paper, we focus on the following dependability-related problems:

- *The termination with probability 1 problem*: Given a probabilistic PN $\mathcal{P}$ and a set of markings S (called the *termination set*), and let Π_μ^S represent the set of computations reaching S from marking μ, the problem is to compute the set $T_{P_r=1}(\mathcal{P}, S) = \{\mu \mid$ the probability of reaching S from marking μ is one, i.e., $Pr(\Pi_\mu^S) = 1\}$.
- *The self-stabilization with probability 1 problem*: In spite of having some non-self-stabilizing computations, in practice a system might be considered

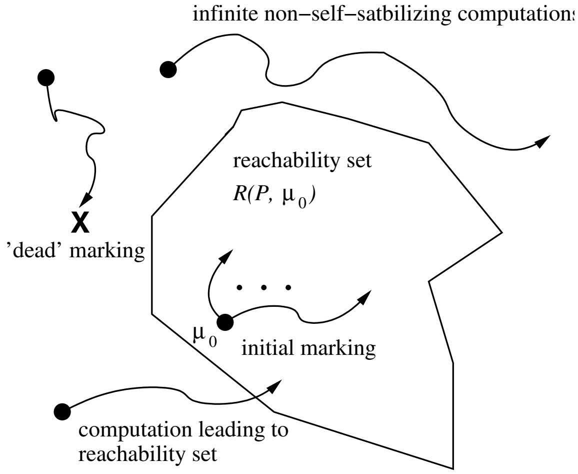

Figure 2. Non-self-stabilizing computations. (See top two paths.)

fault-tolerant if the probability of the system being self-stabilized equals one. Given a probabilistic PN $\mathcal{P}$ with initial marking μ_0, a computation σ from marking μ_1 is said to be *non-self-stabilizing* iff one of the following holds:

(1) σ ($\mu_1 \to \mu_2 \to \cdots \to \mu_m$, for some m) is finite such that μ_m is a 'dead' marking (i.e., μ_m has no immediate successor in $\mathcal{P}$) and $\mu_m \notin R(\mathcal{P}, \mu_0)$, <u>or</u>

(2) σ ($\mu_1 \to \mu_2 \to \cdots \to \mu_i \to \cdots$) is infinite such that $\forall i \geq 1, \mu_i \notin R(\mathcal{P}, \mu_0)$.

See Figure 2. Let Π_μ^{NSS} denote the set of *non-self-stabilizing* computations with respect to $R(\mathcal{P}, \mu_0)$ from marking μ. The *self-stabilization with probability 1 problem* is to compute the set $SS_{P_r=1}(\mathcal{P}, \mu_0) = \{\mu \mid$ the probability of reaching $R(\mathcal{P}, \mu_0)$ from μ is one; or $Pr(\Pi_\mu \setminus \Pi_\mu^{NSS}) = 1$, i.e., $Pr(\Pi_\mu^{NSS}) = 0\}$, where the *difference* $\Pi_\mu \setminus \Pi_\mu^S \equiv \{\sigma \mid \sigma \in \Pi_\mu, \sigma \notin \Pi_\mu^s\}$.

- *The controllability with probability 1 problem:* A *controlled PN* is simply a PN (P, T, φ) with its set of transitions T being partitioned into T_c (the set of *controllable transitions*) and T_u (the set of *uncontrollable transitions*). A *control policy* is a mapping $N^k \to 2^{T_c}$. (What it means is that at each marking, the control policy selects a subset of controllable transitions from which the next transition to fire must come, unless the next transition is an uncontrollable transition.) In order to cope with the situation when the probability-embedded controllable transitions are disabled, we have to make an assumption about the restricted behavior of the pcf-PN under control. As a general setting, e.g., [7, 5], for supervisory control of probabilistic systems, the supervisor will dynamically disable certain set of controllable transitions such that the occurrence probability of disabled transitions becomes zero, whereas the occurrence probability of the remaining enabled transitions, inclusive of uncontrollable transitions, is increased in proportion to their probability in the uncontrolled system. The same assumption is considered herein about probabilistic PNs in that the occurrence probability of the transition t enabled by control policy h at marking μ is given by

