

EXPSPACE lower bounds for the simulation preorder between a communication-free Petri net and a finite-state system[☆]

Śławomir Lasota

Institute of Informatics, Warsaw University

Abstract

We investigate the simulation preorder between finite-state systems and a simple subclass of Petri nets, communication-free nets (known also as BPP-nets). We show EXPSPACE lower bounds for the simulation problems, in both directions, as well as for the simulation equivalence. Our results improve previous PSPACE and co-NP lower bounds, given by Kučera and Mayr in [5]. This paper is a revised and improved version of [8].

Key words: Petri net, BPP-net, simulation preorder and equivalence, trace preorder and equivalence

Fix a finite labeling set \mathbf{Act} . A *transition system* \mathcal{A} is determined by a (possibly infinite) set of states S , an initial state $s_{\mathbf{init}} \in S$, and a family of binary transition relations \xrightarrow{a} over states, one for each $a \in \mathbf{Act}$. We define the simulation preorder in the standard way. Given two transition systems \mathcal{A} and \mathcal{B} , consider a relation R between states of \mathcal{A} and states of \mathcal{B} such that for any $(q_{\mathcal{A}}, q_{\mathcal{B}}) \in R$, any $a \in \mathbf{Act}$, and any state $q'_{\mathcal{A}}$ of \mathcal{A} , with $q_{\mathcal{A}} \xrightarrow{a} q'_{\mathcal{A}}$, there is $q'_{\mathcal{B}}$ such that $q_{\mathcal{B}} \xrightarrow{a} q'_{\mathcal{B}}$ in \mathcal{B} and $(q'_{\mathcal{A}}, q'_{\mathcal{B}}) \in R$. Each such R is called a *simulation relation* between \mathcal{A} and \mathcal{B} . We write $\mathcal{A} \preceq \mathcal{B}$ if the initial states are related by some simulation relation; equivalently, we could require that the initial states are related by *the greatest* simulation relation. We write $\mathcal{A} \simeq \mathcal{B}$, and say that \mathcal{A} and \mathcal{B} are *simulation equivalent*, if $\mathcal{A} \preceq \mathcal{B}$ and $\mathcal{B} \preceq \mathcal{A}$.

We investigate in this paper certain variants of the simulation problems between a finite-state system on one side, and a Petri net on the other, in both directions:

Problem: $\mathbf{FS} \preceq \mathbf{PN}$ ($\mathbf{PN} \preceq \mathbf{FS}$, respectively)

INSTANCE: a finite-state transition system \mathcal{A} and a labelled Petri net \mathcal{N}

QUESTION: $\mathcal{A} \preceq \mathcal{N}$ ($\mathcal{N} \preceq \mathcal{A}$, respectively) ?

[☆]This work has been partially supported by Polish government grant no. N206 008 32/0810.
Email address: `sl@mimuw.edu.pl` (Śławomir Lasota)

It is known that both problems $\text{FS} \preceq \text{PN}$ and $\text{PN} \preceq \text{FS}$ are decidable [3] and EPXSPACE-hard (e.g., the first one is a generalization of the non-termination of Petri nets).

A special case of $\text{FS} \preceq \text{PN}$ ($\text{PN} \preceq \text{FS}$) is obtained by restricting the net \mathcal{N} to be a *BPP-net*, i.e., a net with each transition having precisely one incoming arc. These restricted problems we call $\text{FS} \preceq \text{BPP}$ and $\text{BPP} \preceq \text{FS}$, respectively.

A labeled BPP-net is given by a finite number of control states q_1, \dots, q_n (with a distinguished initial one) and a finite number of transition rules of the following form:

$$q_i \xrightarrow{a} q_{i_1} \dots q_{i_k}, \quad (1)$$

for arbitrary $k \geq 0$. The meaning is that in q_i , an a -labeled transition is possible, spawning k independent systems in control states q_{i_1}, \dots, q_{i_k} . Formally, a BPP-net induces an infinite-state transition system, whose states are finite multisets $q_1^{m_1} \dots q_n^{m_n}$ of control states ($m_i \dots m_n \geq 0$), and transitions are induced by transition rules. A rule (1) induces a transition

$$q_1^{m_1} \dots q_n^{m_n} \xrightarrow{a} q_1^{l_1} \dots q_n^{l_n}$$

if $m_i > 0$ and (l_1, \dots, l_n) is obtained from (m_1, \dots, m_n) by decrementing m_i , and incrementing m_{i_1}, \dots, m_{i_k} . Note that BPP-nets do not exhibit synchronization.

