Decidability of branching bisimulation on normed commutative context-free processes*

Wojciech Czerwiński, Piotr Hofman and Sławomir Lasota

Institute of Informatics, University of Warsaw
{wczerwin,ph209519,sl}@mimuw.edu.pl

Abstract. We investigate normed commutative context-free processes (Basic Parallel Processes). We show that branching bisimilarity admits the *small response property*: in the Bisimulation Game, Duplicator always has a response leading to a process of size linearly bounded with respect to the Spoiler's process. The linear bound is effective, which leads to decidability of branching bisimilarity. For weak bisimilarity, we are able merely to show existence of some linear bound, which is not sufficient for decidability. We conjecture however that the same effective bound holds for weak bisimilarity as well. We believe that further elaboration of novel techniques developed in this paper may be sufficient to demonstrate decidability.

1 Introduction

Bisimulation equivalence (bisimilarity) is a fundamental notion of equivalence of processes, with many natural connections to logic, games and verification [10, 13]. This paper is a continuation of the active line of research focusing on decidability and complexity of decision problems for bisimulation equivalence on various classes of infinite systems [12].

We investigate the class of commutative context-free processes, known also under name Basic Parallel Processes (BPP) [1]. By this we mean the labeled graphs induced by context-free grammars in Greibach normal form, with a proviso that non-terminals appearing on the right-hand side of a productions are assumed to be commutative. For instance, the production $X \longrightarrow a\,Y Z$, written

$$X \xrightarrow{a} Y Z,$$

says that X performs an action a and then executes Y and Z in parallel. Formally, the right-hand side is a multiset rather than a sequence.

Over this class of graphs, we focus on bisimulation equivalence as the primary type of semantic equality of processes. It is known that *strong* bisimulation equivalence is decidable [2] and PSPACE-complete [11, 8]; and is polynomial

* The first author acknowledges a partial support by the Polish MNiSW grant N N206 568640. The remaining two authors acknowledge a partial support by the Polish MNiSW grant N N206 567840.

for *normed* processes [6]. Dramatically less is known about weak bisimulation equivalence, that abstracts from the silent ε-transitions: we only know that it is semi-decidable [3] and that it is decidable in polynomial space over a very restricted class of *totally normed* processes [4]. The same applies to branching bisimulation equivalence, a variant of weak bisimulation that respects faithfully branching of equivalent processes. The only non-trivial decidability result known by now for weak bisimulation equivalence is [14], it applies however to a very restricted subclass.

During last two decades decidability of weak bisimulation over context-free processes became an established long-standing open problem. This paper is a significant step towards solving this problem in confirmative.

It is well known that bisimulation equivalences have an alternative formulation, in terms of Bisimulation Game played between Spoiler (aiming at showing non-equivalence) and Duplicator (aiming at showing equivalence) [13]. One of the main obstacles in proving decidability of weak (or branching) bisimulation equivalence is that Duplicator may do arbitrarily many silent transitions in a single move, and thus the size of the resulting process is hard to bound.

In this paper we investigate branching bisimilarity over normed commutative context-free processes. Our main technical result is the proof of the following *small response property*, formulated as Theorem 1 in Section 3: if Duplicator has a response, then he also has a response that leads to a process of size linearly bounded with respect to the other (Spoiler's) process. Importantly, we obtain an effective bound on the linear coefficient, which enables us to prove (Theorem 2) decidability of branching bisimulation equivalence. The proof of Theorem 1 is quite complex and involves a lot of subtle investigations of combinatorics of BPP transitions, the main purpose being elimination of unnecessary silent transitions.

A major part of the proof works for weak bisimulation equally well (and, as we believe, also for any reasonable equivalence that lies between the two equivalences). However, for weak bisimulation we can merely show *existence* of the linear coefficient witnessing the small response property, while we are not able to obtain any effective bound. Nevertheless we strongly believe (and conjecture) that a further elaboration of our approach will enable proving decidability of weak bisimulation. We plan to pursue this as a future work. In particular, we actually reprove decidability in the subclass investigated in [14].

2 Preliminaries

The commutative context-free processes (known also as Basic Parallel Processes) are determined by the following ingredients (called a *process definition*): a finite set $V = \{X_1, \ldots, X_n\}$ of variables, a finite set A of letters, and a finite set T of transition rules, each of the form $X \xrightarrow{\zeta} \alpha$ where X is a variable, $\zeta \in A \cup \{\varepsilon\}$ and α is a finite multiset of variables.

A *process*, is any finite multiset of variables, thus of the form $X_1^{a_1} \ldots X_n^{a_n}$, and may be understood as the parallel composition of a_1 copies of X_1, ... , and a_n copies of X_n. In particular the *empty process*, denoted ε, when $a_1 = \cdots = a_n = 0$.

For any $W \subseteq V$ we denote by $W^{\otimes}$ the set of all processes where only variables from W occur, that is, $W^{\otimes}$ is the set of all finite multisets over W.

By $\alpha\beta$ we mean the composition of processes α and β, understood as the multiset union. The behavior, i.e., the *transition relation*, is defined by the following extension rule:

$$\text{if } X \xrightarrow{\zeta} \alpha \in T \text{ then } X\beta \xrightarrow{\zeta} \alpha\beta, \text{ for any } \beta \in V^{\otimes}.$$

Remark 1. Commutative context-free processes are precisely labeled communication free Petri nets, where the places are variables and transitions $X \xrightarrow{\zeta} \alpha$ are firing rules. A process $X_1^{a_1} \ldots X_n^{a_n}$ represents the marking with a_i tokens on the place X_i.

The transition relation $\xrightarrow{\varepsilon}$ models silent steps and will be written $\longrightarrow$. We write $\alpha \Longrightarrow \beta$ if a process β can be reached from α by a sequence of $\xrightarrow{\varepsilon}$ transitions. To simplify definitions, we assume that $\alpha \longrightarrow \alpha$ for any α.

Definition 1. *A binary symmetric relation B over processes is a* branching bisimulation *iff for every pair α B β and $\zeta \in A \cup \{\varepsilon\}$ satisfies: if $\alpha \xrightarrow{\zeta} \alpha'$ then $\beta \Longrightarrow \beta'' \xrightarrow{\zeta} \beta'$ such that α B β'' and α' B β'.*

We say that two processes α and β are branching bisimilar, denoted $\alpha \approx \beta$, if there exists a branching bisimulation B such that α B β.

In the proofs we will use the characterization of bisimilarity in terms of Bisimulation Game [10, 13]. The game is played by two players, Spoiler and Duplicator, over an arena consisting of all pairs of processes, and proceeds in rounds. Each round starts with a Spoiler's move followed by a Duplicator's response. In position (α, β), Spoiler chooses one of processes, say α, and one transition $\alpha \xrightarrow{\zeta} \alpha'$. As a response, Duplicator has to do a sequence of transitions of the form $\beta \Longrightarrow \beta'' \xrightarrow{\zeta} \beta'$, and then Spoiler chooses whether the play continues from (α, β'') or (α', β').

If one of players gets stuck, the other wins. Otherwise the play is infinite and in this case it is Duplicator who wins. A well-known fact is that two processes are branching bisimilar iff Duplicator has a winning strategy in the game that starts in these two processes.

