


Flatness and Complexity of Immediate Observation Petri Nets

Mikhail Raskin 

Technical University of Munich, Germany
raskin@in.tum.de

Chana Weil-Kennedy 

Technical University of Munich, Germany
chana.weilkennedy@in.tum.de

Javier Esparza 

Technical University of Munich, Germany
esparza@in.tum.de

Abstract

In a previous paper we introduced immediate observation (IO) Petri nets, a class of interest in the study of population protocols and enzymatic chemical networks. In the first part of this paper we show that IO nets are globally flat, and so their safety properties can be checked by efficient symbolic model checking tools using acceleration techniques, like FAST. In the second part we study Branching IO nets (BIO nets), whose transitions can create tokens. BIO nets extend both IO nets and communication-free nets, also called BPP nets, a widely studied class. We show that, while BIO nets are no longer globally flat, and their sets of reachable markings may be non-semilinear, they are still locally flat. As a consequence, the coverability and reachability problem for BIO nets, and even a certain set-parameterized version of them, are in PSPACE. This makes BIO nets the first natural net class with non-semilinear reachability relation for which the reachability problem is provably simpler than for general Petri nets.

2012 ACM Subject Classification Theory of computation → Distributed computing models; Theory of computation → Concurrency

Keywords and phrases Petri Nets, Reachability Analysis, Parameterized Verification, Flattability

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2020.45

Funding This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme under grant agreement No 787367 (PaVeS).

Acknowledgements We thank Jérôme Leroux and Rupak Majumdar for interesting conversations that put us on the path of flatness and BIO nets. We also thank the reviewers whose comments allowed us to improve this paper, and fix a small mistake in Lemma 18.

1 Introduction

Immediate observation Petri nets (IO nets) model immediate observation population protocols, as introduced by Angluin *et al.* in their seminal paper on the expressive power of population protocols [2]. In an IO net each transition is defined by three places: the source place p_s , the destination place p_d , and the observed place p_o . The transition can move one token from p_s to p_d , provided that p_o is not empty (if $p_s = p_o$, then p_o should contain at least two tokens). In the population protocol interpretation, p_s , p_d , and p_o are three possible states of each of the identical agents executing the protocol, and a transition models an agent in the state p_s observing another agent in the state p_o and switching to the state p_d .



© Mikhail Raskin, Chana Weil-Kennedy, and Javier Esparza;
licensed under Creative Commons License CC-BY

31st International Conference on Concurrency Theory (CONCUR 2020).

Editors: Igor Konnov and Laura Kovács; Article No. 45; pp. 45:1–45:19

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In a previous paper [10] we investigated “many-to-many” versions of the reachability and coverability problems for IO nets, in which we have a set of initial markings and a set of final markings instead of the standard “one-to-one” versions with a single initial marking and a single final marking. The sets we consider are *cubes*, i.e., sets of markings obtained by attaching to each place a lower bound and an upper bound (possibly infinite) for the number of tokens. We showed that while the standard one-to-one problems are PSPACE-hard, they remain in PSPACE in the many-to-many case. This is in strong contrast with general conservative Petri nets (nets in which transitions neither create nor destroy tokens), for which the many-to-many versions of the problems become EXPSPACE-hard or even non-elementary.

In this paper we continue our study of IO nets, and initiate the study of Branching IO nets (BIO nets for short), in which transitions can create or destroy agents. BIO nets deserve study for at least three reasons:

- They are a natural generalization of both IO nets and communication-free nets (aka BPP nets), another very well studied subclass (see e.g. [8, 9, 17, 11, 14, 13, 16]).
- The reachability sets of BIO nets are not necessarily semilinear. In particular, Hopcroft and Pansiot’s well-known example of a Petri net with a non-semilinear reachability set (see [12]) is a BIO net. The classes of unbounded Petri nets for which the reachability problem is demonstrably simpler than for arbitrary Petri nets, like BPP-nets, reversible nets, and IO nets, have semilinear reachability sets. This makes BIO nets ideal to investigate the existence of efficient verification techniques that do not depend on semilinearity.
- BIO nets are a natural model for enzymatic catalytic reactions of the form $A + C \rightarrow C + B_1 + \dots + B_n$ with more than one product. For example, catalase degrades hydrogen peroxide into water and oxygen, a reaction of the form $A + C \rightarrow C + B_1 + B_2$ [7]. Since IO nets have been used to model and analyze enzymatic reactions $A + C \rightarrow C + B$ (see [1, 4, 15]), we expect our results to find a similar application.

In this paper we prove that IO nets are globally flat, in the sense of Leroux and Sutre [14]. In particular, this shows that their reachability relation is semilinear. Since the reachability relation of BIO nets is not semilinear, this result cannot extend to BIO nets. However, we prove that they are locally *pre**-flat, also in the sense of [14]¹. Both global and local flatness allow us to analyze nets applying existing symbolic model checking tools like FAST [5], LASH [6] and TREX [3]. Further, we prove that the many-to-many versions of the reachability and coverability problems for BIO nets are still PSPACE-complete, as for IO nets. To the best of our knowledge, this makes BIO nets the first natural class of nets whose reachability relation is non-semilinear for which these problems have elementary complexity.

Our flatness and complexity results are consequences of two theorems, called the Shortening Theorems for IO and BIO nets. They state that if M is reachable from M' , then M can be reached by a sequence of bounded *accelerated* length, defined as the length of the sequence after exhaustively replacing any subsequence of the form tt by t . In the case of IO nets the accelerated length is independent of the initial and final markings, while for BIO nets it only depends on the final marking. We consider that the Shortening Theorems are also interesting in their own right.

The paper is organized as follows. Section 2 contains preliminaries, and Section 3 defines IO and BIO nets. Section 4 states the Shortening Theorems, and derives our flatness and complexity results for the one-to-one reachability and coverability problems as corollaries.

¹ Actually, the locally flat of [14] are what we call locally *post**-flat. A net is locally *pre**-flat iff its reverse net is locally *post**-flat, and so with respect to reachability questions the difference is immaterial.

The proof of the Shortening Theorem for IO nets is given in Section 5 and our main result, the Shortening Theorem for BIO nets, is proved in Section 6. Finally, we prove in Section 7 that the many-to-many reachability and coverability problems remain in PSPACE.

2 Preliminaries

Multisets. A *multiset* on a finite set E is a mapping $C: E \rightarrow \mathbb{N}$, i.e. for any $e \in E$, $C(e)$ denotes the number of occurrences of element e in C . Let $\{e_1, \dots, e_n\}$ denote the multiset C such that $C(e) = |\{j \mid e_j = e\}|$. Operations on \mathbb{N} like addition or comparison are extended to multisets by defining them component wise on each element of E . Subtraction is allowed in the following way: if C, D are multisets on set E then for all $e \in E$, $(C - D)(e) = \max(C(e) - D(e), 0)$. We call $|C| \stackrel{\text{def}}{=} \sum_{e \in E} C(e)$ the *size* of C , and $\|C\| \stackrel{\text{def}}{=} \{e \mid C(e) > 0\}$ the *support* of C . Given a total order $e_1 \prec e_2 \prec \dots \prec e_n$ on E , a multiset C can be equivalently represented by the vector $(C(e_1), \dots, C(e_n)) \in \mathbb{N}^n$. A set $V \subseteq \mathbb{N}^n$ is *linear* if there is a root $r \in \mathbb{N}^n$ and a set $\{p_1, \dots, p_n\}$ of periods such that $V = \{v + \sum_{i=1}^n \lambda_i p_i \mid \lambda_1, \dots, \lambda_n \in \mathbb{N}\}$, and *semilinear* if it is the union of a finite set of linear sets. A relation on \mathbb{N}^n is semilinear if it is semilinear as a set of \mathbb{N}^{2n} . All these notions extend to sets of multisets.

Place/transition Petri nets with weighted arcs. A *Petri net* N is a triple (P, T, F) consisting of a finite set of *places* P , a finite set of *transitions* T and a *flow function* $F: (P \times T) \cup (T \times P) \rightarrow \mathbb{N}$. A *marking* M is a multiset on P , and we say that a marking M puts $M(p)$ *tokens* in place p of P . The *size* of M , denoted by $|M|$, is the total number of tokens in M . The *preset* $\bullet t$ and *postset* $t \bullet$ of a transition t are the multisets on P given by $\bullet t(p) = F(p, t)$ and $t \bullet(p) = F(t, p)$. A transition t is *enabled* at a marking M if $\bullet t \leq M$, i.e. $\bullet t$ is component-wise smaller or equal to M . If t is enabled then it can be *fired*, leading to a new marking $M' = M - \bullet t + t \bullet$. We let $M \xrightarrow{t} M'$ denote this.

Reachability and coverability. Given $\sigma = t_1 \dots t_n$ we write $M \xrightarrow{\sigma} M_n$ when $M \xrightarrow{t_1} M_1 \xrightarrow{t_2} M_2 \dots \xrightarrow{t_n} M_n$, and call σ a *firing sequence*. We write $M' \xrightarrow{*} M''$ if $M' \xrightarrow{\sigma} M''$ for some $\sigma \in T^*$, and say that M'' is *reachable* from M' . A marking M *covers* another marking M' , written $M \geq M'$ if $M(p) \geq M'(p)$ for all places p . A marking M is *coverable* from M' if there exists a marking M'' such that $M' \xrightarrow{*} M'' \geq M$. The *reachability relation* is the set of pairs of markings (M, M') such that $M \xrightarrow{*} M'$, and we denote it $\xrightarrow{*}$. The sets of predecessors and successors of a set \mathcal{M} of markings of N are $\text{pre}^*(\mathcal{M}) \stackrel{\text{def}}{=} \{M' \mid \exists M \in \mathcal{M}. M' \xrightarrow{*} M\}$ and $\text{post}^*(\mathcal{M}) \stackrel{\text{def}}{=} \{M \mid \exists M' \in \mathcal{M}. M' \xrightarrow{*} M\}$, respectively.