In [5] it was shown that $\text{FS} \preceq \text{BPP}$ is PSPACE-hard and $\text{BPP} \preceq \text{FS}$ is co-NP-hard. We investigate in this paper some further simplifications of the latter two problems.

The asynchronous product of two transition systems \mathcal{A} and \mathcal{B} , denoted $\mathcal{A} \otimes \mathcal{B}$, has as states pairs $(q_{\mathcal{A}}, q_{\mathcal{B}})$ where $q_{\mathcal{A}}$ is a state of \mathcal{A} and $q_{\mathcal{B}}$ a state of \mathcal{B} . The initial state is a pair of initial states. There is a transition $(q_{\mathcal{A}}, q_{\mathcal{B}}) \xrightarrow{a} (q'_{\mathcal{A}}, q'_{\mathcal{B}})$ if either $q_{\mathcal{A}} \xrightarrow{a} q'_{\mathcal{A}}$ in \mathcal{A} and $q'_{\mathcal{B}} = q_{\mathcal{B}}$, or $q_{\mathcal{B}} \xrightarrow{a} q'_{\mathcal{B}}$ in \mathcal{B} and $q'_{\mathcal{A}} = q_{\mathcal{A}}$.

For a finite-state system \mathcal{A} , with state-space S and initial state s_{init} , by $\mathcal{A}!$ we mean the infinite asynchronous product. Its states are all infinite sequences of states from S . There is a transition $(s_i)_{i < \omega} \xrightarrow{a} (s'_i)_{i < \omega}$ in $\mathcal{A}!$ if $s_i \xrightarrow{a} s'_i$ in \mathcal{A} for some i , and $s'_j = s_j$ for all $j \neq i$. The initial state of $\mathcal{A}!$ is given by $s_i = s_{\text{init}}$, for all $i < \omega$. Equivalently, $\mathcal{A}!$ may be understood as the smallest solution of the equation $\mathcal{A}! = \mathcal{A}! \otimes \mathcal{A}$.

Two core problems investigated in this paper are the following:

Problem: $\text{FS} \preceq \text{FS}!$ ($\text{FS}! \preceq \text{FS}$, respectively)

INSTANCE: finite-state transition systems \mathcal{A}, \mathcal{B}

QUESTION: $\mathcal{A} \preceq \mathcal{B}!$ ($\mathcal{B}! \preceq \mathcal{A}$, respectively) ?

It is easy to see that $\text{FS} \preceq \text{BPP}$ and $\text{BPP} \preceq \text{FS}$ subsume $\text{FS} \preceq \text{FS}!$ and $\text{FS}! \preceq \text{FS}$, respectively, as one may easily describe, using a BPP-net, the behavior of $\mathcal{B}!$. It is sufficient to replace each transition $q_{\text{init}} \xrightarrow{a} q$ of \mathcal{B} , outgoing from its initial state q_{init} , by a BPP-transition:

$$q_{\text{init}} \xrightarrow{a} q \ q_{\text{init}}.$$

In this way, intuitively, whenever a transition leaves q_{init} , a fresh copy of \mathcal{B} is spawn, in state q_{init} . All the other transitions in \mathcal{B} remain unchanged.

Therefore any lower bound for the two core problems applies immediately to the more general problems as well. Our main results are:

Theorem 1. $\text{FS} \preceq \text{FS}!$ is *EXPSPACE-hard*.

Theorem 2. $\text{FS}! \preceq \text{FS}$ is *EXPSPACE-hard*, even if both given finite-state systems are deterministic.

Our results are an improvement of the lower bounds for $\text{FS} \preceq \text{BPP}$ and $\text{BPP} \preceq \text{FS}$ given in [5]:

Corollary 1. $\text{FS} \preceq \text{BPP}$, $\text{BPP} \preceq \text{FS}$ and $\text{FS} \simeq \text{BPP}$ are *EXPSPACE-hard*.