For the rest of this paper we assume that each variable X has a sequence of transitions $X \xrightarrow{\zeta_1} \ldots \xrightarrow{\zeta_m} \varepsilon$ leading to the empty process. A process definition that fulfills this requirement is usually called *normed*. By the *norm* of X, denoted $\operatorname{norm}(X)$, we mean the smallest possible number of visible transitions that appears in some sequence as above. Formally speaking, the norm of X is the length of the shortest word $a_1 \ldots a_n \in A^*$ such that

$$X \Longrightarrow \xrightarrow{a_1} \Longrightarrow \ldots \Longrightarrow \xrightarrow{a_n} \Longrightarrow \varepsilon.$$

We additively enhance the definition of norm to processes and write $\operatorname{norm}(\alpha)$ for any $\alpha \in V^{\otimes}$. Note that the norm is *weak* in the sense that silent transitions do not count.

3 Decidability via small response property

It was known before that branching bisimilarity is semi-decidable [3]. A main obstacle for a semi-decision procedure for inequivalence is that commutative context-free processes are not image finite with respect to branching bisimilarity: a priori Duplicator has infinitely many possible responses to a Spoiler's move. The main insight of this paper is that commutative context-free processes *are* essentially image-finite, in the following sense. Define the size of a process as its multiset cardinality: $\text{size}(X_1^{a_1} \dots X_n^{a_n}) = a_1 + \cdots + a_n$. Then Duplicator has always a response of size bounded linearly with respect to a Spoiler's process (cf. Theorem 1 below).

Definition 2. *Let $c \in \mathbb{N}$. By a c-branching bisimulation we mean a relation B defined as in Definition 1 with the additional requirement*

$$\mathit{size}(\beta'), \mathit{size}(\beta'') \le c \cdot \mathit{size}(\alpha'). \tag{1}$$

Let the size d of a process definition be the sum of lengths of all production rules. Our main technical result is an efficient estimation of c, with respect to d:

Theorem 1 (small response property). *For each normed process definition of size d with n variables, branching bisimilarity $\approx$ is a $(2d^{n-1} + d)$-branching bisimulation.*

The proof of Theorem 1 is deferred to Sections 4–6.

In consequence, a Spoiler's winning strategy, seen as a tree, becomes finitely branching. This observation leads directly to decidability:

Theorem 2. *Branching bisimilarity $\approx$ is decidable over normed commutative context-free processes.*

Proof. We sketch two semi-decision procedures (along the lines of [9]): one for branching bisimilarity and the other for $(2d^{n-1} + d)$-branching bisimilarity.

For the positive side we use a standard semi-linear representation, knowing that each congruence, including $\approx$, is semi-linear [5, 7]. The algorithm guesses a base-period representation of a semi-linear set and then checks validity of a Presburger formula that says that this set is a branching bisimulation containing the input pair of processes.

For the negative side, we observe that due to Theorem 1 Duplicator has only finitely many possible answers to each Spoiler's move. Thus, if Spoiler wins then its winning strategy may be represented by a finitely-branching tree. Furthermore, by König Lemma this tree is finite. The algorithm thus simply guesses a finite Spoiler's strategy. This can be done effectively: for given β, β', β'' and ζ it is decidable if $\beta \Longrightarrow_0 \beta'' \xrightarrow{\zeta} \beta'$, as the $\Longrightarrow_0$ relation is effectively semilinear [3]. $\square$

Proof strategy. The rest of this paper is devoted to the proof of Theorem 1. Consider a fixed normed process definition from now on. In Section 4 we define a notion of normal form nf(α) for a process α and provide linear lower and upper bounds on its size:

$$\text{size}(\alpha) \leq \text{size}(\text{nf}(\alpha)) \leq c \cdot \text{size}(\alpha) \tag{2}$$

(the lower bound holds assumed that α is minimal wrt. multiset inclusion in its bisimulation class). However, the linear coefficient c is not bounded effectively. The computable estimation of the coefficient is derived in Section 5. Finally, in Section 6 we show how the bounds (2) are used to prove Theorem 1. Due to space limitations we omit some proofs in Sections 4–6.

As observed e.g. in [14], a crucial obstacle in proving decidability is so called *generating* transitions of the form $X \longrightarrow XY$, as they may be used by Duplicator to reach silently XY^m for arbitrarily large m. A great part of our proofs is an analysis of combinatorial complexity of generating transitions and, roughly speaking, elimination of 'unnecessary' generations.

Weak bisimilarity. Branching bisimilarity is more discriminating than the well known weak bisimilarity. The whole development of Section 4 is still valid if weak bisimilarity is considered in place of branching bisimilarity. Furthermore, except one single case, the entire proof of estimation of the coefficient in Section 5 remains valid too. Interestingly, this single case is obvious under assumptions of [14], thus our proof remains valid for weak bisimilarity over the subclass studied there. We conjecture that the single missing case is provable for weak bisimilarity and thus Theorem 1 holds just as well. This would imply decidability.

4 Normal form by squeezing

In the sequel we often implicitly use the well-known fact that branching bisimilarity is substitutive, i.e., $\alpha \approx \beta$ implies $\alpha\gamma \approx \beta\gamma$.

In this section we develop a framework useful for the proof of Theorem 1 in the following sections. We define a normal form nf(α) of a process α that identifies the bisimulation class of α uniquely. Moreover, we provide estimations of the size of nf(α) relative to the size of α, from both sides, in Corollary 1 and Lemma 11, which culminate this section.

A transition $\alpha \xrightarrow{\zeta} \beta$ is *norm preserving* if $|\alpha| = |\beta|$ and *norm reducing* if $|\alpha| = |\beta| + 1$. In the sequel we will pay special attention to norm preserving ε-transitions. Therefore we write $\alpha \longrightarrow_0 \beta$, respectively $\alpha \Longrightarrow_0 \beta$, to emphasize that the transitions are norm preserving.

Lemma 1. *If $\alpha \Longrightarrow_0 \beta \Longrightarrow_0 \alpha'$ and $\alpha \approx \alpha'$ then $\beta \approx \alpha$.*

We call the transition $\alpha \xrightarrow{\zeta} \beta$ *decreasing* if either $\zeta \in A$ and the transition is norm-reducing; or $\zeta = \varepsilon$ and the transition is norm preserving. Note that every variable has a sequence of decreasing transitions leading to the empty process ε.

Lemma 2 (decreasing response). *Whenever $\alpha \approx \beta$ and $\alpha \stackrel{\zeta}{\longrightarrow} \alpha'$ is decreasing then any Duplicator's matching sequence of transitions from β contains exclusively decreasing transitions.*

Due to Lemma 1, instantiated to single variables, we may assume wlog. that there are no two distinct variables X, Y with $X \Longrightarrow_0 Y \Longrightarrow_0 X$. Indeed, since reachability via the $\Longrightarrow_0$ transitions is decidable [3], in a preprocessing one may eliminate such pairs X, Y. Relying on this assumption, we may define a partial order induced by decreasing transitions.