$$\begin{aligned} Pr(t \mid h \text{ enables } \Gamma_\mu) &= Pr(t \mid t' \in \Gamma_\mu) \\ &= \frac{p_\mu(t)}{\sum_{t' \in \Gamma_\mu} p_\mu(t')} \quad (1) \end{aligned}$$

where Γ_μ is the set of transitions enabled under h at marking μ. Let $\Pi_\mu^{\bar{S}}(h)$ denote the set of the infinite computations from marking μ that always avoid S under the control policy h regardless of how such computations are interleaved with transitions in T_u. The *controllability with probability 1 problem* is that of, given a controlled probabilistic PN and a set S of *forbidden markings*, computing the set $C_{Pr=1}(\mathcal{P}, S) = \{\mu \mid$ there exists a control policy h under which the probability of never reaching a marking in S from μ is one, i.e., $Pr(\Pi_\mu^{\bar{S}}(h)) = 1\}$. Intuitively, $C_{Pr=1}(\mathcal{P}, S)$ represents the set of markings from which the computation can be controlled with probability one to stay away from S. The interested reader is referred to [3] for more about *controlled Petri nets* and the related issues.

3. Valuation-based Dependability Analysis

Given a PN $\mathcal{P}$ =(P, T, φ), the idea of the *valuation method* is to devise a *valuation function* $f : N^k \to N \cup \{\infty\}$, which maps each marking μ to a value in $N \cup \{\infty\}$. Such a value $f(\mu)$ is called the *valuation* of the marking. Furthermore, if the set of markings of zero valuation is *forward-closed*, then the valuation along any Petri net computation is non-increasing, and in many cases, has the tendency to move towards the ground level (i.e., valuation zero). A valuation function f is said to be *monotone* if for every marking μ, if $\mu \overset{t}{\longmapsto} \mu'$ (for some marking μ' and transition t), then $f(\mu) \geq f(\mu')$. It is obvious that if f is monotone and $\mu \overset{\sigma}{\longmapsto} \mu'$ (where $\sigma \in T^*$), then $f(\mu) \geq f(\mu')$.

In this paper, we mainly focus on the following subclass of probabilistic Petri nets named *probabilistic conflict-free Petri nets* (pcf-PNs, for short):

1. $|p^\bullet| \leq 1$, <u>or</u> $\forall t \in p^\bullet$, t and p are on a self-loop (i.e., $t \in (p^\bullet \cap {}^\bullet p)$),where $p^\bullet = \{t \mid \varphi(p,t) > 0\}$ (resp., ${}^\bullet p = \{t \mid \varphi(t,p) > 0\}$) represents the set of output (resp., input) transitions of place p, and
2. with each marking μ we associate a transition probability distribution p_μ such that $p_\mu(t) > 0$ represents the probability of firing the enabled transition t at μ, and $\sum_{t \in T} p_\mu(t) = 1$.

As Condition 1 above indicates, a conflict-free PN requires every place being an input of more than one transition to be in a self-loop with each such transition, hence if a transition becomes enabled, the only way to disable it is to fire itself [6] (i.e., $\forall t, t' \in$ T, $t \neq t'$, $\mu \xmapsto{t} \mu'$ and $\mu \xmapsto{t'}$ implies $\mu' \xmapsto{t'}$.). In words, a pcf-PN is a conflict-free PN extended with conditional probability on the firing of each transition enabled at a given marking.

Example 1 Figure 1 illustrates a conflict-free PN describing a system consisting of m *producers* and one *consumer*. The i-th producer iterates a loop consisting of a sequence of two actions, *produce* (denoted by p_i) followed by *send* (denoted by s_i), whereas the consumer iterates a loop containing the actions of *receive* (denoted by r) and *consume* (denoted by c).