Furthermore, we show that our core problems are in fact as hard as $\text{FS} \preceq \text{BPP}$ and $\text{BPP} \preceq \text{FS}$:

Theorem 3. *There are polynomial-time reductions of $\text{FS} \preceq \text{BPP}$ and $\text{BPP} \preceq \text{FS}$ to $\text{FS} \preceq \text{FS}!$ and $\text{FS}! \preceq \text{FS}$, respectively.*

A trace of \mathcal{A} is any finite sequence $a_1 \dots a_n \in \text{Act}^*$ such that there is a sequence of transitions $s_{\text{init}} \xrightarrow{a_1} s_1 \dots \xrightarrow{a_n} s_n$ in \mathcal{A} , for some states s_1, \dots, s_n . The *trace inclusion* $\mathcal{A} \sqsubseteq \mathcal{B}$ holds if any trace of \mathcal{A} is also a trace of \mathcal{B} . In general the trace inclusion is strictly coarser than the simulation preorder. However, they collapse when the “bigger” system is deterministic. Hence by Theorem 2 we get:

Corollary 2. *Deciding whether $\mathcal{B}! \sqsubseteq \mathcal{A}$ is *EXPSPACE-hard*, for any relation \sqsubseteq which lies between the simulation preorder \preceq and the trace inclusion \sqsubseteq .*

In Sections 1–3 we present the proofs of Theorems 1–3. Theorems 1 and 2 are shown by reduction from two EXPSPACE-complete questions concerning Petri nets: non-termination and place-markability (or control-state reachability), respectively. We prefer to work with nondeterministic counter machines without zero-tests, purely for technical convenience. Our crucial observation is that it is possible to separate the control states of a machine (described by \mathcal{A} , say) from its counters (described by $\mathcal{B}!$) in such a way that existence of a simulation between \mathcal{A} and $\mathcal{B}!$ (or between $\mathcal{B}!$ and \mathcal{A}) corresponds faithfully to an answer to the above-mentioned EXPSPACE-complete questions. In this respect our approach is slightly related to those used in [1] and [10], to obtain the lower bound for LTL model-checking of BPP-nets, and for the simulation of a finite-state system by a product of such systems, respectively. However our technique is different from those used in the cited papers.

Related research. No upper bound, except for decidability, is known for the complexity of the simulation or trace inclusion between a finite-state system and a Petri net (or a BPP-net), in whatever direction. Decidability was shown in [3].

To our knowledge, the complexity of the core problems themselves was not studied before. Related problems are the simulation and trace pre-order between a given finite-state system \mathcal{A} and a composition of given finite-states systems $\mathcal{B}_1 \dots \mathcal{B}_n$:

INSTANCE: finite state transition systems $\mathcal{A}, \mathcal{B}_1, \dots, \mathcal{B}_n$

QUESTION 1: $\mathcal{A} \preceq \mathcal{B}_1 \otimes \dots \otimes \mathcal{B}_n$?

QUESTION 2: $\mathcal{B}_1 \otimes \dots \otimes \mathcal{B}_n \preceq \mathcal{A}$?

QUESTION 3: $\mathcal{B}_1 \otimes \dots \otimes \mathcal{B}_n \subseteq \mathcal{A}$?

Questions 1 and 2 are trivially in EXPTIME; the EXPTIME lower bound for Question 1 was shown in [10]. Question 3 is trivially PSPACE-hard; the PSPACE-algorithm was given in [16] basing on an earlier insight of [14]. Question 2 may be thus answered in PSPACE, when \mathcal{A} is assumed to be deterministic, but the exact complexity of the general case is not known¹. Our core problems may be readily seen as infinitary versions of questions 1 and 2. E.g., $\mathcal{A} \preceq \mathcal{B}!$, for $\mathcal{B} = \mathcal{B}_1 + \dots + \mathcal{B}_n$ the nondeterministic choice of $\mathcal{B}_1, \dots, \mathcal{B}_n$, is the same as $\mathcal{A} \preceq \mathcal{B}_1! \otimes \dots \otimes \mathcal{B}_n!$. It is however not clear whether the core problems subsume the above-mentioned questions. On the other hand, FS \preceq BPP and BPP \preceq FS certainly do subsume questions 1 and 2.