Definition 3. *Let $X >_0 Y$ if there is a sequence of decreasing transitions leading from X to Y. Let $>$ denote an arbitrary fixed total order extending $>_0$.*

In the sequel we assume that there are n variables, ordered $X_1 > X_2 > \ldots > X_n$. Directly from the definition of $>$ we deduce:

Lemma 3 (decreasing transition). *If a decreasing transition $X_1^{a_1} \ldots X_n^{a_n} \stackrel{\zeta}{\longrightarrow} X_1^{b_1} \ldots X_n^{b_n}$ is performed by X_k, say, then $b_1 = a_1, \ldots, b_{k-1} = a_{k-1}$.*

Consider a norm preserving silent transition $X \longrightarrow_0 \delta$. If X appears in δ, i.e. $\delta = X\bar{\delta}$, we call the transition *generating*. We use the name *generating* also for a general transition $\alpha \longrightarrow_0 \beta$ as a single transition is always performed by a single variable.

Lemma 4 (decreasing transition cont.). *If a decreasing transition as in Lemma 3 is not generating then $b_k = a_k - 1$.*

Following [14], we say that X *generates* Y if $X \Longrightarrow_0 XY$. Thus if $X \longrightarrow X\bar{\delta}$ then X generates every variable that appears in $\bar{\delta}$. In particular, X may generate itself. Note that each generated variable is of norm 0. More generally, we say that α generates β if $\alpha \Longrightarrow_0 \alpha\beta$. This is the case precisely iff every variable occurring in β is generated by some variable occurring in α. As a direct corollary of Lemma 1 we get ($\sqsubseteq$ stands for the multiset inclusion of processes):

Lemma 5. *If α generates β then $\alpha \approx \alpha\bar{\beta}$ for any $\bar{\beta} \sqsubseteq \beta$.*

Lemma 5 will be especially useful in the sequel, as a tool for eliminating unnecessary transitions and thus decreasing the size of a resulting process.

A process $X_1^{a_1} \ldots X_n^{a_n}$ may be equivalently presented as a sequence of exponents $(a_1 \ldots a_n) \in \mathbb{N}^n$. In this perspective, $\sqsubseteq$ is the point-wise order. The sequence presentation $(a_1 \ldots a_n) \in \mathbb{N}^n$ induces additionally the lexicographic order on processes, denoted $\preceq$. We will exploit the fact that this order is total, and thus each bisimulation class exhibits the least element. (A *bisimulation class* of a process α is the set of all processes β with $\beta \approx \alpha$.)

The sequence presentation allows us to speak naturally of *prefixes* of a process: the *k-prefix* of $X_1^{a_1} \ldots X_n^{a_n}$ is the process $X_1^{a_1} \ldots X_k^{a_k}$, for $k = 0 \ldots n$.

We now go to one of the crucial notions used in the proof: *unambiguous processes* and their *greatest extensions.*

Definition 4 (unambiguous processes). *A process $X_1^{a_1} \dots X_n^{a_n}$, is called* k-unambiguous *if for every $1 \le i \le k$, $\alpha, \beta \in \{X_{i+1}, \dots, X_n\}^{\otimes}$ and $b, c \in \mathbb{N}$ such that*

$$X_1^{a_1} X_2^{a_2} \dots X_{i-1}^{a_{i-1}} X_i^b \alpha \approx X_1^{a_1} X_2^{a_2} \dots X_{i-1}^{a_{i-1}} X_i^c \beta$$

we have either $b, c \ge a_i$ or $b = c$. When $k = n$ we write simply unambiguous*.*

Note that being k-unambiguous is a property of the k-prefix: a process is k-unambiguous iff its k-prefix is so.

Example 1. Consider following process definition:

$$\begin{array}{lll} X_1 \xrightarrow{a} X_1 & X_2 \xrightarrow{b} X_3 & X_3 \xrightarrow{b} \varepsilon \\ X_1 \longrightarrow \ \varepsilon & X_2 \longrightarrow X_3 & X_3 \longrightarrow \varepsilon \end{array}$$

and an order $X_1 > X_2 > X_3$ on variables. We observe that $X_1^2 \approx X_1$, therefore the process X_1^2 is not (1-)unambiguous. On the other hand $X_1 \not\approx \alpha$ for any $\alpha \in \{X_2, X_3\}^{\otimes}$ (because neither X_2 nor X_3 can perform an a transition), so X_1 is unambiguous. Furthermore $X_1 X_2 \approx X_1 X_3^2$, hence $X_1 X_2$ is not (2-)unambiguous. Finally we observe that $X_1 X_3^2 \not\approx X_1 X_3$. Therefore $X_1 X_3^2$ is unambiguous, but also $X_1 X_3$ is so. □

Note that a prefix of a k-unambiguous process is k-unambiguous as well. Moreover, k-unambiguous processes are downward closed wrt. $\sqsubseteq$: whenever $\alpha \sqsubseteq \beta$ and β is k-unambiguous, then α is k-unambiguous as well.

Directly by Definition 4, if $\gamma = X_1^{a_1} \dots X_{k-1}^{a_{k-1}}$ is $(k-1)$-unambiguous then it is automatically k-unambiguous (in fact j-unambiguous for any $j \ge k$). According to the sequence presentation, this corresponds to extending the process $(a_1 \dots a_{k-1})$ with $a_k = 0$. We will be especially interested in the greatest value of a_k possible, as formalized in the definition below.

Definition 5 (the greatest extension). *The* greatest k-extension *of a $(k-1)$-unambiguous process $\gamma = X_1^{a_1} \dots X_{k-1}^{a_{k-1}} \in \{X_1 \dots X_{k-1}\}^{\otimes}$ is that process among k-unambiguous processes γX_k^a that maximizes a.*

Clearly the greatest extension does not need exist in general, as illustrated below.

Example 2. Consider the processes from Example 1. The process X_1 is the greatest 1-extension of the empty process as X_1^2 is not 1-unambiguous. X_1 is also its own greatest 2-extension. Furthermore, X_1 does not have the greatest 3-extension. Indeed, $X_1 X_3^a$ is not bisimilar to $X_1 X_3^b$, for $a \ne b$, therefore $X_1 X_3^a$ is 3-unambiguous for any a. □

Definition 6 (unambiguous prefix). *By an* unambiguous prefix *of a process $X_1^{a_1} \dots X_n^{a_n}$ we mean any k-prefix $X_1^{a_1} \dots X_k^{a_k}$ that is k-unambiguous, for $k = 0 \dots n$. The* maximal unambiguous prefix *is the one that maximizes k.*

Example 3. For the process definition from Example 1, the maximal unambiguous prefix of $X_1 X_2^2$ is X_1, and the maximal unambiguous prefix of $X_1^2 X_2$ is the empty process. □

The following lemma is a crucial observation underlying our subsequent development.