Global and local flatness. A net $N = (P, T, F)$ is *globally flat* if there exist transition words $w_1, w_2, \dots, w_k \in T^*$ such that for every two markings M', M , if $M' \xrightarrow{*} M$, then there exist $j_1, \dots, j_k \geq 0$ satisfying $M' \xrightarrow{w_1^{j_1} \dots w_k^{j_k}} M$. Observe that the words w_1, w_2, \dots, w_k are independent of both M and M' . A net $N = (P, T, F)$ is *locally pre*-flat* (resp. *locally post*-flat*) if for every M (resp. M') there exist transition words $w_1, w_2, \dots, w_k \in T^*$ such that for every M' (resp. M) satisfying $M' \xrightarrow{*} M$ there exist $j_1, \dots, j_k \geq 0$ such that $M' \xrightarrow{w_1^{j_1} \dots w_k^{j_k}} M$. The locally flat Petri nets of [14] correspond to our *post*-flat* nets.




Figure 1 Examples of IO and BIO nets.

3 Immediate Observation and Branching Immediate Observation Nets

We recall the definition of immediate observation nets (IO nets), as introduced in [10], and extend it to branching immediate observation nets (BIO nets).

► **Definition 1.** A transition t of a Petri net is an immediate observation transition (IO transition) if there are places p_s, p_d, p_o , not necessarily distinct, such that $\bullet t = \{p_s, p_o\}$ and $t\bullet = \{p_d, p_o\}$. We call p_s, p_d, p_o the source, destination, and observed places of t , respectively. A Petri net is an immediate observation net (IO net) if all its transitions are IO transitions.

A transition t of a Petri net is a branching IO transition (BIO transition) if there is $k \geq 0$ and places $p_s, p_{d_1}, \dots, p_{d_k}, p_o$, not necessarily distinct, such that $\bullet t = \{p_s, p_o\}$ and $t\bullet = \{p_{d_1}, \dots, p_{d_k}, p_o\}$. A Petri net is a branching IO net (BIO net) if all its transitions are BIO transitions.


In the following examples, we allow ourselves to consider IO and BIO nets containing transitions with no observed place. To make the net a formally correct IO or BIO net, it suffices to add an extra marked place which acts as observed place for these transitions.

► **Example 2.** Figure 1a shows an IO net taken from the literature on population protocols [2]. Intuitively, it models a protocol allowing a crowd of undistinguishable agents that can only interact in pairs to decide whether they are at least 3. Initially all agents are in state p_1 , modelled by tokens in place p_1 . If two agents in state p_1 interact, one of them moves to state p_2 (transition t_1). If two agents in state p_2 interact, one of them moves to p_3 (transition t_2). Finally, an agent in state p_3 can “attract” all other agents to state p_3 (transitions t_3 and t_4). Given a marking M_0 with tokens only in p_1 , if $M_0(p_1) \geq 3$ and the pairs of tokens that interact next are chosen uniformly at random, then eventually all tokens reach p_3 .

Figure 1b shows a BIO net representing a client server interaction. If the server S observes a client C , it creates a worker W , which creates a response R and terminates. The client C “leaves” after observing a response. Responses may expire.

IO nets are *conservative*, i.e. there is no creation or destruction of tokens, while BIO nets are not. The next example, taken from [12], shows that BIO nets may have non-semilinear sets of reachable markings.

► **Example 3** ([12]). Consider the BIO net N of Figure 2, with states p, q, c_1, c_2, c_3 and initial marking $M_0 = (1, 0, 0, 0, 1)$. The set of markings reachable from M_0 in N is characterized by the condition $(\mathbf{p} = 1 \wedge \mathbf{q} = 0 \wedge 0 < \mathbf{c}_2 + \mathbf{c}_3 \leq 2^{\mathbf{c}_1}) \vee (\mathbf{p} = 0 \wedge \mathbf{q} = 1 \wedge 0 < 2\mathbf{c}_2 + \mathbf{c}_3 \leq 2^{\mathbf{c}_1+1})$, where \mathbf{c} denotes the number of tokens in some place c . Informally, one token cycles between p and q , putting a new token in c_1 at every new cycle. When p is marked, tokens in c_3 can move to c_2 , and when q is marked, tokens in c_2 can move to c_3 while doubling their number (see Lemma 2.8 of [12]). Clearly the reachability relation of this BIO net is not semilinear.



■ **Figure 2** A non-flat BIO net.

4 Shortening Theorems

We introduce the main results of our paper, called the Shortening Theorems. We use them to prove flatness results for IO and BIO nets, and to extend complexity results of [10] for the reachability and coverability problems of IO nets to the (much harder) case of BIO nets. The Shortening Theorems themselves are proved in Sections 5 and 6, respectively.

First, we introduce a measure of the length of firing sequences that abstracts from the number of times a transition is consecutively executed.

► **Definition 4.** Let N be a Petri net, and let σ be a firing sequence. Let k_1, \dots, k_m be the unique positive natural numbers such that $\sigma = t_1^{k_1} t_2^{k_2} \dots t_m^{k_m}$ and $t_i \neq t_{i+1}$ for every $i = 1, \dots, m-1$. We say that σ has accelerated length m , and let $|\sigma|_a$ denote the accelerated length of σ .

The Shortening Theorems for IO and BIO show that a firing sequence leading from M' to M can be shortened to a sequence of bounded accelerated length. For IO nets the bound only depends on the net, not on the markings M or M' :

► **Theorem 5 (IO Shortening).** Let N be an IO net with n places, and let M', M be two markings of N . If $M' \xrightarrow{*} M$, then $M' \xrightarrow{\sigma} M$ for some σ of accelerated length $|\sigma|_a \leq (n^3 + 1)^n$.

Example 3 shows that for BIO nets the bound cannot be independent of both M and M' :

► **Example 6.** Recall the BIO net of Example 3 with states p, q, c_1, c_2, c_3 . It is easy to see that for $j \geq 1$ the marking $M_j \stackrel{\text{def}}{=} (1, 0, j, 0, 2^j)$ is reachable only via the firing sequence

$$(t_1 t_2 t_3 t_4) (t_1^2 t_2 t_3^2 t_4) \dots (t_1^i t_2 t_3^i t_4) \dots (t_1^j t_2 t_3^j t_4).$$

This sequence has accelerated length $4j$, which depends on the target marking M_j .

However, we can still obtain a bound independent of M' :

► **Theorem 7 (BIO Shortening).** Let N be a BIO net with n places, let M', M be two markings of N , and let $|M'| = m'$, $|M| = m$. Let $m_d := \max_{t \in T} |t^\bullet - \bullet t|$ denote the maximum number of tokens created by a transition of N . If $M' \xrightarrow{*} M$, then $M' \xrightarrow{\sigma} M$ for some σ of accelerated length $|\sigma|_a \leq 2^n (m+1)^n (n+1)^n$. Further, the intermediate markings along σ have size at most $(m' + 2^n (m+1)^n (n+1)^n (m+n) m_d) m_d^n$.

4.1 Flatness and complexity results

The Shortening Theorems lead easily to our flatness and complexity results:

► **Theorem 8.** IO nets are globally flat. BIO nets are locally pre*-flat, but neither globally flat nor locally post*-flat.

Proof.

- (a) We show that IO nets are globally flat. Let $N = (P, T, F)$ be an IO net with n places and $T = \{t_1, \dots, t_m\}$, and let $K = (n^3 + 1)^n$. By Theorem 5, for every two markings M' and M of N there is a firing sequence $t_{i_1}^{j_1} \dots t_{i_K}^{j_K}$ leading from M' to M . Since every such sequence belongs to the regular language $(t_1^* t_2^* \dots t_m^*)^K$, the words $w_1, w_2, \dots, w_{m \cdot K}$ given by $w_i = t_{((i-1) \bmod m) + 1}$ for every $1 \leq i \leq m \cdot K$ witness that N is globally flat.
- (b) We show that BIO nets are locally *pre**-flat. Let $N = (P, T, F)$ be a BIO net with n places and $T = \{t_1, \dots, t_m\}$, let M be a marking of N with $|M| = m$, and let $K = 2^n(m + 1)^n(n + 1)^n$. By Theorem 7, for every marking M' of N there is a firing sequence $t_{i_1}^{j_1} \dots t_{i_K}^{j_K}$ leading from M' to M . Proceed now as for (a).
- (c) We show that BIO nets are not locally *post**-flat, and so also not globally flat. Consider the BIO net of Figure 2 with states p, q, c_1, c_2, c_3 . Recall that for all $j \geq 1$, M_0 only reaches the marking $M_j \stackrel{\text{def}}{=} (1, 0, j, 0, 2^j)$ via $(t_1 t_2 t_3 t_4)(t_1^2 t_2^2 t_3^2 t_4) \dots (t_1^j t_2^j t_3^j t_4) \dots (t_1^j t_2^j t_3^j t_4)$. So in order to reach M_j it is necessary to fire j times a sequence of the form $t_1^{k_1} t_2^{k_2} t_3^{k_3} t_4^{k_4}$, which proves the result. \blacktriangleleft