The EXPTIME lower bound was shown for certain variants of questions 1 and 2, where both sides are a composition of finite-state systems which moreover synchronize on certain actions [7], and hiding operator is allowed [13]. In [11] the EXPSPACE lower bound was given for the latter variant of trace equivalence.

As usually, the bisimulation equivalence is more tractable than the simulation equivalence (preorder). For instance, the bisimulation equivalence between a finite-state system and a BPP-net is in PTIME [4].

Game-theoretic characterization. A simulation relation witnessing $\mathcal{A} \preceq \mathcal{B}$ may be seen as a winning strategy of Duplicator in the *simulation game* over \mathcal{A} and \mathcal{B} . The game is played by two players, Spoiler and Duplicator, and proceeds in rounds. The first round starts in the configuration $(q_{\mathcal{A}}, q_{\mathcal{B}})$ consisting of initial states of \mathcal{A} and \mathcal{B} , respectively. Each round consists of one Spoiler's move, followed by one Duplicator's move. Spoiler chooses a label $a \in \mathbf{Act}$ and a state $q'_{\mathcal{A}}$ in \mathcal{A} such that $q_{\mathcal{A}} \xrightarrow{a} q'_{\mathcal{A}}$. Then, Duplicator answers by choosing a state $q'_{\mathcal{B}}$ in \mathcal{B} such that $q_{\mathcal{B}} \xrightarrow{a} q'_{\mathcal{B}}$ in \mathcal{B} . Afterwards, the next round continues from the new configuration $(q'_{\mathcal{A}}, q'_{\mathcal{B}})$. One of the players wins, if a configuration is reached such that the other player has no possible move. Otherwise, the play

¹Our preliminary research on this problem leads to the EXPTIME lower bound. Thus, surprisingly, simulation preorder is computationally harder than the trace inclusion.

is infinite – in this case Duplicator wins unconditionally. $\mathcal{A} \preceq \mathcal{B}$ if and only if Duplicator has a winning strategy in the simulation game over \mathcal{A} and \mathcal{B} .

1. EXPSPACE lower bound for FS \preceq FS!

We provide a reduction from non-termination problem for nondeterministic counter machines without zero-tests. Instead, one could equivalently consider Petri nets, or vector addition systems with states. EXPSPACE-hardness of non-termination follows from [9]; it is EXPSPACE-complete [12]. The intuition is that the protocol between Spoiler and Duplicator during the simulation game enables them to imitate jointly a computation of a given counter machine.

A counter machine without zero-tests has a set of states Q , a distinguished initial state $q_{\text{init}} \in Q$, a set of counters $C = \{c_1, \dots, c_k\}$, and a finite set of increment or decrement transitions of the following form:

in state q , increment (decrement) c_i and go to state q' .

A counter can always be incremented, but can only be decremented if it is greater than zero. A computation of \mathcal{M} is a (finite or infinite) sequence of states q_i and counter valuations $v_i : C \rightarrow \mathbb{N}$

$$(q_0, v_0) \quad (q_1, v_1) \quad (q_2, v_2) \quad \dots$$


consistent with transitions of \mathcal{M} . We assume that \mathcal{M} starts in q_{init} , with all counters set to 0: $q_0 = q_{\text{init}}$, $v_0(c_i) = 0$ for $i \leq k$. We say that \mathcal{M} *terminates* if all computations of \mathcal{M} are finite.

Assume, wlog, that each state of \mathcal{M} has precisely two outgoing transitions. (If there is only one, add a duplicate identical transition. If there are more than two, organize them into a binary tree, and use increments of an additional artificial counter on the branches. If there is no transition, add two decrementing self-loops. Note that this construction does not affect termination.)

We describe now a construction that, for a given machine \mathcal{M} , yields two finite-state systems \mathcal{A} , \mathcal{B} such that \mathcal{M} does not terminate if and only if $\mathcal{A} \preceq \mathcal{B}$!. During the simulation game, \mathcal{A} will keep track of the actual control state of \mathcal{M} while the values of counters will be stored in \mathcal{B} !.


Let q be an arbitrary state of \mathcal{M} . To simplify the presentation assume that the two transitions from q are the following: increment c_i and go to state q_1 ; decrement c_j and go to q_2 (all other possibilities, i.e., two increments or two decrements, are dealt with in exactly the same way).