Given a conflict-free PN $\mathcal{P}$=(P, T, φ) and a set of markings S, the following valuation function f will be used throughout the rest of this paper: $f(\mu)$ is defined to be the length of the shortest path from μ to a marking in S; if μ cannot reach S, $f(\mu)$ is ∞. Notice that $\forall \mu \in S, f(\mu) = 0$ (i.e., S defines the set of markings of zero valuation). What follows is another way to view such a valuation function. We partition N^k into a sequence of disjoint sets of markings $U_0, U_1, ..., U_\infty$ such that

$$U_0 = S$$

$$U_1 = (pred(U_0)) - U_0$$

...

$$U_i = (pred(U_{i-1})) - (\cup_{j=0,\ldots,i-1} U_j), i \geq 1$$

$$U_\infty = N^k - (\cup_{j \geq 0} U_j)$$

It is not hard to see that $f(\mu) = i$ iff $\mu \in U_i$.

Before getting into the details of our analysis, we require a lemma concerning conflict-free PNs as well as the valuation function defined above.

Lemma 1 *(from Lemma 3.1 in [11]) Given a conflict-free PN $\mathcal{P}$ and a forward-closed set S, let f be the valuation function based upon the shortest path criterion defined above. The following hold:*

(1) f is always monotone,

(2) For an arbitrary μ and a path $\mu \xmapsto{\delta} \mu''$, if $\mu \xmapsto{\sigma} \mu'$ ($\mu' \in S$) is one of the shortest paths reaching S and $Tr(\delta) \cap Tr(\sigma) \neq \emptyset$, then $f(\mu'') < f(\mu)$. What this statement says is that if δ uses some transition(s) belonging to the shortest path to S, then δ will constitute a drop in valuation.

Proof: For the sake of completeness, we provide a proof sketch in the following. For (1), it suffices to show that $\mu \xmapsto{t} \mu_1$ implies $f(\mu) \geq f(\mu_1)$. Consider the following cases:

(a) ($f(\mu) = \infty$:) In this case, $f(\mu_1) = \infty$.

(b) ($0 < f(\mu) < \infty$:) Let $\mu \xmapsto{\sigma} \mu'$ be one of the shortest paths reaching S. Consider two cases:

(i) If $t \notin Tr(\sigma)$, $\exists \mu''$ such that $\mu' \xmapsto{t} \mu''$ and $\mu'' \in S$ (since S is forward-closed). Clearly, $\mu_1 \xmapsto{\sigma} \mu''$, and $f(\mu_1) \leq f(\mu)$ follows.

(ii) If $t \in Tr(\sigma)$, let $\mu \xmapsto{\sigma_1} \mu_2 \xmapsto{t} \mu_3 \xmapsto{\sigma_2} \mu'$ for some μ_2, μ_3, and t is not in σ_1. Due to PN being conflict-free, $\mu_1 \xmapsto{\sigma_1} \mu_3 \xmapsto{\sigma_2} \mu'$; hence, $f(\mu_1) \leq f(\mu) - 1$.

Figure 3 illustrates the monotone property of the valuation function for paths in a conflict-free PN.

Now consider (2). Due to $Tr(\delta) \cap Tr(\sigma) \neq \emptyset$, there must exist a first occurrence of transition, say $t \in T$, along δ as well as σ such that (see Figure 4)

- $\sigma = \sigma_1 t \sigma_2$ and $\mu \xmapsto{\sigma_1} \mu_1 \xmapsto{t} \mu_2 \xmapsto{\sigma_2} \mu'$, for some $\sigma_1, \sigma_2 \in T^*$, and μ_1, μ_2,and
- $\delta = \delta_1 t \delta_2$ and $\mu \xmapsto{\delta_1} \bar{\mu} \xmapsto{t} \bar{\mu}_1 \xmapsto{\delta_2} \mu''$, for some $\delta_1, \delta_2 \in T^*$, and $\bar{\mu}, \bar{\mu}_1$.

