

On Time-Average Limits in Deterministic and Stochastic Petri Nets*

Tomáš Brázdil
brazdil@fi.muni.cz

Ľuboš Korenčiak
korenciak@fi.muni.cz

Jan Krčál
krcal@fi.muni.cz

Jan Křetínský
jan.kretinsky@fi.muni.cz

Vojtěch Řehák
rehak@fi.muni.cz

Faculty of Informatics, Masaryk University, Botanická 68a, 60200 Brno, Czech Republic

ABSTRACT

In this poster paper, we study performance of systems modeled by deterministic and stochastic Petri nets (DSPN). As a performance measure, we consider long-run average time spent in a set of markings. Even though this measure often appears in DSPN literature, its existence has never been considered. We provide a DSPN model of a simple communication protocol in which the long-run average time spent in a fixed marking is *not* well-defined due to a highly unstable behavior of the model. Further, we introduce a syntactical restriction on DSPN which preserves most of the modeling power yet guarantees existence of the long-run average.

Categories and Subject Descriptors

G.3 [Mathematics of Computing]: Probability and statistics—*Stochastic processes*

General Terms

Performance, Theory

Keywords

deterministic and stochastic Petri nets, performance analysis, steady-state distribution, time-average limits

1. INTRODUCTION

Petri nets are widely used in modelling of concurrent systems, such as network components. A Petri net is specified by a finite set of *places*, a finite set of *transitions*, and non-negative numbers $i_{t,p}$ and $o_{t,p}$ for each transition t and place p . Configuration of a Petri net is called *marking* and assigns to each place a non-negative number of tokens. A transition t is *enabled* if every place p contains at least $i_{t,p}$ tokens. *Firing* an enabled transition t consist of removing $i_{t,p}$ tokens from each place p , and adding $o_{t,p}$ tokens to each place p . The net starts in an initial marking and changes its marking by subsequent firing of enabled transitions.

For the sake of performance evaluation, Petri nets are turned into a continuous-time model as follows. Whenever a transition t becomes enabled, its *time to fire* is set to a real

*The work has been supported by the Czech Science Foundation, grant No. P202/12/P612.




Figure 1: A DSPN example of an M/D/1/2 queue.

number. As long as t is enabled, its time to fire then decreases with the time flow until it reaches zero and the transition fires. If there are more transitions with zero time to fire, the non-determinism is resolved according to fixed priorities of transitions¹. In deterministic and stochastic Petri nets (DSPN) [3], there are transitions of three types with time to fire set as follows. For an *immediate* transition, it is set to 0; for a *deterministic* transition t , it is set to a constant $d_t > 0$; and for an *exponential* transition t , it is set randomly according to the exponential distribution with rate $r_t > 0$.

To illustrate the definitions, a model of an M/D/1/2 queue with memoryless arrival of jobs, one deterministic server, and capacity of the queue 2 is shown in Figure 1. Places are drawn as circles; immediate, deterministic, and exponential transitions are drawn as thin bars, thick black bars, and thick void bars, respectively (with their $d_t = 6$ or $r_t = 1$). Assume that the deterministic and the exponential transitions are enabled with time to fire 6 and 5, respectively. The exponential transition is fired first after 5 time units and the deterministic transition has at that moment time to fire 1.

In this paper, we focus on the long-run average behavior and discuss a crucial quantity, the *time-average limit* [2]

$$r(f) = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t f(X(u)) du,$$

of a function f that assigns a real number to each marking. Here, $X(u)$ denotes the marking of the net at time $u \geq 0$. Observe that in our example the utilization of the server can be expressed as a time-average limit by setting $f(m) = 1$ if the marking m has a token in P_2 and $f(m) = 0$, otherwise.

The time-average limits in DSPN have been widely studied. Yet, no one has addressed the existence of the limit $r(f)$. If the limit does not exist, its approximation, e.g. using simulation, makes no sense. In this paper, we show a simple example of a DSPN where the time-average limit $r(f)$ does not exist. We also provide a sufficient (but not necessary) syntactical restriction on DSPN guaranteeing existence of the limit $r(f)$.

¹The firing policy we consider is *race with enabling memory*.





Figure 2: DSPN without the time-average limit.