\mathcal{A} will contain the transitions drawn on the left-hand side below, while the \mathcal{B} system will contain the transitions drawn on the right-hand side:



As a whole, \mathcal{A} will contain such transitions for each state q . Hence states of \mathcal{A} will include the states of \mathcal{M} , and additional three temporary states for each state q of \mathcal{M} . The initial state is q_{init} .


Σ stands for the whole alphabet hence the state of \mathcal{B} with the Σ -loop is unconditionally winning for Duplicator. The Spoiler plays first with the **choice**-move, but the actual choice between transition 1 and 2 is done as a response of Duplicator. Then in his second move, Spoiler has to respect the Duplicator's choice, otherwise he loses.² Finally, in the third step, Spoiler checks whether the transition of \mathcal{M} chosen by Duplicator is doable (this is relevant only in case of decrement, as increments are always doable). To answer this moves, in the \mathcal{B} system there is also a small component for each counter c_i :



To make all the right-hand side components into a single finite-state \mathcal{B} , we simply collapse state **ch** and states $\mathbf{zero}_1, \dots, \mathbf{zero}_k$ into one initial state **ch**. It is irrelevant which copy of \mathcal{B} will be used by Duplicator to decide between 1 and 2. The very construction of \mathcal{A} guarantees that in each position in the game, at most one copy of \mathcal{B} is in **ch**₁ or **ch**₂ state, as each Spoiler's choice move is

²This construction is an adaptation of the “Duplicator’s choice” technique of [2, 15]. While mostly used in the context of bisimulation, it was also applied for the simulation equivalence in [6].

immediately followed by his 1 or 2 move. The whole \mathcal{B} system will look like:



The states of \mathcal{B} are $\{\text{ch}, \text{ch}_1, \text{ch}_2\} \cup \{\text{one}_i\}_{i=1\dots k} \cup \{\text{the loop state}\}$. The initial state is ch . $\Sigma = \{\text{choice}, 1, 2\} \cup \{\text{inc}_i, \text{dec}_i\}_{i=1\dots k}$. As explained above, precisely one of copies of \mathcal{B} will be used for making choices by Duplicator at each stage of the game; all the other copies will only be used to store the values of the counters.

We have: \mathcal{M} does not terminate iff Duplicator wins, i.e., iff $\mathcal{A} \preceq \mathcal{B}!$.


Indeed, if \mathcal{M} has an infinite computation, Duplicator's strategy is to follow precisely this computation. If Spoiler does not respect a Duplicator choice between 1 and 2, it is immediately punished by Duplicator in \mathcal{B} by going to the loop state. Otherwise, assume that all computations of \mathcal{M} are finite. Even if Spoiler has no actual initiative in the game, it is guaranteed that after a finite number of steps of \mathcal{M} no further transition is doable. This means that Duplicator has no response to Spoiler's check (note that this can only take place in case of decrement). This completes the proof.

2. EXPSPACE lower bound for $\text{FS}! \preceq \text{FS}$

We will give a reduction from the control state reachability of a nondeterministic counter machine without zero-tests (this is equivalent to markability of a given place of a Petri net, an EXPSPACE-complete problem [9, 12]). As before we assume that in each state of \mathcal{M} there are precisely two transitions, except for the distinguished state q_f , from which there are no outgoing transitions (hence \mathcal{M} always stops when it enters q_f). Observe that the transformation sketched in the previous section, applied to all states except for q_f , does not affect control-state reachability of q_f .


Given such an \mathcal{M} , we construct finite-state systems \mathcal{A} and \mathcal{B} such that \mathcal{M} has a finite computation that ends in q_f iff $\mathcal{B}! \not\preceq \mathcal{A}$.

Now the Spoiler plays in one of the (infinitely many) copies of \mathcal{B} and Duplicator responds in \mathcal{A} . Let q be an arbitrary state of \mathcal{M} different from q_f , and let the two transitions from q be the following: increment c_i and go to state q_1 ; decrement c_j and go to q_2 (again, the other cases are dealt analogously). \mathcal{A} will contain the following transitions:



Thus states of \mathcal{A} include the states of \mathcal{M} , an additional Σ -loop state, and two temporary states for each state of \mathcal{M} , except for q_f . Similarly as before, q_{init} will be the initial state in \mathcal{A} . Finally, there is no transition from q_f in \mathcal{A} .