► **Theorem 9.** *The reachability and coverability problems for BIO nets are PSPACE-complete.*


Proof. Reachability and coverability are PSPACE-complete for IO nets [10], and IO nets are a subclass of BIO nets, so the problems stay PSPACE-hard for BIO nets. By Savitch's theorem it suffices to show that the problems are in NPSPACE. Consider first the reachability problem. By the Shortening Theorem, given a BIO net with n places and two markings M and M' we can guess a firing sequence leading from M to M' , if one exists, using space $\log(f(n, m, m', m_d))$, where $f(n, m, m', m_d)$ is the exponential bound of the Shortening Theorem. So the reachability problem is in NPSPACE. For coverability, we reduce it to reachability in the usual way. Let M be the marking we want to cover. For each place p , we add a “destroying transition” τ_p with preset $\bullet t = \{p\}$ and postset $t \bullet = \emptyset$. It is easy to see that for every marking M' , the modified net N' has a firing sequence from M' to M iff N has a firing sequence from M' to some marking covering M . \blacktriangleleft

5 Shortening Theorem for IO nets

The proof of Theorem 5 is based on a result of [10] called the Pruning Lemma. We briefly introduce some notions required to state the lemma, and then the lemma itself. More details can be found in [10].

Trajectories and histories. Since the transitions of IO nets do not create or destroy tokens, we can give tokens identities. Given a firing sequence, each token of the initial marking follows a *trajectory* through the places of the net until it reaches the final marking of the sequence. The trajectories of the tokens between given source and target markings constitute a *history*.

Fix an IO net N . A *trajectory* of an IO net N is a sequence $\tau = p_1 \dots p_k$ of places. We let $\tau(i)$ denote the i -th place of τ . The i -th *step* of τ is the pair $\tau(i)\tau(i + 1)$. A *history* H of length h is a multiset of trajectories of length h . Given an index $1 \leq i \leq h$, the i -th *marking* of H , denoted M_H^i , is defined as follows: for every place p , $M_H^i(p)$ is the number of trajectories $\tau \in H$ such that $\tau(i) = p$. The markings M_H^1 and M_H^h are the *initial* and *final* markings of H , and we write $M_H^1 \xrightarrow{H} M_H^h$. A history H of length $h \geq 1$ is *realizable* if there exist transitions t_1, \dots, t_{h-1} and numbers $k_1, \dots, k_{h-1} \geq 0$ such that



■ **Figure 3** A realizable history of the IO net of Figure 1a before and after pruning.

- $M_H^1 \xrightarrow{t_1^{k_1}} M_H^2 \cdots M_H^{h-1} \xrightarrow{t_{h-1}^{k_{h-1}}} M_H^h$, where for every t we define $M' \xrightarrow{t^0} M$ iff $M' = M$.
- For every $1 \leq i \leq h-1$, there are exactly k_i trajectories $\tau \in H$ such that $\tau(i)\tau(i+1) = p_s p_d$, where p_s, p_d are the source and target places of t_i , and all other trajectories $\tau \in H$ satisfy $\tau(i) = \tau(i+1)$. Moreover, there is at least one trajectory τ in H such that $\tau(i)\tau(i+1) = p_o p_o$, where p_o is the observed place of t_i .

We say that $t_1^{k_1} \cdots t_{h-1}^{k_{h-1}}$ realizes H . Intuitively, at a step of a realizable history only one transition occurs, although perhaps multiple times, for different tokens. From the definition of realizable history we immediately obtain:

- $M' \xrightarrow{*} M$ iff there exists a realizable history with M' and M as initial and final markings.
- Every firing sequence that realizes a history of length h has accelerated length at most h .

► **Example 10.** Figure 3a shows a realizable history of the IO net of Figure 1a. It consists of six trajectories. The initial and final markings are $(5, 0, 1)$ and $(1, 0, 5)$. The history is realized by the firing sequence $t_3 t_1 t_1 t_3 t_2 t_4$.

Bunches and Pruning Lemma. A *bunch* is a multiset of trajectories with the same length and the same initial and final place. The Pruning Lemma states that every realizable history containing a bunch of trajectories from p to p' of size larger than the number of places n can be “pruned”, meaning that the bunch can be replaced by a smaller one, also leading from p to p' , while keeping the history realizable. (Notice, however, that the smaller bunch cannot always be chosen as a sub-multiset of the original one.)

► **Lemma 11 (Pruning Lemma).** *Let N be an IO net with n places. Let H be a realizable history of N containing a bunch $B \subseteq H$ of size larger than n . There exists a bunch B' of size at most n with the same initial and final places as B , such that the history $H' \stackrel{\text{def}}{=} H - B + B'$ (where $+$ and $-$ denote multiset addition and subtraction) is also realizable in N .*

► **Example 12.** The realizable history H of Figure 3a, leading from $(5, 0, 1)$ to $(1, 0, 5)$, has a bunch B of size $4 \geq n$ from p_1 to p_3 . Figure 3b shows a history H' , leading from $(4, 0, 1)$ to $(1, 0, 4)$, resulting from the application of the Pruning Lemma to H and B . The new bunch B' from p_1 to p_3 given by the Pruning Lemma is drawn in dashed trajectories. Notice that the trajectory of B' that passes through p_2 does not appear in B . The firing sequence $t_3 t_1 t_3 t_4$ realizes H' .

Proof of the Shortening Theorem. We need a Boosting Lemma, which states that duplicating a trajectory of a history of an IO net preserves realizability. Intuitively, duplicating a trajectory corresponds to adding a “shadow” to a token, that follows the token wherever it goes. Since an enabled IO transition can move arbitrarily many tokens from its source place to its destination place, the shadow token can always follow the primary token. A formal proof of the lemma is given in the Appendix.

► **Lemma 13** (Boosting Lemma). *Let H be a realizable history of an IO net containing a trajectory τ . The history $H + \langle \tau \rangle$ is also realizable.*

► **Theorem 5** (IO Shortening). *Let N be an IO net with n places, and let M', M be two markings of N . If $M' \xrightarrow{*} M$, then $M' \xrightarrow{\sigma} M$ for some σ of accelerated length $|\sigma|_a \leq (n^3 + 1)^n$.*

Proof sketch. We explain our proof strategy for the IO Shortening Theorem. Given $M' \xrightarrow{*} M$, we take a history H such that $M' \xrightarrow{H} M$. Repeatedly applying the Pruning Lemma, we construct another realizable history \tilde{H} such that $\tilde{T}_{p,q} = \min\{n, T_{p,q}\}$ for every two places p and q , where $T_{p,q}$ and $\tilde{T}_{p,q}$ denote the number of trajectories of H and \tilde{H} leading from p to q . Using the fact that H has at most n^3 trajectories, we show that \tilde{H} can be chosen so that its length is bounded by $(n^3 + 1)^n$. We are not done yet, because in general \tilde{H} does not lead from M' to M , we only have $\tilde{M}' \xrightarrow{\tilde{H}} \tilde{M}$ for markings \tilde{M}', \tilde{M} such that $\tilde{M}' \leq M'$ and $\tilde{M} \leq M$. In the last step we use the Boosting Lemma to add trajectories to \tilde{H} without increasing its length, yielding a realizable history \bar{H} of the same length as \tilde{H} , but satisfying $M' \xrightarrow{\bar{H}} M$. Finally, we extract from \bar{H} a sequence $M' \xrightarrow{\sigma} M$ of accelerated length at most $(n^3 + 1)^n$. The full proof can be found in the Appendix. ◀

6 Shortening Theorem for BIO nets

The proof of the BIO Shortening Theorem (Theorem 7) is very involved. It follows the proof outline of Theorem 5: Given a firing sequence, consider a history H realized by it, construct an equivalent “small” history H' , and extract from H' a sequence of short accelerated length. However, since BIO nets can create and destroy tokens, trajectories must be generalized to branching trajectories, which are trees of places; intuitively, the tree captures the cascade of tokens created by a token of the initial marking.


We fix a BIO net $N = (P, T, F)$ with n places, and let $m_d := \max_{t \in T} |t^\bullet - \bullet t|$ denote the maximum number of tokens created by a transition.

Branching trajectories. A *branching trajectory* of N is a nonempty, directed tree β whose nodes are labeled with places of P . A node labeled by p is called a *p-node*. The i -th level of β , denoted by $\beta(i)$, is the (possibly empty) set of nodes of β at distance $(i - 1)$ from the root. We let $M_\beta(i)$ denote the multiset of places labeling the nodes of $\beta(i)$. Observe that $M_\beta(i)$ is a marking. We say that β has *length* l if $\beta(l) \neq \emptyset$ and $\beta(l + 1) = \emptyset$.

Histories and realizable histories. A *history* H of length l is a forest of branching trajectories of length at most l . We use histories to describe a behaviour from an initial marking; the history contains a branching trajectory for each token of the initial marking.