2. INSTABILITY IN DSPN

In Figure 2, there is an example of a DSPN model with unstable behavior in the long-run, not having a time-average limit. It is a model of a simple producer-consumer system operated by one exponential and three deterministic transitions using buffers B_t and B_c . In addition, on the right, there is a controller with immediate transitions switching a token between modes M_1 and M_2 .

Let us explain the behaviour of the model. The transition t_p produces a token exactly every time unit and places it to the transport buffer B_t . Due to the parallel competition of an exponential and a deterministic transition, the token is transported into the consumption buffer B_c in *at most* 1 time unit. The consumption of a token by the transition t_c takes again exactly 1 time unit. The consumption buffer B_c can hold at most two tokens. When a third token appears in the buffer, all three tokens are removed by switching the modes. The subsequent immediate transition returns two tokens back, re-enabling t_c .

The question to be analyzed is what is the average time spent in the mode M_1 , i.e. the limit $r(f)$ where f equals 1 if there is a token in M_1 and 0, otherwise. This time-average limit $r(f)$ does not exist (for most runs). Let us explain the behavior on the following figure.



At the top, there is a time axis with thick and thin marks showing when the production and transport transitions are fired, respectively. The time spent transporting is written above (and also indicated by crosses). In the middle, the value of time to fire of the consumption transition t_c is plotted. Note that t_c is fired only when the graph reaches zero level. The bullets on the zero level mark re-enabling of t_c due to a switch of modes (caused by a third token being transported into B_c). Observe that the modes are switched whenever the current transport takes less time than it has ever taken. The lower the current minimal transport time is, the more time it takes to set a new minimum. Hence, the stays in a single mode get longer and longer. Due to this fact, the partial time-averages oscillate and never reach a limit (as indicated in the plot at the bottom).

Formal proof of the non-existence of the time-average limit follows from [1]. Intuitively, it stems from the absolute precision of the deterministic transitions causing that along each run of the model the beginning of the n -th consumption gets closer and closer to the end of the n -th production but never overtakes it. Since this behavior is rarely to be expected in reality, the validity of such a model is at least questionable. However, we can restrict DSPN in such a way that this form of instability does not appear.

3. CONDITIONS ON STABILITY

We say that a transition t *initiates* a transition u if there is a reachable marking m such that t is enabled in m and firing t in m (potentially followed by firing a sequence of immediate transitions) leads to a marking where u is enabled and either $u = t$ or u was not enabled in m (or in markings passed by the immediate transitions). For example, t_3 initiates t_3 in Figure 1, as from the marking with one token in both P_1 and P_2 we reach by firing t_3 and immediately t_2 the marking with one token in both P_0 and P_2 where t_3 is still enabled.

We say that a DSPN *initiates increasingly* if there is a function f assigning to each deterministic transition its level $l \in \mathbb{N}_0$ such that there is at most one deterministic transition with level 0 (that can initiate other transitions arbitrarily) and all other deterministic transitions with positive levels can initiate only deterministic transitions of strictly higher levels. Note that the DSPN of Figure 2 does not initiate increasingly as both t_p and t_c initiate themselves.

THEOREM 3.1. *For a DSPN with finitely many reachable markings that initiates increasingly, the time-average limit $r(f)$ almost surely exist for any function f .*

The proof can be obtained as a straightforward modification of the proof of Theorem 4 in [1]. The markings and their enabled transitions correspond to states and their scheduled events in Generalized semi-Markov processes in [1], resp.

4. CONCLUSIONS

We show that a widely studied formalism of DSPN does not necessarily have a basic time-average limits which is caused by its unrealistic behavior. We provide syntactical restrictions on DSPN upon which this unrealistic behavior disappears and the time-average limits almost surely exist.

5. REFERENCES

- [1] T. Brázdil, J. Krčál, J. Křetínský, and V. Řehák. Fixed-delay Events in Generalized Semi-Markov Processes Revisited. In *Proceedings of CONCUR'11*, volume 6901 of *LNCS*, pages 140–155. Springer, 2011.
- [2] P.J. Haas. *Stochastic Petri Nets: Modelling, Stability, Simulation*. Springer, 2002.
- [3] M. Ajmone Marsan and G. Chiola. On Petri nets with deterministic and exponentially distributed firing times. In *Advances in Petri Nets 1987*, volume 266 of *LNCS*, pages 132–145. Springer, 1987.