Lemma 6. *Let* $\gamma \in \{X_1 \dots X_{k-1}\}^{\otimes}$ *be* $(k-1)$*-unambiguous and assume that* γX_k^a *is its greatest* k*-extension. Let* $b > a$ *and let* $\alpha, \beta \in \{X_{k+1}, \dots, X_n\}^{\otimes}$ *be arbitrary processes such that*

$$\gamma\, X_k^b\, \beta \approx \gamma\, X_k^a\, \alpha.$$

Then for any decreasing transition $X_k^b\, \beta \stackrel{\zeta}{\longrightarrow} X_k^{b'}\, \beta'$*, that gives rise to a Spoiler's move*

$$\gamma\, X_k^b\, \beta \stackrel{\zeta}{\longrightarrow} \gamma\, X_k^{b'}\, \beta'$$

there are some $\alpha', \alpha'' \in \{X_{k+1}, \dots, X_n\}^{\otimes}$ *and a sequence* $\alpha \Longrightarrow_0 \alpha'' \stackrel{\zeta}{\longrightarrow} \alpha'$ *of transitions that gives rise to a Duplicator's response*

$$\gamma\, X_k^a\, \alpha \Longrightarrow_0 \gamma\, X_k^a\, \alpha'' \stackrel{\zeta}{\longrightarrow} \gamma\, X_k^a\, \alpha',$$

as required by Definition 1.

Note 1. According to the assumptions, γX_k^a is an unambiguous prefix of $\gamma X_k^a \alpha$. The crucial consequence of the lemma is that Duplicator has a response that preserves γX_k^a being a prefix, as only α is engaged in the response.

Proof. Consider a Duplicator's response (all transition are necessarily decreasing by Lemma 2):

$$\gamma\, X_k^a\, \alpha \Longrightarrow_0 \gamma''\, X_k^{a''}\, \alpha'' \stackrel{\zeta}{\longrightarrow} \gamma'\, X_k^{a'}\, \alpha' \tag{3}$$

where $\gamma', \gamma'' \in \{X_1 \dots X_{k-1}\}^{\otimes}$ and $\alpha', \alpha'' \in \{X_{k+1}, \dots, X_n\}^{\otimes}$. Wlog. we may assume that

$$\gamma''\, X_k^{a''}\, \alpha'' \not\approx \gamma'\, X_k^{a'}\, \alpha' \tag{4}$$

as otherwise lemma holds trivially. A fast observation is that

$$\gamma\, X_k^a \preceq \gamma'\, X_k^{a'}. \tag{5}$$

Indeed, suppose $\gamma'\, X_k^{a'} \prec \gamma\, X_k^a$. Knowing $\gamma\, X_k^{b'} \beta' \approx \gamma'\, X_k^{a'} \alpha'$ and $b' \geq a$ we get to a contradiction with the fact that $\gamma\, X_k^a$ is k-unambiguous.

Our aim is to demonstrate that Duplicator has a matching response (3) that uses only transition rules of variables $X_{k+1} \dots X_n$; in particular, by Lemma 3 this will imply $\gamma'\, X_k^{a'} = \gamma\, X_k^a$. We will describe below a transformation of the Duplicator's response to the required form.

Assume that some of variables $X_1 \dots X_k$ was engaged in (3) and let X_i be the greatest of them wrt. $>$. By (5) and by Lemma 4 we learn that at least one of transitions performed by some X_i must be generating, say

$$X_i \longrightarrow X_i \delta. \tag{6}$$

We will show how to remove one of these transitions from (3) but still preserve the bisimulation class of processes appearing along (3), and thus keep satisfying the requirements of Definition 1.

All variables that appear in δ are necessarily of norm 0, and thus they may participate later in the sequence (3) only with further norm preserving ε-transitions. Informally speaking, we consider the tree of norm preserving ε-transitions initiated by (6), that are performed along (3), say:

$$X_i \Longrightarrow_0 X_i^j \delta', \tag{7}$$

for some $j \geq 0$ and $\delta' \in \{X_{i+1} \dots X_n\}^{\otimes}$.

Formally, the sequence (7) is defined by the following coloring argument. As a process may contain many occurrences of the same variable we consider variable occurrences as independent entities. Assume that every variable occurrence in $\gamma\, X_k^a$ has been initially colored by a unique color. Assume further that colors are inherited via transitions: every transition in (3) is colored with the color of the occurrence of its left-hand side variable that is engaged; and likewise are colored all the right-hand side variables occurrences. The sequence (7) contains all transitions colored with the color of (6).

The sequence (7) forms a subsequence of (3). There can be many such sequences, but at least one witnesses $j > 0$, by (5) and by the choice of X_i as the greatest wrt. $>$. Let us focus on removing this particular subsequence from (3).

As δ' is generated by X_i, by Lemma 5 we obtain $X_i \approx X_i^j \delta'$. By our assumption (4) we deduce that the sequence (7) can not contain the last transition of (3). Thus, by substitutivity of $\approx$, the sequence (3), after removing transitions (7), yields a process bisimulation equivalent to that yielded by (3). By continuing in the same manner we arrive finally at the Duplicator's response that does not engage variables $X_1 \dots X_k$ at all. This completes the proof. □

Lemma 7 (squeezing out). *Let $\gamma \in \{X_1 \dots X_{k-1}\}^{\otimes}$ be $(k-1)$-unambiguous and assume that γX_k^a is its greatest k-extension. Then for some $\delta \in \{X_{k+1} \dots X_n\}^{\otimes}$ it holds:*

$$\gamma\, X_k^{a+1} \approx \gamma\, X_k^a\, \delta. \tag{8}$$

Definition 7. *If a $(k-1)$-unambiguous process $\gamma \in \{X_1 \dots X_{k-1}\}^{\otimes}$ has the greatest k-extension, say γX_k^a, then the variable X_k is called* γ-squeezable *and any $\delta \in \{X_{k+1} \dots X_n\}^{\otimes}$ satisfying* (8) *is called a* γ-squeeze *of X_k.*

By the very definition, X_k has a γ-squeeze only if it is γ-squeezable. Lemma 7 shows the opposite: a γ-squeezable X_k has a γ-squeeze, that may depend in general on γ and k. The squeeze is however not uniquely determined and in fact X_k may admit many different γ-squeezes. In the sequel assume that for each $(k-1)$-unambiguous $\gamma \in \{X_1 \dots X_{k-1}\}^{\otimes}$ and X_k, some γ-squeeze of X_k is chosen; this squeeze will be denoted by $\delta_{k,\gamma}$.

Definition 8 (squeezing step). *For a given process α, assuming it is not n-unambiguous, let γ be its maximal unambiguous prefix. Thus there is $k \leq n$ such that*

$$\alpha = \gamma\, X_k^a\, \delta,$$

$\gamma \in \{X_1 \dots X_{k-1}\}^{\otimes}$, $\delta \in \{X_{k+1} \dots X_n\}^{\otimes}$, and $\gamma\, X_k^a$ is not k-unambiguous. Note that a is surely greater than 0. We define squeeze(α) by

$$\text{squeeze}(\alpha) = \gamma\, X_k^{a-1}\, \delta_{k,\gamma}\, \delta.$$

Otherwise, i.e. when α is n-unambiguous, for convenience put squeeze$(\alpha) = \alpha$.

By Lemma 7 and by substitutivity of $\approx$ we conclude that $\alpha \approx \text{squeeze}(\alpha)$ and if α is not unambiguous then $\text{squeeze}(\alpha) \prec \alpha$.