Given a history H of length h and an index $1 \leq i \leq h$, the i -th level of H is the set $H(i) = \bigcup_{\beta \in H} \beta(i)$, and the *i -th marking of H* , denoted M_H^i , is the multiset $M_H^i = \sum_{\beta \in H} M_\beta(i)$. The markings M_H^1 and M_H^h are called the *initial* and *final* markings of H , and we write $M_H^1 \xrightarrow{H} M_H^h$. If the length of H is longer than the length of its branching trajectories, the final marking of H is the zero marking. Two histories are *equivalent* if they have the same initial and final markings.

A history H of length $h \geq 1$ is *realizable* if there exist transitions $t_1, \dots, t_{h-1} \in T$ and numbers $k_1, \dots, k_{h-1} \geq 0$ such that for every $1 \leq i \leq h - 1$ the set $H(i)$ can be partitioned into two sets:



■ **Figure 4** A decorated realizable history of a BIO net.

- A set $H_a(i)$ of exactly k_i nodes labeled by the source place of t_i . We call these nodes *active* nodes. Given a particular active node, say v , the multiset of labels of its children is the (possibly empty) multiset $\{p_{d_1}, \dots, p_{d_k}\}$ of destination places of t_i .
- A set $H_p(i)$ of nodes, each of them with exactly one child, carrying the same label as their parents. We call these nodes *passive* nodes. This set must contain at least one node labeled by the place p_o observed by t_i .

We say that the sequence $t_1^{k_1} \dots t_{h-1}^{k_{h-1}}$ realizes H . It follows easily from the definitions that

$$M_H^1 \xrightarrow{t_1^{k_1}} M_H^2 \dots M_H^{h-1} \xrightarrow{t_{h-1}^{k_{h-1}}} M_H^h \text{ holds (where } M \xrightarrow{t^0} M' \text{ iff } M = M').$$

From this definition we easily obtain:

- $M \xrightarrow{*} M'$ iff there exists a realizable history with M and M' as initial and final markings.
- Every firing sequence that realizes a history of length h has accelerated length at most h .

► **Example 14.** - Figure 4 shows a realizable history H of a BIO net with places $\{p, q, r\}$. H consists of six branching trajectories: β_c , four copies of β_1 , and β_2 . The initial and final markings are $(6, 0, 0)$ and $(0, 1, 1)$. The transition t_i executed at step i is

$$\begin{aligned} t_1 &= p \xrightarrow{p} \{q, p\} & t_2 &= p \xrightarrow{p} \{2q, p\} & t_3 &= p \xrightarrow{q} \emptyset & t_4 &= q \xrightarrow{q} \{r, q\} \\ t_5 &= r \xrightarrow{p} \emptyset & t_6 &= p \xrightarrow{q} \{2q\} & t_7 &= t_4 & t_8 &= q \xrightarrow{r} \emptyset \end{aligned}$$

where $t = x \xrightarrow{y} m$ denotes that x is the source place, y the observed place, and m the multiset of destination places of t . The firing sequence that realizes H is $t_1^5 t_2^2 t_3^4 t_4 t_5 t_6^4 t_7 t_8^8$. While the final marking of H is produced by β_c only, β_c is not realizable on its own. For example, the r -node of β_c at level 5 is destroyed in the next step by the firing of t_5 , but t_5 can only occur if there is at least one token in place p ; this token is supplied by β_1 or β_2 . We can think of β_1 and β_2 as branching trajectories that eventually become extinct, but before extinction provide tokens that need to be observable to fire some transitions.

Cargo, fuel, and smoke of a history. A decoration \widehat{H} of a history H consists of the history H itself and a partition of the nodes of H into *cargo*, *fuel*, and *smoke* nodes. Figure 4 shows not only a history H but also a decoration \widehat{H} . Cargo nodes are white, grey nodes are fuel,

and black nodes are smoke. Before giving the formal definition of a decoration, let us provide some intuition. Think of the sequence of markings of a history as the sequence of states of a ship. All nodes of the final marking are cargo, they are what the ship “delivers” in the end. At any other marking, the cargo nodes are the “causal predecessors” of the final cargo nodes. Every decoration has the same cargo nodes, they only differ in the partition of the other nodes into fuel and smoke. Intuitively, a decoration reserves the right to use fuel nodes to fire transitions (a p -node can be “used” to fire a transition that observes p), and commits to never using a smoke node or its descendants. The most conservative decoration (which always exists) is the one that declares all non-cargo nodes as fuel. Our first goal will be to show that every history has an equivalent *fuel-efficient* history that delivers the same cargo but admits a low-fuel decoration.

Formally, a *decoration* of H is a partition of the nodes of H into *cargo*, *fuel*, and *smoke* nodes satisfying the following conditions:

- A node of H is a *cargo node* iff it has at least one descendant in $H(l)$.
- All descendants of smoke nodes are smoke nodes.
- For every place p and level i , if $H(i)$ contains smoke p -nodes, then it also contains fuel p -nodes. (“No smoke without fuel”. Intuitively, the smoke p -nodes are not needed because the fuel p -nodes can be used instead.)

A *decorated history* is a pair consisting of H and a decoration of H . Observe that along all paths cargo comes before fuel, and fuel before smoke. Graphically, white nodes (if any) come before grey nodes (if any), and grey nodes before black nodes (if any).

Every history is equivalent to a fuel-efficient history. We prove that every realizable history has an equivalent realizable history with a *fuel-efficient* decoration, defined as follows:

► **Definition 15.** Let \widehat{H} be a decorated history. A place p is wasteful at level i if $\widehat{H}(i)$ contains more than n fuel p -nodes. A place p is wasteful in \widehat{H} if it is wasteful at some level; otherwise p is fuel-efficient in \widehat{H} . Finally, \widehat{H} is fuel-efficient if all places are fuel-efficient.

► **Example 16.** Since $n = 3$, in the decorated history of Figure 4 place p is wasteful at levels 1 to 6, and q is wasteful at levels 3 to 8. The history is not fuel-efficient.

The proof is based on a Replacement Lemma, which plays the same role as the combination of the Pruning and Boosting Lemmas for IO nets. We start by introducing a definition.


► **Definition 17.** The (p, i) -bunch of H , denoted $B_p(i)$, is the set of subtrees of H rooted at the p -nodes of $H(i)$.

Loosely speaking, the Replacement Lemma shows that if i is the earliest level at which p is wasteful, then the bunch $B_p(i)$ of trajectories can be replaced so that the new history has a decoration where p is not wasteful anymore. The lemma shows how to do this while ensuring that the histories before and after the replacement are equivalent. Repeated applications of the Replacement Lemma yield a fuel-efficient history.

Formally, given a history B'_p with p -nodes as roots and with the same number of trees as $B_p(i)$, we let $H[B'_p/B_p(i)]$ denote the result of replacing each tree of $B_p(i)$ by a different tree of B'_p . For this we assume that $B_p(i)$ and B'_p have been enumerated in some way, and the j -th tree of $B_p(i)$ is replaced by the j -th tree of B'_p . We state the Replacement Lemma:

► **Lemma 18 (Replacement Lemma).** Let \widehat{H} be a decoration of a realizable history H such that p is wasteful, and i is the earliest level at which p is wasteful. There exists a history B'_p such that $H' \stackrel{\text{def}}{=} H[B'_p/B_p(i)]$ is realizable, equivalent to H , and has a decoration whose fuel-efficient places contain all fuel-efficient places of \widehat{H} and p .

Proof sketch. We describe the history B'_p , illustrating the construction on the decorated history of Figure 4. In this example p is already wasteful at level $i = 1$, and $B_p(1) = H$. So all of H is replaced by the bunch B'_p , shown in Figure 5.



■ **Figure 5** Result of replacing $B_p(1)$ in the history of Figure 4.

In order to describe B'_p we need some notions. We call smoke and fuel nodes *transportation* nodes. Given a decorated history \widehat{H} , let $last(p)$ denote the last level i such that $\widehat{H}(i)$ contains a transportation p -node. A *place-level* is a pair (q, j) , where q is a place and j is a level of H . A *path* of place-levels is a concatenation of “steps” of two types: “doing nothing steps” from (r, l) to $(r, l + 1)$ such that $l < last(r)$, and “transportation steps” from (r, l) to $(s, l + 1)$ such that some transportation r -node of $\widehat{H}(l)$ has an s -child in $\widehat{H}(l + 1)$. We say that (q, j) is *reachable* from (p, i) if there is a path from (p, i) to (q, j) , and let $\mathcal{R}_{p,i}$ be the set of all place-levels (q, j) reachable from (p, i) . In our example we have $\mathcal{R}_{p,1} = \{(p, 1), \dots, (p, 6), (q, 3), \dots, (q, 8)\}$. (Observe that $(r, 5)$ does not belong to $\mathcal{R}_{p,1}$, because its parent is a cargo node.)

B'_p is the union of three sets of branching trajectories, B_c , B_f , and B_s (where c, f, s stand for cargo, fuel, and smoke):


- B_c contains all branching trajectories of $B_p(i)$ rooted at a cargo node. (In Figure 5, B_c is the singleton set $\{\beta_c\}$.) The decoration of B_c is chosen so that it conserves the cargo nodes of \widehat{H} . Intuitively, B_c ensures that H' delivers the same cargo as H .
- B_f contains a branching trajectory β_q for every q such that $(q, j) \in \mathcal{R}_{p,i}$ for some j . (In Figure 5, B_f contains the two trees β_p and β_q .) Intuitively, these trajectories guarantee that the new set $\mathcal{R}_{p,i}$ of \widehat{H}' is a superset of the old one, and so that any transition firing that relies on observing some place q at level j can still occur, because (q, j) is still reachable from (p, i) .