Again, due to PN $\mathcal{P}$ being conflict-free, it is not hard to see that $\bar{\mu}_1 \xmapsto{\sigma_1} \bar{\mu}_2$, for some $\bar{\mu}_2$, and $\bar{\mu}_2 \xmapsto{\sigma_2} \bar{\mu}'$ for some $\bar{\mu}' \in S$. Hence $\bar{\mu}_1 \xmapsto{\sigma_1 \sigma_2} \bar{\mu}' \in S$. (See Figure 4) Hence, $f(\bar{\mu}_1) \leq |\sigma_1 \sigma_2| = f(\mu) - 1$. Using the result of (1) and the fact that $\bar{\mu}_1 \xmapsto{\delta_2} \mu''$, (2) follows. ■

It should be noticed that being conflict-free plays an important role in Lemma 1 as the following example indicates.

Example 2 Consider a non-conflict-free PN $\mathcal{P}$ shown in Figure 5. (In $\mathcal{P}$, t and t' are in conflict with each other.) The initial marking is $(1, 0, 0, 0)$ (i.e., one token in p_1 while p_2, p_3, and p_4 are empty) and suppose

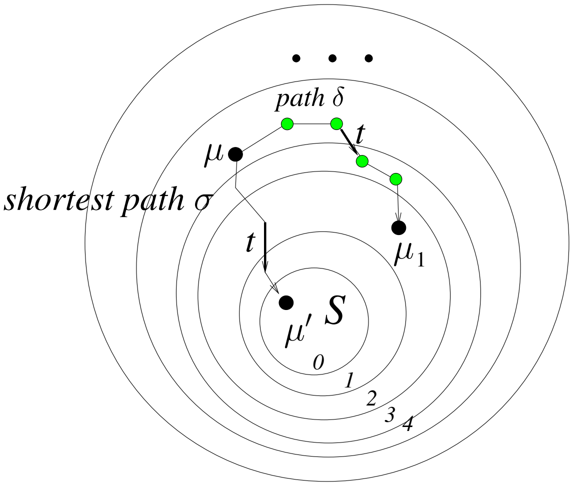

Figure 3. Paths and the associated valuations.

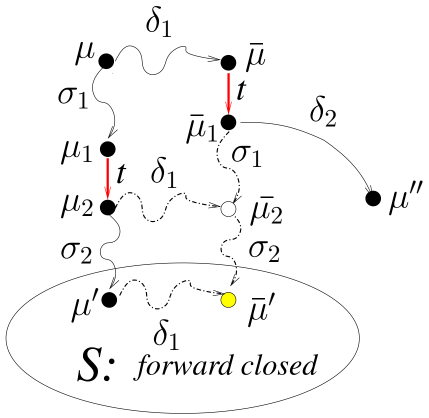

Figure 4. Illustration of the proof of Lemma 1.

$S = \{(0,0,0,1)\}$, which is clearly forward-closed as none of the transitions is enabled in S. From the easy fact that $f((1,0,0,0)) = 1$, $(1,0,0,0) \xmapsto{t'} (0,1,0,0)$ and $f(0,1,0,0) = 2$, the monotonicity property does not hold for $\mathcal{P}$.

Now we are in a position to reason about the list of problems mentioned in Section 2 for pcf-PNs in the framework of the valuation method.

Theorem 1 *Given a pcf-PN $\mathcal{P}$=(P,T,φ,p), and a forward-closed termination set S, if there exists no path from μ leading to 'dead' marking beyond S, then $\mu \in T_{P_r=1}(\mathcal{P}, S)$ iff $\mu \in pred^*(S)$.*

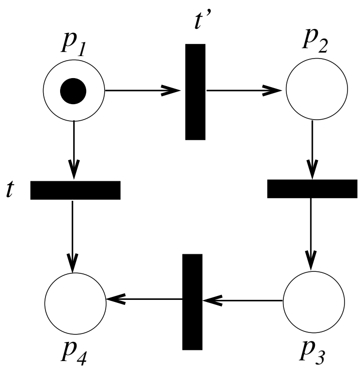

Figure 5. A non-conflict-free Petri net.