The \mathcal{B} system will contain the following transitions:



Thus each copy of \mathcal{B} either remembers a single increment of a counter, or is used to choose between 1 and 2 – it is irrelevant in which copy Spoiler performs this choice. The initial state of \mathcal{B} is ch . $\Sigma = \{1, 2\} \cup \{\text{inc}_i, \text{dec}_i\}$. \mathcal{A} is constructed so that at each stage of the game, Duplicator narrows down the possible Spoiler moves. Spoiler may either choose between 1 or 2 only, which corresponds to a choice of the next transition of \mathcal{M} to be executed; or may execute the appropriate single counter operation. As long as Spoiler correctly follows a run of \mathcal{M} in that way, Duplicator can only stay passive. Hence, if q_f is reachable by a finite computation, this computation may be used as a Spoiler's winning strategy. If q_f is not reachable, no matter what computation is chosen

by Spoiler, finite or infinite, the computation chosen never reaches q_f , and hence Duplicator wins.

Remark 1. Both \mathcal{A} and \mathcal{B} in the above reduction are deterministic. Since \mathcal{A} is deterministic, all relations between the simulation preorder and the trace inclusion collapse, and therefore the hardness result applies to all such relations.

3. Reductions of $\text{FS} \preceq \text{BPP}$ and $\text{BPP} \preceq \text{FS}$

TODO ...

Acknowledgments. The author is grateful to an anonymous referee for the valuable comments, especially for a suggestion of simplification of the construction in the proof of Thm 1 in Section 1.

References

- [1] P. Habermehl, On the complexity of the linear-time μ -calculus for Petri nets, in: Proc. ICATPN'97, vol. 1248 of LNCS, 1997.
- [2] P. Jančar, High undecidability of weak bisimilarity for Petri nets, in: Proc. CAAP'95, vol. 915 of LNCS, 1995.
- [3] P. Jančar, F. Moller, Checking regular properties of Petri nets, in: Proc. CONCUR'95, vol. 962 of LNCS, 1995.
- [4] M. Kot, Z. Sawa, Bisimulation equivalence of a BPP and a finite-state system can be decided in polynomial time, ENTCS 138 (2005) 49–60.
- [5] A. Kučera, R. Mayr, Simulation preorder over simple process algebras, Information and Computation 173(2) (2002) 184–198.
- [6] A. Kučera, R. Mayr, Why is simulation harder than bisimulation?, in: Proc. CONCUR'02, vol. 2421 of LNCS, 2002.
- [7] F. Laroussinie, P. Schnoebelen, The state-explosion problem from trace to bisimulation equivalence, in: Proc. FoSSaCS'00, vol. 1784 of LNCS, 2000.
- [8] S. Lasota, EXPSPACE lower bounds for the simulation preorder between a communication-free Petri net and a finite-state system, Information Processing Letters 109 (2009) 850–855.
- [9] R. Lipton, The reachability problem requires exponential space, Tech. Rep. 62, Department of Computer Science, Yale University (1976).
- [10] A. Muscholl, I. Walukiewicz, A lower bound on web services composition, Logical Methods in Computer Science 4 (2:5) (2008) 1–14.
- [11] A. Rabinowich, Complexity of equivalence problems for concurrent systems of finite agents, Information and Computation 139(2) (1997) 111–129.

- [12] C. Rackoff, The covering and boundedness problems for vector addition systems, *Theor. Comput. Sci.* 6 (1) (1978) 223–231.
- [13] Z. Sawa, Equivalence checking of non-flat systems is exptime-hard, in: *Proc. CONCUR'03*, vol. 2761 of LNCS, 2003.
- [14] S. K. Shukla, H. B. Hunt, I. Daniel, J. Rosenkrantz, R. E. Stearns, On the complexity of relational problems for finite state processes, in: *Proc. ICALP'96*, vol. 1099 of LNCS, 1996.
- [15] J. Srba, Strong bisimilarity of simple process algebras: Complexity lower bounds, *Acta Informatica* 39 (2003) 469–499.
- [16] A. Valmari, A. Kervinen, Alphabet-based synchronisation is exponentially cheaper, in: *Proc. CONCUR'02*, vol. 2421 of LNCS, 2002.