Let us now define β_q . (Figure 6 shows β_q for the history of Figure 5.) Let $first(q)$ be the smallest j such that $(q, j) \in \mathcal{R}_{p,i}$. There is a shortest path from (p, i) to $(q, first(q))$, and each step of the path corresponds to doing nothing or to executing a transition once. (In Figure 6 we have $(p, i) = (p, 1)$, $(q, first(q)) = (q, 3)$, and the path corresponds to doing nothing in the first step, and then firing t_2 .) Let δ_q be the corresponding branching trajectory. (In Figure 6, δ_q is the tree contained in the blue area.) First we append a path to each leaf of δ_q : If the leaf is, say, an r -node at level j , then we append to it a

45:12 Flatness and Complexity of Immediate Observation Petri Nets

path of r -nodes from level j to level $last(r)$. (Red area of Figure 6.) Then, we append to the end of each path a *destroyer*, i.e., a tree that makes the token disappear. We choose for this any subtree of \widehat{H} rooted in a transportation node of $(r, last(r))$. (Green area of Figure 6; in order to destroy a p -node we first transform it into two q -nodes by firing t_6 , wait while t_7 is fired in another part of the history, and then destroy the q -nodes by firing t_8 twice. The two q -nodes are destroyed by firing t_8 twice.) The decoration of β_q is chosen so that there is a fuel path rooted in (p, i) containing q -nodes from levels $first(q)$ to $last(q)$, and the rest is smoke.

- B_s contains $|B_p(i)| - |B_c| - |B_f|$ copies of a tree of smoke nodes β_s , consisting of a path of p -nodes, leading from level i to level $last(p)$, appended with a destroyer. Intuitively, this is smoke added to ensure that $H(i) = H'(i)$.



■ **Figure 6** Illustration of the construction of the set B_f of trees.

This concludes the description of B'_p . There are at most $|B_f| \leq n$ fuel nodes per level in B'_p , so p is fuel-efficient. The proof that H' is realizable, equivalent to H , and has a decoration in which there are no new wasteful places can be found in the appendix. ◀

Repeated applications of the Replacement Lemma yield the existence of a fuel-efficient decoration \widehat{H}' of a history H' equivalent to H .

► **Example 19.** Applying the Replacement Lemma to p and $i := 1$ and the decorated history \widehat{H} of Figure 4 yields the decorated history \widehat{H}' of Figure 5. Like H , it leads from $(6, 0, 0)$ to $(0, 1, 1)$. It is realized by $t_1 t_2 t_3 t_4 t_5 t_6^5 t_7 t_8^{12}$. Place p is no longer wasteful in \widehat{H}' , and in fact all places are fuel-efficient.

The next step of the proof is the Unique Footprint Lemma. Loosely speaking, it shows that for every history there exists an equivalent history in which any two levels differ in the cargo, the fuel, or the *support* of the smoke. This allows us to bound the length of the history. We need a preliminary lemma. Let $\widehat{H}_c(i)$, $\widehat{H}_f(i)$, $\widehat{H}_s(i)$ denote the multisets of cargo, fuel, and smoke nodes of $\widehat{H}(i)$. Intuitively, the Smoke Irrelevance lemma shows that we can always deliver the same cargo using the same fuel *independently* of the initial amount of smoke.

► **Lemma 20 (Smoke Irrelevance Lemma).** *Let \widehat{H} be a realizable decorated history of length h , and let μ be any multiset of places such that $\|\mu\| \subseteq \|\widehat{H}_s(1)\|$. There exists a realizable decorated history \widehat{H}' of length h such that $\widehat{H}'_s(1) = \mu$, and $\widehat{H}'_c(i) = \widehat{H}_c(i)$ and $\widehat{H}'_f(i) = \widehat{H}_f(i)$ for every level $1 \leq i \leq h$.*

Proof sketch. Rename $\nu \stackrel{\text{def}}{=} \widehat{H}_s(1)$ for clarity. To construct \widehat{H}' , start with \widehat{H} , and do the following for every place $p \in P$. If $\mu(p) \leq \nu(p)$, then delete $\nu(p) - \mu(p)$ smoke p -nodes from $\widehat{H}(1)$ as well as all their descendants (which are all smoke nodes by definition). If $\mu(p) > \nu(p)$,

then add to \widehat{H} $(\mu(p) - \nu(p))$ copies of an arbitrary tree β of smoke nodes of \widehat{H} rooted in $(p, 1)$. This tree exists because $p \in \|\mu\|$, and so $p \in \|\nu\|$. The addition of the copies of β maintains the “no smoke without fuel” property, because it was already fulfilled in \widehat{H} by the nodes of β . The smoke nodes of $\widehat{H}'(1)$ thus constructed are labelled by μ , and fuel and cargo nodes are neither added nor removed. The proof that \widehat{H}' is realizable can be found in the Appendix. \blacktriangleleft

► **Definition 21.** *Given a level $\widehat{H}(i)$ of a decorated history, define its footprint as the triple $(\widehat{H}_c(i), \widehat{H}_f(i), \|\widehat{H}_s(i)\|)$ (that is, we only take the support of $\widehat{H}_s(i)$, not $\widehat{H}_s(i)$ itself).*

► **Lemma 22 (Unique Footprint Lemma).** *Every realizable history has an equivalent fuel-efficient decorated history in which every level has a different footprint.*

Proof. Let \widehat{H} be a realizable decorated history. By the Replacement Lemma, we can assume w.l.o.g. that \widehat{H} is fuel-efficient. Assume further that \widehat{H} has minimal length h , i.e., every equivalent decorated history that is also fuel-efficient has length at least h . We claim that every level of \widehat{H} has a different footprint. Assume this is not the case. Then there exist two indices $1 \leq i < j \leq h$ such that $(\widehat{H}_c(i), \widehat{H}_f(i), \|\widehat{H}_s(i)\|) = (\widehat{H}_c(j), \widehat{H}_f(j), \|\widehat{H}_s(j)\|)$. The truncated history $\widehat{H}(j)\widehat{H}(j+1)\dots\widehat{H}(h)$ is clearly realizable. Since $\|\widehat{H}_s(i)\| = \|\widehat{H}_s(j)\|$, we can apply the Smoke Irrelevance Lemma with $\mu := \widehat{H}_s(i)$ and obtain a decorated history \widehat{H}' of length $h - j + 1$ such that $(\widehat{H}'_c(i), \widehat{H}'_f(i), \widehat{H}'_s(i)) = (\widehat{H}'_c(1), \widehat{H}'_f(1), \widehat{H}'_s(1))$ (notice: now $\widehat{H}'_s(i) = \widehat{H}'_s(1)$, instead of only $\|\widehat{H}'_s(i)\| = \|\widehat{H}'_s(1)\|$). But this implies $\widehat{H}(i) = \widehat{H}'(1)$, and so the concatenation $H(1)\dots H(i-1)\widehat{H}'(1)\dots H'(h-j+1)$ is also a realizable history. By the Smoke Irrelevance Lemma we have $\widehat{H}'_c(h-j+1) = \widehat{H}_c(h)$. Since the last levels of a decorated history only contain cargo nodes, this implies $\widehat{H}'(h-j+1) = \widehat{H}(h)$, and so the concatenation is equivalent to H . Further, since \widehat{H}' has the same cargo and fuel nodes as $\widehat{H}(j)\widehat{H}(j+1)\dots\widehat{H}(h)$, the concatenation is also fuel-efficient, contradicting that \widehat{H} has minimal length. \blacktriangleleft

We are equipped to prove the Shortening Theorem.

► **Theorem 7 (BIO Shortening).** *Let N be a BIO net with n places, let M', M be two markings of N , and let $|M'| = m', |M| = m$. Let $m_d := \max_{t \in T} |t^\bullet - \bullet t|$ denote the maximum number of tokens created by a transition of N . If $M' \xrightarrow{*} M$, then $M' \xrightarrow{\sigma} M$ for some σ of accelerated length $|\sigma|_a \leq 2^n(m+1)^n(n+1)^n$. Further, the intermediate markings along σ have size at most $(m' + 2^n(m+1)^n(n+1)^n(m+n)m_d)m_d^n$.*

Proof. We first prove the bound on the accelerated length. By the Unique Footprint Lemma, there is a history H such that $M' \xrightarrow{H} M$ and H has a decoration \widehat{H} where every level has a different footprint. So the length of \widehat{H} is bounded by the number of possible footprints of the histories leading from M' to M . Since, by definition, the number of cargo nodes cannot decrease from a level to the next, and the last level consists of only cargo, every level has between 0 and m cargo nodes per place. Since \widehat{H} is fuel-efficient, every level has between 0 and n fuel nodes per place. Finally, there are at most 2^n possible supports in a net with n places. So the number of footprints, and so the length of \widehat{H} , and the accelerated length of any firing sequence realizing \widehat{H} , is at most $2^n(m+1)^n(n+1)^n$.