We have the following characterization of unambiguous processes:

Lemma 8. *A process α is n-unambiguous if and only if it is the least one in its bisimulation class wrt. $\preceq$.*

Lemma 7, applied in a systematic manner sufficiently many times on a process α, yields a kind of normal form, as stated in Lemma 9 below. A process α we call shortly $\sqsubseteq$-minimal if there is no $\beta \sqsubset \alpha$ with $\beta \approx \alpha$.

Definition 9 (normal form). *For any process α let nf(α) denote the unambiguous process obtained by consecutive alternating applications of the following two steps:*

- *the squeezing step: replace α by squeeze(α),*
- *the $\sqsubseteq$-minimization step: replace α by any $\sqsubseteq$-minimal $\bar{\alpha} \sqsubseteq \alpha$ with $\bar{\alpha} \approx \alpha$.*

As $\alpha \approx \text{squeeze}(\alpha)$ then $\alpha \approx \text{nf}(\alpha)$ and thus using Lemma 8 we conclude that bisimulation equivalence is characterized by syntactic equality of normal forms:

Lemma 9. *$\alpha \approx \beta$ if and only if nf$(\alpha) =$ nf(β).*

Finally we are able to formulate lower and upper bounds on the size of $\text{nf}(\alpha)$, with respect to the size of α, that will be crucial for the proof of Theorem 1. The first one applies uniquely to $\sqsubseteq$-minimal processes.

Lemma 10. *If α is $\sqsubseteq$-minimal then size$(\alpha) \leq$ size$(\bar{\alpha})$, for any $\bar{\alpha} \sqsubseteq$ squeeze(α) such that $\bar{\alpha} \approx$ squeeze(α).*

Corollary 1 (lower bound). *If α is $\sqsubseteq$-minimal then size(nf$(\alpha)) \geq$ size(α).*

Lemma 11 (upper bound). *There is a constant c, depending only on the process definition, such that size(nf$(\alpha)) \leq c \cdot$ size(α) for any process α.*

Concerning the upper bound, in the following section we demonstrate a sharper result, with the constant c estimated effectively.

5 Small normal form

Denote the size of the process definition by d.

Lemma 12 (upper bound). *For any α, size(nf$(\alpha)) \leq d^{n-1} \cdot$ size(α).*

Lemma 12 follows immediately from Lemma 13 that says that squeezing does not increase a weighted measure of size, defined as:

$$d\text{-size}(X_1^{a_1} \dots X_n^{a_n}) = a_1 \cdot d^{n-1} + a_2 \cdot d^{n-2} + \dots + a_{n-1} \cdot d + a_n.$$

Lemma 13. *For every k and $(k-1)$-unambiguous $\gamma \in \{X_1 \dots X_{k-1}\}^{\otimes}$, if X_k is γ-squeezable then it has a γ-squeeze δ with d-size$(\delta) \leq d$-size(X_k).*

Indeed, Lemma 13 implies d-size(nf(α)) $\leq d$-size(α) and then Lemma 12 follows:

$$\text{size}(\text{nf}(\alpha)) \leq d\text{-size}(\text{nf}(\alpha)) \leq d\text{-size}(\alpha) \leq d^{n-1} \cdot \text{size}(\alpha).$$

Before embarking on the proof of Lemma 13, we formulate a slight generalization of Lemma 6 from Section 4. For two processes $\alpha, \beta \in \{X_1 \dots X_l\}^{\otimes}$ we say that α *is l-dominating* β if α is bisimilar to some $\alpha' \sqsupseteq \beta$.

Lemma 14. *Let α be an arbitrary process, $\beta_1 \in \{X_1 \dots X_l\}^{\otimes}$ be m-unambiguous and $\beta_2 \in \{X_{l+1} \dots X_n\}^{\otimes}$ such that $\alpha \approx \beta_1\beta_2$. Let $\alpha \xrightarrow{\zeta} \alpha'$ be an arbitrary decreasing transition such that the l-prefix of α' is l-dominating β_1. Then there is a sequence of transitions $\beta_2 \Longrightarrow_0 \beta_2'' \xrightarrow{\zeta} \beta_2'$ that gives rise to a Duplicator's response*

$$\beta_1\beta_2 \Longrightarrow_0 \beta_1\beta_2'' \xrightarrow{\zeta} \beta_1\beta_2',$$

as required by Definition 1.

Lemma 14 is proved in exactly the same way as Lemma 6. Recalling Lemma 6 observe that it is indeed a special case of Lemma 14: $\gamma X_k^{b'}$ is surely k-dominating γX_k^a as $b' \geq b - 1 \geq a$.

Now we return to the proof of Lemma 13, by induction on k. For $k = n$ it trivially holds. Fix $k < n$ and assume the lemma for all greater values of k. Fix a $(k-1)$-unambiguous $\gamma \in \{X_1 \dots X_{k-1}\}^{\otimes}$ and consider its greatest k-extension γX_k^a. The proof is split into three cases:

- $a > 0$,
- $a = 0$ and X_k has a γ-squeeze δ such that $X_k \Longrightarrow_0 \delta$,
- $a = 0$ and X_k has no γ-squeeze δ such that $X_k \Longrightarrow_0 \delta$.

In the rest of this section we prove the last case only. The other cases are omitted due to space limitations.

Simplifying assumption. Variables $X_{k+1} \dots X_n$ may be split into those generated by X_k, an those not generated by X_k. A simple but crucial observation is that the order $>$ on variables $X_{k+1} \dots X_n$ may be rearranged, without losing generality, so that all variables generated by X_k are smaller than all variables not generated by X_k. Clearly, if we provide a γ-squeeze of X_k for the rearranged order, it is automatically a γ-squeeze of X_k for the initial order.

Thus for some $l \geq k$ we know that variables $X_{l+1} \dots X_n$ are all generated by X_k, and all the remaining variables $X_{k+1} \dots X_l$ are not generated by X_k. To emphasize this we will write $[\alpha \cdot \beta]$ for the composition of α and β, instead of $\alpha\beta$, whenever we know that $\alpha \in \{X_{k+1} \dots X_l\}^{\otimes}$ and $\beta \in \{X_{l+1} \dots X_n\}^{\otimes}$.

Lemma 15. *No γ-squeeze of X_k contains a variable generated by X_k.*

Proof. Assume the contrary, that is,

$$\gamma\, X_k \approx \gamma\, \delta'\, Y, \tag{9}$$

with $\delta'\, Y \in \{X_{k+1} \dots X_n\}^{\otimes}$ and Y generated by X_k. Consider the Bisimulation Game for $\gamma X_k \approx \gamma\, \delta'\, Y$ and an arbitrary sequence of $\longrightarrow_0$ transitions $Y \Longrightarrow_0 \varepsilon$ from Y to the empty process ε, giving rise to the sequence of Spoiler's moves

$$\gamma\, \delta'\, Y \Longrightarrow_0 \gamma\, \delta'.$$

By Lemma 14 we know that there is a Duplicator's response that does not engage γ at all:

$$\gamma\, X_k \Longrightarrow_0 \gamma\, \omega,$$

i.e. $X_k \Longrightarrow_0 \omega$. Now substituting $\gamma\omega$ in place of $\gamma\delta'$ in (9) we obtain a γ-squeeze of X_k

$$\gamma\, X_k \approx \gamma\, \omega\, Y,$$

such that $X_k \longrightarrow_0 X_k Y \Longrightarrow_0 X_k\, \omega\, Y$. This is in contradiction with the assumption that no γ-squeeze is reachable from X_k by $\Longrightarrow_0$. Thus the claim is proved. □

Using Lemma 15 we deduce that the normal form $\text{nf}(\gamma\, X_k) = \gamma\, \delta$ contains no variable generated by X_k, i.e., $\text{nf}(\gamma\, X_k) = \gamma\, [\delta \cdot \varepsilon]$. We will show that the weighted size of δ satisfies the required bound.