Proof: We let $f(\mu) = 0, \forall \mu \in S$. Clearly, if $\mu \notin pred^*(S)$, then there exists no firing sequence σ such that $\mu \xmapsto{\sigma} \mu' \in S$; hence $\mu \notin T_{P_r=1}(\mathcal{P}, S)$. The *only-if part* follows. Now we show the *if* part, i.e., $\mu \in pred^*(S) \Longrightarrow \mu \in T_{P_r=1}(\mathcal{P}, S)$. Suppose, in contrast, that $\mu \in pred^*(S)$, yet $\mu \notin T_{P_r=1}(\mathcal{P}, S)$; i.e., $Pr(\Pi_\mu^S) < 1$ or $Pr(\Pi_\mu \setminus \Pi_\mu^S) > 0$. Since no 'dead' markings reachable from μ beyond S, every computation $\sigma \in \Pi_\mu \setminus \Pi_\mu^S$ is infinite. According to Lemma 1, there shall exist a marking μ_1, and computations σ_1 and σ_2 such that $\mu \xmapsto{\sigma_1} \mu_1 \xmapsto{\sigma_2}$, $f(\mu) \geq f(\mu_1) > 0$, and the valuation along $\mu_1 \xmapsto{\sigma_2}$ remains $f(\mu_1)$. Let $\mu_1 \xmapsto{t_1 t_2 \cdots t_i} \mu'$ be one of the shortest paths reaching some marking in S. If $t_1 \in Tr(\sigma_2)$, then the valuation along σ_2 must eventually drop below $f(\mu_1)$ (Lemma 1), which is a contradiction. Now, consider the case for which $t_1 \notin Tr(\sigma_2)$. For the path of σ_2, say $\mu_1 \xmapsto{x_1} \mu_2 \xmapsto{x_2} \cdots \mu_n \xmapsto{x_n} \cdots$, let X_i be the random variable assuming the value of the i-th transition, namely x_i, along the path, and $T_{\mu_i}^S$ the set of transitions being able to transfer μ_i to markings in S when fired. Since $t_1 \notin Tr(\sigma_2)$, t_1 would have been enabled infinitely often without being fired along σ_2. Hence, let $C = \max_{\sigma \in \Pi_\mu \setminus \Pi_\mu^S,\ \sigma = \sigma_1 \sigma_2} Pr(\sigma_1)$, we have

$$\begin{aligned}
Pr(\Pi_\mu \setminus \Pi_\mu^S) &= \sum_{\sigma \in \Pi_\mu \setminus \Pi_\mu^S,\ \sigma = \sigma_1 \sigma_2} Pr(\sigma_1) Pr(\sigma_2) \\
&\leq C \cdot \sum_{\sigma \in \Pi_\mu \setminus \Pi_\mu^S,\ \sigma = \sigma_1 \sigma_2} Pr(\sigma_2) \\
&= C \cdot \lim_{n \to \infty} \prod_{i=1}^{n} Pr(X_i \notin (T_{\mu_i}^S \bigcup \{t_1\})) \\
&\leq C \cdot \lim_{n \to \infty} \prod_{i=1}^{n} Pr(X_i \neq t_1) \\
&\leq C \cdot 0 \ \text{(as } Pr(X_i \neq t_1) = 1 - P_{\mu_i}(t_1) < 1) \\
&= 0.
\end{aligned}$$

– again a contradiction. Our theorem follows. ■

Theorem 2 *Given a pcf-PN $\mathcal{P}$=(P,T,φ,p) and an initial marking μ_0, $\mu \in SS_{P_r=1}(\mathcal{P}, \mu_0)$ iff $\mu \in pred^*(R(\mathcal{P}, \mu_0))$.*