Let us now prove the token bound. To bound the number of smoke nodes in each level, we apply the following operation. Replace every largest tree of smoke nodes (since the children of smoke nodes are smoke, this means trees rooted at smoke nodes whose parents are cargo or fuel) by the tree β_s defined as in the Replacement Lemma: β_s is a path of smoke p -nodes ending at level $last(p)$, appended by a p -destroyer tree. This maintains realizability, because

(by the “no smoke without fuel” property in \widehat{H}), it does not decrease the support of the multiset of places of any level. We call \widehat{H}' the resulting realizable history with decorated nodes. Note that the “no smoke without fuel” property may not hold in \widehat{H}' , so it is not formally a decorated history, but it is sufficient to conclude the proof. \widehat{H}' has the following property: smoke p -nodes can only create other nodes (which, by definition, are also smoke) at the level $last(p)$, and it can create at most m_d of them.

At all other levels j of \widehat{H}' , only cargo and fuel nodes can create nodes. There are at most $h' \leq 2^n(m+1)^n(n+1)^n$ levels, and at most $(m+n)$ cargo and fuel nodes per place. Each transition has a unique source place, and all the nodes are added to the initial m' nodes corresponding to the tokens of M' . Thus there are at most $m' + h'(m+n)m_d$ nodes at the first level $last(p)$ in which a smoke node creates nodes. At most all of the nodes are smoke, so at most $(m' + h'(m+n)m_d)m_d$ nodes are created. There are at most n levels $last(p)$, which each create at most the total amount of nodes times m_d nodes. Thus at every level of the history there are at most $(m' + h'(m+n)m_d)m_d^n$ nodes, concluding the proof. ◀

7 Many-to-many reachability and coverability

In [10] we prove that many-to-many versions of the reachability and coverability problems for IO nets are PSPACE-complete. We extend this result to BIO nets, which requires to use not only the Shortening Theorem itself, but also the lemmas conducting to its proof.

We recall some definitions of [10]. A set \mathcal{C} of markings of a net $N = (P, T, F)$ is a *cube* if there exist mappings $L: P \rightarrow \mathbb{N}$ and $U: P \rightarrow \mathbb{N} \cup \infty$ such that $M \in \mathcal{C}$ if and only if $L \leq M \leq U$. Abusing language, we identify \mathcal{C} with the pair (L, U) . Observe that cubes can be infinite sets of markings. The *cube-reachability (coverability)* consists of deciding, given a net N and cubes $\mathcal{C}, \mathcal{C}'$ of \mathbb{N} , whether there exist markings $M \in \mathcal{C}$ and $M' \in \mathcal{C}'$ such that M is reachable (coverable) from M' .

► **Theorem 23.** *The cube-reachability and cube-coverability problems for BIO nets are PSPACE-complete.*

Proof. PSPACE-hardness follows from PSPACE-hardness for IO nets. We show that the problems are in NPSPACE and apply Savitch’s theorem. Cube-coverability from \mathcal{C}' to $\mathcal{C} = (L, U)$ reduces to cube-reachability from \mathcal{C}' to the cube (L, U'') such that $U''(p) = \infty$ for all p , so it suffices to consider cube-reachability from $\mathcal{C}' = (L', U')$ to $\mathcal{C} = (L, U)$. For each place p with upper bound $U(p) = \infty$ in \mathcal{C} , add a “destroying transition” τ_p to N with preset $\bullet\tau_p = \{p\}$ and postset $\tau_p\bullet = \emptyset$. We guess a marking M of size m satisfying $M(p) = L(p)$ if $U(p) = \infty$, and $L(p) \leq M(p) \leq U(p)$ if $U(p) < \infty$. This reduces the problem to checking if M is reachable in the modified net from some marking of \mathcal{C}' . By Lemma 20, only the footprint of a marking matters for knowing whether it can reach marking M . We pick M' in \mathcal{C}' of size $m' \leq m + n^2 + \max(|L'|, n)$. The summands correspond to the cargo, fuel and smoke nodes of the initial marking of a fuel-efficient decorated history given by the Replacement Lemma if $M' \xrightarrow{*} M$ holds, where $\max(|L'|, n)$ is enough smoke nodes so that $M' \in \mathcal{C}'$ and any set of places is covered. By Theorem 9, $M' \xrightarrow{*} M$ can be checked in PSPACE. ◀

8 Conclusion

We have shown that immediate observation Petri nets are globally flat, allowing the use of existing efficient verification tools. We have also studied branching immediate observation nets, which are simultaneously a generalisation of IO nets, and of the Basic Parallel Processes model. The class of BIO nets significantly extends the expressive power of both IO nets and BPP nets, bringing together process creation and (restricted) cross-process interaction via

a simple and natural definition. While such an extension does not preserve global flatness, we have proven that local flatness is still preserved, and many-to-many reachability and coverability problems are still in PSPACE.

As BIO nets combine PSPACE-verifiable reachability and non-semilinear reachability relation, the further study of the structure of this reachability relation seems of interest. For instance, we plan to obtain the bounds on the size of the pre- and post- image of a marking, provided that these images are finite. It is also worth noting that the results of this paper still hold (up to a slight alteration of the Shortening Theorem bounds) if we define BIO transitions via the constraint $|\bullet t - t \bullet| \leq 1$. This is equivalent to extending BIO transitions with the possibility of multiple observations and the absence of a source place.

References

- 1 David Angeli, Patrick De Leenheer, and Eduardo D Sontag. A petri net approach to the study of persistence in chemical reaction networks. *Mathematical biosciences*, 210(2):598–618, 2007.
- 2 Dana Angluin, James Aspnes, David Eisenstat, and Eric Ruppert. The computational power of population protocols. *Distributed Computing*, 20(4):279–304, 2007.
- 3 Aurore Annichini, Ahmed Bouajjani, and Mihaela Sighireanu. TRex: A tool for reachability analysis of complex systems. In *Lecture Notes in Computer Science*, volume 2102, pages 368–372, 2001.
- 4 Paolo Baldan, Nicoletta Cocco, Andrea Marin, and Marta Simeoni. Petri nets for modelling metabolic pathways: a survey. *Natural Computing*, 9(4):955–989, 2010.
- 5 Sébastien Bardin, Alain Finkel, Jérôme Leroux, and Laure Petrucci. FAST: Fast acceleration of symbolic transition systems. In *Lecture Notes in Computer Science*, volume 2725, pages 118–121, 2003.
- 6 Bernard Boigelot. The LASH toolset homepage, 2014. URL: <http://www.montefiore.ulg.ac.be/~boigelot/research/lash/index.html>.
- 7 P. Chelikani, I. Fita, and P.C. Loewen. Diversity of structures and properties among catalases. *Cell. Mol. Life Sci.*, 61, 2004.
- 8 Søren Christensen, Yoram Hirshfeld, and Faron Moller. Decomposability, decidability and axiomatisability for bisimulation equivalence on basic parallel processes. In *LICS*, pages 386–396, 1993. URL: <https://www.wikidata.org/entity/Q59557027>.
- 9 Javier Esparza. Petri nets, commutative context-free grammars, and basic parallel processes. *Fundam. Inform.*, 31(1):13–25, 1997.
- 10 Javier Esparza, Mikhail Raskin, and Chana Weil-Kennedy. Parameterized analysis of immediate observation petri nets. In *Lecture Notes in Computer Science*, volume 11522, pages 365–385, 2019.
- 11 Laurent Fribourg. Petri nets, flat languages and linear arithmetic. In *WFLP*, pages 344–365, 2000.
- 12 John E. Hopcroft and Jean-Jacques Pansiot. On the reachability problem for 5-dimensional vector addition systems. *Theor. Comput. Sci.*, 8:135–159, 1979.
- 13 Slawomir Lasota. EXPSPACE lower bounds for the simulation preorder between a communication-free petri net and a finite-state system. *Inf. Process. Lett.*, 109(15):850–855, 2009.
- 14 Jérôme Leroux and Grégoire Sutre. Flat counter automata almost everywhere! In *ATVA*, volume 3707 of *Lecture Notes in Computer Science*, pages 489–503. Springer, 2005.
- 15 Wolfgang Marwan, Annegret Wagler, and Robert Weismantel. Petri nets as a framework for the reconstruction and analysis of signal transduction pathways and regulatory networks. *Natural Computing*, 10(2):639–654, 2011.
- 16 Ernst W. Mayr and Jeremias Weihmann. Complexity results for problems of communication-free petri nets and related formalisms. *Fundam. Inform.*, 137(1):61–86, 2015.
- 17 Hsu-Chun Yen. On reachability equivalence for BPP-nets. *Theor. Comput. Sci.*, 179(1-2):301–317, 1997.

A Shortening Theorem for IO nets

► **Lemma 13** (Boosting Lemma). *Let H be a realizable history of an IO net containing a trajectory τ . The history $H + \langle \tau \rangle$ is also realizable.*

Proof sketch. Let h be the length of H , and let $t_1^{k_1} \cdots t_{h-1}^{k_{h-1}}$ be a realization of H . For every $1 \leq i \leq h-1$ define k'_i as follows: if $\tau(i) = \tau(i+1)$, then $k'_i \stackrel{\text{def}}{=} k_i$; if $\tau(i) \neq \tau(i+1)$, then $k'_i \stackrel{\text{def}}{=} k_i + 1$. We claim that $t_1^{k'_1} \cdots t_{h-1}^{k'_{h-1}}$ is a realization of $H + \tau$. The proof is by induction on h .

Assume $h = 1$. Then H is realizable by t^0 for any transition t , and so is $H + \tau$.