Consider the Bisimulation Game for $\gamma\, X_k \approx \gamma\, \delta$ and the Spoiler's move from the smallest variable occurring in δ wrt. $>$, say X_m. Process δ contains no variable generated by X_k, hence $m \leq l$. Thus $\delta = \delta'\, X_m$, and let the Spoiler's move be induced by a decreasing non-generating transition $X_m \xrightarrow{\zeta} \omega$:

$$\gamma\, \delta'\, X_m \xrightarrow{\zeta} \gamma\, \delta'\, \omega.$$

By Lemma 14 we know that there is a Duplicator's response that does not engage γ. As no γ-squeeze of X_k is reachable from X_k by $\Longrightarrow_0$, the response has necessarily the following form

$$\gamma\, X_k \Longrightarrow_0 \gamma\, X_k\, \eta \xrightarrow{\zeta} \gamma\, \sigma\, \eta,$$

where η is generated by X_k:

$$X_k \Longrightarrow_0 X_k\, \eta \quad \text{and} \quad X_k \xrightarrow{\zeta} \sigma,$$

as otherwise at some point in the $\Longrightarrow_0$ sequence a γ-squeeze would appear. We obtain $\gamma\, \sigma\, \eta \approx \gamma\, \delta'\, \omega$ and thus

$$\text{nf}(\gamma\, \sigma\, \eta) \;=\; \text{nf}(\gamma\, \delta'\, \omega). \tag{10}$$

From the last equality we will deduce how the sizes of $\mathrm{nf}(\gamma\,\sigma)$ and $\mathrm{nf}(\gamma\,\delta')$ are related, in order to conclude that the weighted size of δ is as required.

Let's inspect the l-prefix of the left processes in (10). Process η can not contribute to that prefix of the normal form, thus if we restrict to the l-prefixes we have the equality

$$l\text{-prefix}(\mathrm{nf}(\gamma\,\sigma\,\eta)) \;=\; l\text{-prefix}(\mathrm{nf}(\gamma\,\sigma)). \tag{11}$$

Similarly, let's inspect the m-prefix of the right process in (10). Again, ω can not contribute to that prefix of the normal form, thus if we restrict to the m-prefixes we have the equality

$$m\text{-prefix}(\mathrm{nf}(\gamma\,\delta'\,\omega)) \;=\; m\text{-prefix}(\mathrm{nf}(\gamma\,\delta')).$$

As $\gamma\,\delta$ is the normal form, the process $\gamma\,\delta'$ is unambiguous and thus clearly $\mathrm{nf}(\gamma\,\delta') = \gamma\,\delta'$. Substitute this to the last equality above:

$$m\text{-prefix}(\mathrm{nf}(\gamma\,\delta'\,\omega)) \;=\; m\text{-prefix}(\gamma\,\delta') \;=\; \gamma\,\delta'. \tag{12}$$

Using induction assumption we obtain $d\text{-size}(\mathrm{nf}(\gamma\,\sigma)) \leq d\text{-size}(\gamma\,\sigma)$. As $m \leq l$, by (10), (11) and (12) we conclude that

$$d\text{-size}(\gamma\,\delta') \leq d\text{-size}(\mathrm{nf}(\gamma\,\sigma)) \leq d\text{-size}(\gamma\,\sigma)$$

and thus $d\text{-size}(\delta') \leq d\text{-size}(\sigma)$. By the last inequality together with $\mathrm{size}(\sigma) \leq d-1$ and $\sigma \in \{X_{k+1} \dots X_n\}^{\otimes}$ we get the required bound on weighted size of δ:

$$\begin{aligned} d\text{-size}(\delta) = d\text{-size}(\delta') + d\text{-size}(X_m) \leq d\text{-size}(\sigma) + d^{n-m} \leq \\ (d-1)\,d^{n-k-1} + d^{n-m} \leq d^{n-k} = d\text{-size}(X_k). \end{aligned}$$

6 Proof of the small response property

Now we show how Theorem 1 follows from the estimations given in Corollary 1 and Lemma 12. We will need a definition and two lemmas.

We write $\alpha \overset{\approx}{\Longrightarrow}_0 \beta$ if $\alpha \Longrightarrow_0 \beta$ and $\alpha \approx \beta$. A process α is called $\overset{\approx}{\Longrightarrow}_0$-*minimal* if there is no $\beta \prec \alpha$ with $\alpha \overset{\approx}{\Longrightarrow}_0 \beta$.

Lemma 16. *For any α there is a $\overset{\approx}{\Longrightarrow}_0$-minimal process $\bar{\alpha}$ with $\alpha \overset{\approx}{\Longrightarrow}_0 \bar{\alpha}$ of size bounded by $size(\bar{\alpha}) \leq size(nf(\alpha))$.*

Lemma 17. *If α is $\overset{\approx}{\Longrightarrow}_0$-minimal and $\alpha \overset{\approx}{\Longrightarrow}_0 \beta$ then $\alpha \sqsubseteq \beta$.*

Proof of Theorem 1. Consider $\alpha \approx \beta$, a Spoiler's move $\alpha \xrightarrow{\zeta} \alpha'$ and a Duplicator's response: $\beta \Longrightarrow_0 \beta_1 \xrightarrow{\zeta} \beta_2$, with $\alpha \approx \beta_1$ and $\alpha' \approx \beta_2$. The basic idea of the proof is essentially to eliminate some unnecessary generation done by transitions $\beta \Longrightarrow_0 \beta_1$.

As the first step we apply Lemma 16 to β, thus obtaining a sequence of transitions $\beta \Longrightarrow_0 \bar{\beta}$, for some $\overset{\approx}{\Longrightarrow}_0$-minimal process $\bar{\beta}$, in order to consider the pair $(\alpha, \bar{\beta})$ instead of (α, β). Knowing $\alpha \approx \bar{\beta}$ we obtain a Duplicator's response

$$\beta \Longrightarrow_0 \bar{\beta} \Longrightarrow_0 \beta_1' \xrightarrow{\zeta} \beta_2' \tag{13}$$

with $\alpha \approx \beta_1'$ and $\alpha' \approx \beta_2'$. Note that by Lemma 17 we know $\bar{\beta} \sqsubseteq \beta_1'$.