Proof: Clearly $R(\mathcal{P}, \mu_0)$ is forward-closed. We let the valuations of those markings in $R(\mathcal{P}, \mu_0)$ be zero. If $\mu \notin pred^*(R(\mathcal{P}, \mu_0))$, none of the (finite or infinite) computations can reach $R(\mathcal{P}, \mu_0)$; hence $\mu \notin SS_{P_r=1}(\mathcal{P}, \mu_0)$. The *only-if-part* follows. Now we prove the *if-part*, i.e., $\mu \in pred^*(R(\mathcal{P}, \mu_0)) \implies \mu \in SS_{P_r=1}(\mathcal{P}, \mu_0)$. Following an argument similar to the proof of Theorem 1 with $\Pi_\mu \setminus \Pi_\mu^S$ being substituted for Π_μ^{NSS}, we have Pr(Π_μ^{NSS})=0; or the probability of reaching $R(\mathcal{P}, \mu_0)$ from μ is one. Hence $\mu \in SS_{P_r=1}(\mathcal{P}, \mu_0)$. ■

Theorem 3 *Given a controlled pcf-PN $\mathcal{P}$=(P, $T_u \cup T_c$, φ,p) and a forward-closed set S of forbidden markings, $\mu \in C_{Pr=1}(\mathcal{P}, S)$ iff $\mu \notin pred^*(S)$.*

Proof: It is obvious that if $\mu \notin pred^*(S)$, any computation from μ (regardless of whether it is controlled or not) never encounters S. It suffices to show that $\mu \in pred^*(S) \implies \mu \notin C_{Pr=1}(\mathcal{P}, S)$; i.e., any computation has the tendency to move towards S with nonzero probability, in spite of the presence of a control policy. Suppose, in contrast, that $\mu \in pred^*(S)$, yet $\mu \in C_{Pr=1}(\mathcal{P}, S)$; i.e., there exists a control policy h such that $Pr(\Pi_\mu^{\bar{S}}(h)) = 1$. It is clear that any $\sigma \in \Pi_\mu^{\bar{S}}(h)$ can be decomposed into $\delta_0\sigma_1\delta_1 \cdots \sigma_m\delta_m\sigma_{m+1}$ such that $\delta_0, \delta_1, ..., \delta_m \in (T_u)^*$, $\sigma_1, ..., \sigma_m \in (T_c)^*$, and σ_{m+1} is an infinite computation consisting of transitions from T_c. (As transitions in T_u cannot be disabled by the control policy, one may view the $\delta_0, \delta_1, ..., \delta_m$ segments as the steps performed by an 'adversary' trying to force the computation into S. This explains why the infinite suffix computation σ_{m+1} is assumed to use transitions in T_c only.) In words, σ can be decomposed into $\mu \overset{\sigma_1}{\longmapsto} \mu_1 \overset{\sigma_2}{\longmapsto}$ for some marking μ_1, and computations σ_1 and σ_2 such that $f(\mu) \geq f(\mu_1) > 0$, and the valuation along $\mu_1 \overset{\sigma_2}{\longmapsto}$ remains $f(\mu_1)$. Using an argument parallels to the proof of Theorem 1 with $\Pi_\mu \setminus \Pi_\mu^S$ being substituted for $\Pi_\mu^{\bar{S}}(h)$ and $p_{\mu_i}(t_1)$ being substituted for $\frac{p_{\mu_i}(t_1)}{\sum_{t \in \Gamma_{\mu_i}} p_{\mu_i}(t)}$ (refer to equation (1)) in the derivation of occurrence probability of σ_2, the path σ_2 must eventually enter S under the control policy h – a contradiction. ■

The following known result serves as a vehicle for us to compute the forward and backward reachability sets, which play a vital role in using the valuation method as our earlier theorems show.