Assume that the induction property holds for some $h \geq 1$, and let H be of length $h+1$, realizable by $t_1^{k_1} \cdots t_h^{k_h}$. By induction, the history $H + \tau$ truncated of its last step is realizable by $t_1^{k'_1} \cdots t_{h-1}^{k'_{h-1}}$. If $\tau(h) \neq \tau(h+1)$ in H , then since $\tau \in H$ and H is realizable, $\tau(h)\tau(h+1) = p_s p_d$ for p_s and p_d the source and destination places of t_h . Additionally, there are $k_h - 1$ other trajectories τ' such that $\tau'(h)\tau'(h+1) = p_s p_d$, and there is at least one trajectory τ' such that $\tau'(h)\tau'(h+1) = p_o p_o$. Thus $t_1^{k'_1} \cdots t_{h-1}^{k'_{h-1}} t_h^{k_h+1}$ realizes $H + \tau$. If $\tau(h) = \tau(h+1)$ in H , then $H + \tau$ is realized by $t_1^{k'_1} \cdots t_{h-1}^{k'_{h-1}} t_h^{k_h}$. ◀

► **Theorem 5** (IO Shortening). *Let N be an IO net with n places, and let M', M be two markings of N . If $M' \xrightarrow{*} M$, then $M' \xrightarrow{\sigma} M$ for some σ of accelerated length $|\sigma|_a \leq (n^3 + 1)^n$.*

Proof. Let H be a realizable history such that $M' \xrightarrow{H} M$, and let h be the length of H . For every two places p, q , let $B_{p,q}$ denote the bunch of all trajectories of H leading from p to q , and let $T_{p,q} = \text{size}(B_{p,q})$. Applying the Pruning Lemma to all bunches $B_{p,q}$ such that $T_{p,q} \geq n$, we obtain a new realizable history \tilde{H} satisfying

$$\tilde{T}_{p,q} = \min\{n, T_{p,q}\} \quad \text{for every } p, q \in P. \quad (1)$$

So \tilde{H} has $\sum_{p,q \in P} \tilde{T}_{p,q} \leq n^3$ trajectories. Let $M_{\tilde{H}}^1 \xrightarrow{t_1^{k_1}} M_{\tilde{H}}^2 \cdots M_{\tilde{H}}^{h-1} \xrightarrow{t_{h-1}^{k_{h-1}}} M_{\tilde{H}}^h$ be a realization of \tilde{H} . Since \tilde{H} has at most n^3 trajectories, we have $M_{\tilde{H}}^i(p) \leq n^3$ for every $p \in P$ and $1 \leq i \leq n$. If $h \geq (n^3 + 1)^n$, then there are $1 \leq i \neq j \leq h$ such that $M_{\tilde{H}}^i = M_{\tilde{H}}^j$, and the history \tilde{H}' obtained by “cutting out” the fragment of \tilde{H} between $M_{\tilde{H}}^i$ and $M_{\tilde{H}}^j$ is also realizable. (Formally, \tilde{H}' is the result of replacing every trajectory $\tau \in \tilde{H}$ by $\tau(1) \cdots \tau(i)\tau(j+1) \cdots \tau(h)$.) So w.l.o.g. we can assume $\tilde{h} < (n^3 + 1)^n$.

Since \tilde{H} is realizable, we have $\tilde{M}' \xrightarrow{\tilde{H}} \tilde{M}$ for some markings \tilde{M}', \tilde{M} . We examine the relation between M' and \tilde{M}' , and between M and \tilde{M} . For every place p , the initial (final) number of tokens of p in H is equal to the number of trajectories of H of starting in p (ending in p), and similarly for \tilde{H} . So we have

$$\begin{aligned} M'(p) &= \sum_{q \in P} T_{p,q} & \text{and} & & M(p) &= \sum_{q \in P} T_{q,p} \\ \tilde{M}'(p) &= \sum_{q \in P} \tilde{T}_{p,q} & \text{and} & & \tilde{M}(p) &= \sum_{q \in P} \tilde{T}_{q,p}. \end{aligned}$$

Further, for every place $p \in P$:

(a) $\tilde{M}'(p) \leq M'(p)$, and $\tilde{M}(p) \leq M(p)$.

Follows immediately from $\tilde{T}_{p,q} \leq T_{p,q}$ for every $q \in P$ (Equation 1).

(b) If $\tilde{M}'(p) = 0$ then $M'(p) = 0$, and if $\tilde{M}(p) = 0$ then $M(p) = 0$.

If $\tilde{M}'(p) = 0$ then $\tilde{T}_{p,q} = 0$ for every $q \in P$. So, by Equation 1, $\tilde{T}_{p,q} = T_{p,q}$ for every $q \in P$, and so $M'(p) = \sum_{q \in P} T_{p,q} = \sum_{q \in P} \tilde{T}_{p,q} = \tilde{M}'(p) = 0$. The proof for the target markings is analogous.

Let \overline{H} be the history obtained from \tilde{H} as follows: For every $p, q \in P$, if $\tilde{T}_{p,q} > 0$ then pick a trajectory $\tau \in B_{p,q}$, and set $\overline{B}_{p,q} = \tilde{B}_{p,q} + (\tilde{T}_{p,q} - T_{p,q} - 1) \cdot \tau$. By the Boosting Lemma, \overline{H} is realizable, and so there are markings $\overline{M}', \overline{M}$ such that $\overline{M}' \xrightarrow{\overline{H}} \overline{M}$. Further, by (a) and (b) above we have $\overline{T}_{p,q} = T_{p,q}$ for every $p, q \in P$, and so for every $p \in P$:

$$\overline{M}'(p) = \sum_{q \in P} \overline{T}_{p,q} = \sum_{q \in P} T_{p,q} = M'(p)$$

So we get $M' \xrightarrow{\overline{H}} M$. Since \tilde{H} and \overline{H} have the same length, we get $\bar{h} < (n^3 + 1)^n$. So every firing sequence realizing \overline{H} has accelerated length at most $(n^3 + 1)^n$, and we are done. \blacktriangleleft

B Shortening Theorem for BIO nets

We give ourselves a few more definitions to help in the proofs. We call smoke and fuel nodes *transportation* nodes. Given a decorated history \hat{H} , let $last(p)$ denote the last level i such that $\hat{H}(i)$ contains a transportation p -node. A *place-level* is a pair (q, j) , where q is a place and j is a level of H . A *path* of place-levels is a concatenation of “steps” of two types: “doing nothing steps” from (r, l) to $(r, l + 1)$, and “transportation steps” from (r, l) to $(s, l + 1)$ such that some transportation r -node of $\hat{H}(l)$ that has an s -child in $\hat{H}(l + 1)$. We say that (q, j) is *reachable* from (p, i) if there is a path from (p, i) to (q, j) , and let $\mathcal{R}_{p,i}$ be the set of all place-levels (q, j) reachable from (p, i) .

► **Lemma 18** (Replacement Lemma). *Let \hat{H} be a decoration of a realizable history H such that p is wasteful, and i is the earliest level at which p is wasteful. There exists a history B'_p such that $H' \stackrel{\text{def}}{=} H[B'_p/B_p(i)]$ is realizable, equivalent to H , and has a decoration whose fuel-efficient places contain all fuel-efficient places of \hat{H} and p .*

Proof. We first construct $H' \stackrel{\text{def}}{=} H[B'_p/B_p(i)]$ and show that it is realizable and equivalent to H . Then, we define a decoration \hat{H}' of H' , and show that it realizes the condition of the lemma.

Construction of H' . We define B'_p as the union of three sets of branching trajectories, B_c , B_f , and B_s (where c, f, s stand for cargo, fuel, and smoke):

- B_c contains all branching trajectories of $B_p(i)$ rooted at a cargo node.
- B_f contains a branching trajectory β_q for every $q \in \mathcal{R}_{p,i}$.
We define β_q . Let $first(q)$ be the smallest j such that $(q, j) \in \mathcal{R}_{p,i}$. Notice that $first(q) \leq last(q)$ for all $q \in P$, since by definition of reachability there exists a transportation q -node in level $first(q)$. There is a shortest path from (p, i) to $(q, first(q))$, and each step of the path corresponds to doing nothing or to executing a transition once. Let δ_q be the corresponding branching trajectory. First we append a path to each leaf of δ_q : If the leaf is, say, an r -node at level j , then we append to it a path of r -nodes from level j to level $last(r)$. Then, we append to the end of each path a *destroyer*, i.e., a tree that makes the token disappear. We choose for this any subtree γ_r of \hat{H} rooted in a transportation node of $(r, last(r))$.
- B_s contains $|B_p(i)| - |B_c| - |B_f|$ copies of a tree β_s , consisting of a path of p -nodes, leading from level i to level $last(p)$, appended with a destroyer γ_p .

We define the replacement $H' = H[B'_p/B_p(i)]$: we replace the trees of $B_p(i)$ with a cargo root in \hat{H} by the same tree in B_c , we replace some trees of $B_p(i)$ with a fuel root in \hat{H} by the trees of B_f (in any order), and the rest of the trees of $B_p(i)$ by the trees of B_s . This is

well-defined because the trees of B'_p all have p -nodes as root, there are no more than n trees in B_f and more than n trees with fuel roots in $B_p(i)$ since p is wasteful at i , and there are as many trees overall in B'_p as in $B_p(i)$.