As the second step extend (13) with any sequence $\beta_2' \overset{\approx}{\Longrightarrow}_0 \bar{\beta}_2'$ leading to a $\sqsubseteq$-minimal process $\bar{\beta}_2' \sqsubseteq \beta_2'$. Our knowledge may be outlined with the following diagram (the subscript in $\Longrightarrow_0$ is omitted):

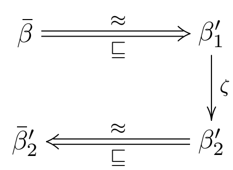

Both left-most processes in the diagram are size bounded. Indeed, Corollary 1 applied to $\bar{\beta}$ and $\bar{\beta}_2'$ yields

$$\text{size}(\bar{\beta}) \leq \text{size}(\text{nf}(\alpha)) \quad \text{and} \quad \text{size}(\bar{\beta}_2') \leq \text{size}(\text{nf}(\alpha')).$$

Then applying Lemma 12 to α and α' we obtain:

$$\text{size}(\bar{\beta}) \leq \text{size}(\alpha) \cdot d^{n-1} \quad \text{and} \quad \text{size}(\bar{\beta}_2') \leq \text{size}(\alpha') \cdot d^{n-1}. \tag{14}$$

As the third and the last step of the proof, we claim that β_1' and β_2' may be replaced by processes of size bounded, roughly, by the sum of sizes of $\bar{\beta}$ and $\bar{\beta}_2'$.

Claim. There are some processes $\beta_1'' \approx \beta_1'$ and $\beta_2'' \approx \beta_2'$ such that

$$\bar{\beta} \Longrightarrow_0 \beta_1'' \xrightarrow{\zeta} \beta_2'' \tag{15}$$

and

$$\text{size}(\beta_1''), \text{size}(\beta_2'') \leq \text{size}(\bar{\beta}) + \text{size}(\bar{\beta}_2') + d. \tag{16}$$

The claim is sufficient for Theorem 1 to hold, by inequalities (14). Thus to complete the proof we only need to demonstrate the claim. The idea underlying the proof of the claim is illustrated by the following diagram:

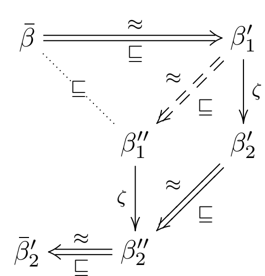

We use a coloring argument, similarly as in the proof of Lemma 6. Let us color uniquely every variable occurrence in β'_1 and let every transition preserve the color of the left-hand side variable. Obviously at most size($\bar{\beta}'_2$) of these colors will be still present in $\bar{\beta}'_2$, name them *surviving colors*. Let the $\beta'_1 \xrightarrow{\zeta} \beta'_2$ transition be performed due to a transition rule $X \xrightarrow{\zeta} \delta$, color this particular X, say, brown.

Let β''_1 consists of all variables which either belong to $\bar{\beta}$ or are colored surviving or brown color. Thus clearly $\bar{\beta} \sqsubseteq \beta''_1 \sqsubseteq \beta'_1$. One easily observes that after the brown transition $X \xrightarrow{\zeta} \delta$ from β''_1 we get β''_2 such that $\bar{\beta}'_2 \sqsubseteq \beta''_2 \sqsubseteq \beta'_2$, because all surviving colored variables are still present. By Lemma 1 one has $\beta''_1 \approx \beta'_1$ and $\beta''_2 \approx \beta'_2$.

Finally we obtain the size estimation $\text{size}(\beta''_1) \leq \text{size}(\bar{\beta}) + \text{size}(\bar{\beta}'_2) + 1$ as in β''_1 there can be at most $\text{size}(\bar{\beta}'_2) + 1$ surviving and brown colored variables that do not belong to $\bar{\beta}$. This easily implies the estimation for size(β''_2). □

Acknowledgments. We are grateful to the reviewers for valuable comments, in particular for encouraging us to restrict the paper to branching bisimilarity.

References

1. S. Christensen. *Decidability and Decomposition in process algebras.* PhD thesis, Dept. of Computer Science, University of Edinburgh, UK, 1993.
2. S. Christensen, Y. Hirshfeld, and F. Moller. Bisimulation equivalence is decidable for Basic Parallel Processes. In *CONCUR*, pages 143–157, 1993.
3. J. Esparza. Petri nets, commutative context-free grammars, and Basic Parallel Processes. *Fundam. Inform.*, 31(1):13–25, 1997.
4. S. Fröschle and S. Lasota. Normed processes, unique decomposition, and complexity of bisimulation equivalences. *Electr. Notes Theor. Comp. Sci.*, 239:17–42, 2009.
5. Y. Hirshfeld. Congruences in commutative semigroups. Technical report, University of Edinburgh, LFCS report ECS-LFCS-94-291, 1994.
6. Y. Hirshfeld, M. Jerrum, and F. Moller. A polynomial-time algorithm for deciding bisimulation equivalence of normed Basic Parallel Processes. *Mathematical Structures in Computer Science*, 6(3):251–259, 1996.
7. P. Jancar. Decidability questions for bismilarity of Petri nets and some related problems. In *STACS*, pages 581–592, 1994.
8. P. Jancar. Strong bisimilarity on Basic Parallel Processes is PSPACE-complete. In *LICS*, pages 218–227, 2003.
9. S. Lasota. Decidability of performance equivalence for Basic Parallel Processes. *Theoretical Computer Science*, 360:172–192, 2006.
10. R. Milner. *Communication and Concurrency.* Prentice Hall, 1995.
11. J. Srba. Strong bisimilarity and regularity of Basic Parallel Processes is PSPACE-hard. In *STACS*, pages 535–546, 2002.
12. J. Srba. *Roadmap of Infinite results*, volume Vol 2: Formal Models and Semantics. World Scientific Publishing Co., 2004.
13. C. Stirling. The joys of bisimulation. In *MFCS*, pages 142–151, 1998.
14. C. Stirling. Decidability of weak bisimilarity for a subset of Basic Parallel Processes. In *FoSSaCS*, pages 379–393, 2001.

A Proofs missing in Section 4

Proof of Lemma 1. Immediate using Definition 1. If Spoiler plays from α, Duplicator uses its response from α', precomposed with $\beta \Longrightarrow_0 \alpha'$. On the other hand, if Spoiler plays from β, Duplicator moves $\alpha \Longrightarrow_0 \beta$ and then copies the Spoiler's transition. $\Box$

Proof of Lemma 2. Follows from the following simple observations: $\approx$ is norm preserving; for $a \neq \varepsilon$, the transition relation $\xrightarrow{a}$ may decrease the norm by at most one; the transition relation $\xrightarrow{\varepsilon}$ never decreases the norm. $\Box$

Proof of Lemma 7. By δ, δ', etc. we denote below processes from $\{X_{k+1} \ldots X_n\}^{\otimes}$.