Lemma 2 *(Howell et al. [4]) Given a conflict-free PN $\mathcal{P}$=(P, T, φ) and a marking μ_0, we can construct in nondeterministic polynomial time a system of linear inequalities $\mathcal{L}(\mathcal{P}, \mu_0, \mu)$ (of size bounded by a polynomial in the size of $\mathcal{P}$) such that $\mu \in R(\mathcal{P}, \mu_0)$ iff $\mathcal{L}(\mathcal{P}, \mu_0, \mu)$ has an integer solution. Furthermore, $\mathcal{L}(\mathcal{P}, \mu_0, \mu)$ remains linear even if μ_0 and μ are replaced by variables. (The reader is referred to Lemma 4.3 in [4] for a detailed description of the system of linear inequalities associated with $\mathcal{L}(\mathcal{P}, \mu_0, \mu)$.)*

What Lemma 2 says is that checking reachability for conflict-free PNs can be equated with solving the integer linear programming problem, which is known to be in NP. It is important to point out that, as μ_0 and μ can be regarded as variables, the forward and backward reachability sets are readily expressible in terms of integer linear programming.

Theorems 1-3, in conjunction with Lemma 2, immediately yield the following result.

Theorem 4 *Given a pcf-PN $\mathcal{P}$ and a forward-closed set S expressible in integer linear programming, the following sets are computable: $T_{P_r=1}(\mathcal{P}, S)$, $SS_{P_r=1}(\mathcal{P}, \mu_0)$ (assuming $\mathcal{P}$ satisfies the condition stated in Theorem 1), and $C_{P_r=1}(\mathcal{P}, S)$.*

4. Summary and directions for future research

We have demonstrated, through a *valuation-based* strategy, effective procedures for solving problems associated with a number of dependability-related properties such as *termination with probability 1*, *self-stabilization with probability 1*, and *controllability with probability 1* in a unified framework. One direction of future research is to see whether the valuation-based strategy has applications to other subclasses of probabilistic Petri nets. Another issue that deserves further investigation is the enhancement of probabilistic models and valuation methods for describing and reasoning about different dependability properties of real-time systems, as many real-world systems are of real-time nature.

References

[1] Clarke, E., Grumberg, O. and Long, D., Verification Tools for Finite-State Concurrent Systems, *Lecture Notes in Computer Science* **803**, Springer-Verlag, 124–175, 1994.

[2] Dijkstra, E., Self-stabilizing systems in spite of distributed control, *C.ACM* **17**, 643–644, 1974.

[3] Holloway, L. and Krogh, B., Controlled Petri nets: a tutorial survey, *Discrete Event Systems, Lecture Notes in*

Control and Information Sciences **199**, G. Cohen and J.-P. Quadrat (eds.), 158–168, Springer-Verlag, 1994.

[4] Howell, R., Rosier, L. and Yen, H., Normal and sinkless Petri nets, *Journal of Computer and System Sciences* **46**, 1–26, 1993.

[5] Kumar, R. and Garg, V. K., Control of stochastic discrete event systems modeled by Probabilistic languages, *IEEE Transaction on Automatic Control* **46(4)**, 593–606, 2001.

[6] Landweber, L. and Robertson, E., Properties of conflict-free and persistent Petri nets, *JACM* **25(3)**, 352–364, 1978.

[7] Lawford M. and Wonham, W. M., Supervisory control of probabilistic discrete event systems, *Proc. 36th Midwest Symp. Circuits Systems*, 327–331, 1993.

[8] Marsan, M. A., Balbo, G., Conte, G., Donatelli, S., and Franceschinis, G., Modeling with generalized stochastic petri nets, John Wiley & Sons, 1995.

[9] Murata, T., Petri nets: properties, analysis and applications, *Proc. of the IEEE* **77(4)**, 541–580, 1989.

[10] Peterson, J., Petri Net Theory and the Modeling of Systems, Prentice Hall, Englewood Cliffs, NJ, 1981.

[11] Yen, H. A valuation-based analysis of conflict-free Petri nets, *Systems and Control Letters* **45(5)**, 387-395, 2002.