History H' is realizable and equivalent. History H' is equivalent to history H : the trees added in $B'_p \setminus B_c$ all end in destroyers, and the other trees of H' were already in H , so H' has the same final marking. In case $i = 1$, the number of p -nodes in $H(i)$ and $H'(i)$ is the same so H' has the same initial marking.

History H is realizable, and we note $t_1^{k_1} \dots t_{h-1}^{k_{h-1}}$ a sequence that realizes it, for some transitions $t_1, \dots, t_{h-1} \in T$ and numbers $k_1, \dots, k_{h-1} \geq 0$. We show that H' is realizable using the same transitions but different numbers $l_1, \dots, l_{h-1} \geq 0$. Let $1 \leq j \leq h-1$. Let $H'_p(j)$ be the set of nodes of $H'(j)$ which have exactly one child with the same label, and let $H'_a(j)$ be the rest. We claim that for every node v' in $H'_a(j)$ with label r and multiset of children labels c , there exists a node v in $H_a(j)$ with label r and multiset of children labels c . By realizability of H this entails that v' is labeled with the source place p_s of t_j , and the multiset of labels of its children is the multiset $\{p_{d_1}, \dots, p_{d_k}\}$ of destinations of t_j .

Now to show our claim. Let v' a node of $H'_a(j)$. If v' is not a node of the subtree B'_p , or if v' is a node of B_c , then we are done. Let us assume this is not the case, i.e. $v' \in B_f \cup B_s$.

- If v' is in a tree β_s , then it is in a a destroyer (since v' is not in $H'_p(j)$) and so it is in a copy of a subtree of \widehat{H} .
- Assume v' is in a tree β_q for some $q \in \mathcal{R}_{p,i}$. If v' is in a destroyer then it is in a copy of a subtree of \widehat{H} , we are done. Otherwise, v' is in the tree δ_q induced by the shortest path ρ_q from (p, i) to $(q, \text{first}(q))$ in H . Since v' is not passive, i.e. $v' \notin H'_p(j)$, and by definition of how a path induces a tree, there is an r -node v of $H(j)$ with the same children as v' .

We now show that the set $H'_p(j)$ contains a node labeled by the place p_o observed by t_j . If there is a node labeled p_o in $H_p(j)$ that is not in $B_p(i)$, then it is also in $H'_p(j)$ and we are done. Let us assume that the only nodes of $H_p(j)$ labeled p_o are in $B_p(i)$. If there is a cargo node labeled p_o in $\widehat{H}_p(j)$ then it is also in $H'_p(j)$ so we are done. Otherwise there exists a transportation node v labeled p_o in $\widehat{H}_p(j)$, and $j \leq \text{last}(p_o)$ by definition. Since v is in $B_p(i)$, either v is in a tree with a cargo root, or place-level (p_o, j) is reachable from (p, i) . If v is in a tree of $B_p(i)$ with a cargo root, it is also in $B_c \subseteq B'_p$. Otherwise $(p_o, j) \in \mathcal{R}_{p,i}$, and therefore by construction there is a node in B'_p labeled p_o at every level between $\text{first}(p_o)$ and $\text{last}(p_o)$, in particular at j .

Decoration of H' . Let \widehat{H}' be the following decoration of H' .

We start with the nodes of B_f and B_s . In each tree β_q in B_f , constructed around the tree induced by a shortest path ρ_q from (p, i) to $(q, \text{first}(q))$, we let the nodes along the path ρ_q be fuel nodes, along with the nodes along one branch from $(q, \text{first}(q))$ to $(q, \text{last}(q))$. All the other nodes of β_q are defined as smoke nodes. We let all the nodes of the trees β_s be smoke nodes.

The rest of the nodes of H' are decorated in two steps. First, we set \widehat{H}' to be equal to \widehat{H} on the nodes of $H' \setminus (B_f \cup B_s)$, which is possible because $H' \setminus B'_p = H \setminus B_p(i)$ and the trajectories of B_c are trajectories of $B_p(i)$. Then, we do the following “re-decoration”. Let (q, j) be a place level reachable from (p, i) in H' . If there are any fuel nodes labeled q in $(H' \setminus B_f)(j)$, redecorate them and all their descendants as smoke nodes in \widehat{H}' . Do this for every (q, j) reachable from (p, i) .

The order of “cargo then fuel then smoke” is respected along the branching trajectories of \widehat{H}' because they are respected in B'_p , and there are no more than n trees with a fuel root in B'_p while there are more than n in $B_p(i)$. The cargo nodes in \widehat{H}' are well defined, as the cargo nodes of \widehat{H}' are the cargo nodes of \widehat{H} .

The smoke/fuel partition of \widehat{H}' is well defined: First, remark that the last level index $last(p)$ at which there is a transportation p -node in \widehat{H}' is equal to $last(p)$ in \widehat{H} , for any place p by construction. Let v be a smoke q -node at level $\widehat{H}'(j)$, for some q and j . We check that the “no smoke without fuel” condition is fulfilled. If (q, j) is reachable from (p, i) in H then there exists a fuel q -node in $\widehat{H}'(j)$ provided by β_q , since $j \leq last(q)$ by virtue of v being smoke. If (q, j) is not reachable from (p, i) in H , then there is no subtree of $B_p(i)$ rooted in a transportation p -node with a descendant labeled q . Therefore in \widehat{H}' , node v is not in B_f . Since it is also not in B_s , whose trees are only p nodes until $last(p)$, v is in either a tree of B_c or in no tree of B'_p , and therefore v exists also in \widehat{H} as a smoke node. Since the smoke/fuel partition of \widehat{H} is well defined, there exists a fuel q -node v' in $\widehat{H}(j)$. Since (q, j) is not reachable from (p, i) in H , v' is either in B_c or not part of $B_p(i)$ and so v' is also in $\widehat{H}'(j)$.

For every place-level (q, j) in H' reachable from (p, i) , there are at most n fuel q -nodes in $\widehat{H}'(j)$. Indeed, by definition, the only fuel nodes labeled q in $\widehat{H}'(j)$ are in $B_f(j)$. By definition of $B_f(j)$, the only fuel nodes labeled q in $B_f(j)$ are in the trees β_r for some $r \in \mathcal{R}_{p,i}$. There are at most n such trees, and in each tree there is at most one fuel node per level.

Therefore there are no wasteful places q at some level j such that (q, j) is reachable from (p, i) in H' . In particular, p is fuel-efficient since i is the earliest level at which p is wasteful in H . If there is a wasteful place-level in \widehat{H}' , then it is unreachable from (p, i) in H' . By definition of H' , this means that it is also a wasteful place-level in $H \setminus B'_p$ and thus in H . Thus the fuel-efficient places of \widehat{H}' contain all the fuel-efficient places of \widehat{H} , as well as the place p . ◀

We remind the reader that $\widehat{H}_c(i)$, $\widehat{H}_f(i)$, $\widehat{H}_s(i)$ denote the multiset of cargo, fuel, and smoke nodes of $\widehat{H}(i)$.

► **Lemma 20** (Smoke Irrelevance Lemma). *Let \widehat{H} be a realizable decorated history of length h , and let μ be any multiset of places such that $\|\mu\| \subseteq \|\widehat{H}_s(1)\|$. There exists a realizable decorated history \widehat{H}' of length h such that $\widehat{H}'_s(1) = \mu$, and $\widehat{H}'_c(i) = \widehat{H}_c(i)$ and $\widehat{H}'_f(i) = \widehat{H}_f(i)$ for every level $1 \leq i \leq h$.*

Proof. Rename $\nu := \widehat{H}_s(1)$ for clarity. To construct \widehat{H}' , start with \widehat{H} , and do the following for every place $p \in P$. If $\mu(p) \leq \nu(p)$, then delete $\nu(p) - \mu(p)$ smoke p -nodes from $\widehat{H}(1)$ as well as all their descendants (which are all smoke nodes by definition). If $\mu(p) > \nu(p)$, then add to \widehat{H} $(\mu(p) - \nu(p))$ copies of an arbitrary tree β of smoke nodes of \widehat{H} rooted in $(p, 1)$. This tree exists because $p \in \|\mu\|$, and so $p \in \|\nu\|$. The addition of the copies of β maintains the “no smoke without fuel” property, because it was already fulfilled in \widehat{H} by the nodes of β .

The smoke nodes of $\widehat{H}'(1)$ thus constructed are labelled by μ , and fuel and cargo nodes are neither added nor removed. We prove that \widehat{H}' is realizable. Let $t_1^{k_1} \dots t_{h-1}^{k_{h-1}}$ be a sequence that realizes \widehat{H} , for some transitions $t_1, \dots, t_{h-1} \in T$ and numbers $k_1, \dots, k_{h-1} \geq 0$.

Removing trees of smoke nodes from \widehat{H}' does not affect realizability: if there is a smoke p_o -node labeled by the observed place of t_i in some level $\widehat{H}(i)$ with a child labeled the same in $\widehat{H}(i+1)$, then there is also a pair of such fuel p_o -node in $\widehat{H}(i)$ and $\widehat{H}(i+1)$ by property of smoke nodes. This pair of fuel nodes is still in \widehat{H}' because we only remove trees of smoke nodes. Removing the trees translates as decreasing the iterations of some transitions in the realizing sequence of \widehat{H} . The trees of smoke nodes that we add to \widehat{H}' also do not affect realizability: they only increase the iterations of the transitions in the realizing sequence, as in the proof of the Replacement Theorem. ◀