As a is the maximal extension of γ, there is some $b > a$ and some processes δ, δ' such that

$$\gamma\, X_k^b\, \delta \approx \gamma\, X_k^a\, \delta'.$$

Consider an arbitrary sequence of decreasing transitions

$$X_k^b\, \delta \xrightarrow{\zeta_1} \ldots \xrightarrow{\zeta_m} X_k^{a+1}.$$

By Lemma 6 there is a sequence of matching (necessarily decreasing) transitions

$$X_k^a\, \delta' \xrightarrow{\psi_1} \ldots \xrightarrow{\psi_l} X_k^a\, \delta'',$$

for some δ'', such that

$$\gamma\, X_k^{a+1} \approx \gamma\, X_k^a\, \delta''.$$

This completes the proof. $\Box$

Proof of Lemma 8. If α is not unambiguous then it is not the least one in its bisimulation class wrt. $\preceq$ as $\alpha \approx \text{squeeze}(\alpha)$ and $\text{squeeze}(\alpha) \prec \alpha$.

On the other hand, assume α is not the least process in its bisimulation class. That is, for some $i \leq n$ we have $\alpha = \gamma X_i^a \bar{\alpha}$ and there is some $\beta = \gamma X_i^b \bar{\beta} \approx \alpha$ with $b < a$. Thus, according to the definition, α is not unambiguous. $\Box$

Proof of Lemma 10. If α is unambiguous the proof is trivial therefore assume otherwise. According to Definition 8, let $\alpha = \gamma X_k^a \delta$ and let

$$\text{squeeze}(\alpha) = \gamma\, X_k^{a-1}\, \delta_{k,\gamma}\, \delta. \tag{17}$$

Consider any $\bar{\alpha} \sqsubseteq \text{squeeze}(\alpha)$ such that $\bar{\alpha} \approx \text{squeeze}(\alpha)$. First we observe that γ is necessarily a $(k-1)$-prefix of $\bar{\alpha}$ as α is $(k-1)$-unambiguous and $\alpha \approx \bar{\alpha}$. Therefore

$$\bar{\alpha} = \gamma\, X_k^{b-1}\, \bar{\delta}_{k,\gamma}\, \bar{\delta}$$

for some $b \leq a$ and $\bar{\delta}_{k,\gamma} \sqsubseteq \delta_{k,\gamma}$ and $\bar{\delta} \sqsubseteq \delta$. We observe that $\bar{\delta}_{k,\gamma}$ is necessarily non-empty, as α is $\sqsubseteq$-minimal and $\alpha \approx \bar{\alpha}$. For $\text{size}(\alpha) \leq \text{size}(\bar{\alpha})$ it is thus sufficient to demonstrate that

$$b = a \quad \text{and} \quad \bar{\delta} = \delta.$$

Towards a contradiction assume the opposite, i.e., either $b < a$, or $\bar{\bar{\delta}} \sqsubset \delta$. As $\alpha \approx \bar{\alpha}$, i.e.,

$$\gamma\, X_k^a \;\approx\; \gamma\, X_k^{b-1}\, \bar{\delta}_{k,\gamma}\, \bar{\bar{\delta}},$$

knowing that $a > b-1$ we deduce that the process γX^b may not be k-unambiguous. Thus we may apply squeeze(_) to $\gamma X^b \bar{\bar{\delta}}$ to obtain

$$\text{squeeze}(\gamma\, X_k^b\, \bar{\bar{\delta}}) = \gamma\, X_k^{b-1}\, \delta_{k,\gamma}\, \bar{\bar{\delta}}.$$

By Lemma 1 applied to

$$\text{squeeze}(\alpha) = \gamma\, X_k^{a-1}\, \delta_{k,\gamma}\, \delta \Longrightarrow_0 \gamma\, X_k^{b-1}\, \delta_{k,\gamma}\, \bar{\bar{\delta}} \Longrightarrow_0 \gamma\, X_k^{b-1}\, \bar{\delta}_{k,\gamma}\, \bar{\bar{\delta}} = \bar{\alpha}$$

we deduce $\text{squeeze}(\alpha) \approx \gamma\, X_k^{b-1}\, \delta_{k,\gamma}\, \bar{\bar{\delta}}$, i.e.,

$$\text{squeeze}(\alpha) \;\approx\; \text{squeeze}(\gamma\, X_k^b\, \bar{\bar{\delta}}).$$

Since always $\alpha \approx \text{squeeze}(\alpha)$ we obtain

$$\alpha = \gamma\, X_k^a\, \delta \;\approx\; \text{squeeze}(\alpha) \;\approx\; \text{squeeze}(\gamma\, X_k^b\, \bar{\bar{\delta}}) \;\approx\; \gamma\, X_k^b\, \bar{\bar{\delta}},$$

with either $b < a$ or $\bar{\bar{\delta}} \sqsubset \delta$, thus contradicting the $\sqsubseteq$-minimality of α. This completes the proof. □

Proof of Lemma 11. Let α be an arbitrary process. We claim that the size of $\text{nf}(\alpha)$ is bounded by:

$$\text{size}(\text{nf}(\alpha)) \leq \text{size}(\alpha) \cdot \text{size}(\delta_{k_1,\gamma_1}) \cdot \ldots \cdot \text{size}(\delta_{k_n,\gamma_n}) \tag{18}$$

for some unambiguous processes $\gamma_1 \ldots \gamma_n$. Indeed, let γ_k be the $(k-1)$-unambiguous process witnessing the squeezing step for X_k (if any). The size of the process, during all squeezing steps for X_k, increases at most $\text{size}(\delta_{k_k,\gamma_k})$ times.

However, in general, there may be infinitely many different processes $\delta_{k,\gamma}$ used in the squeezing steps for different processes α, as there may be in general infinitely many unambiguous processes γ. We will argue that for the purpose of estimating the size of $\text{nf}(\alpha)$ for all processes α, it is sufficient to take into account only a finite subset of unambiguous processes. We will rely on the following simple observation. Let $\gamma, \gamma' \in \{X_1 \ldots X_{k-1}\}^{\otimes}$, for some $k \leq n$, be both $(k-1)$-unambiguous and $\gamma \sqsubseteq \gamma'$, respectively. Let the greatest k-extensions of γ and γ' be γX_k^a and $\gamma X_k^{a'}$. The exponents necessarily satisfy $a \geq a'$. The crucial observation is that whenever $a = a'$ then every γ-squeeze, like $\delta_{k,\gamma}$, is also a γ'-squeeze. Indeed:

$$\gamma\, X_k^{a+1}\, \delta \;\approx\; \gamma\, X_k^a\, \delta_{k,\gamma}\, \delta \;\;\text{implies}\;\; \gamma'\, X_k^{a+1}\, \delta \;\approx\; \gamma'\, X_k^a\, \delta_{k,\gamma}\, \delta,$$

since $\approx$ is substitutive. In other words: one may safely assume $\delta_{k,\gamma'} = \delta_{k,\gamma}$ whenever $\gamma \sqsubseteq \gamma'$ and $a \leq a'$.

Now we easily obtain the estimation. For every $k \in \{1 \ldots n\}$, consider all pairs (γ, a), where $\gamma \in \{X_1 \ldots X_{k-1}\}^{\otimes}$ is any $(k-1)$-unambiguous process that exhibits the greatest extension γX_k^a (note that only such processes γ witness a squeezing step). Choose those among them that are minimal wrt. $\sqsubseteq$ on the first coordinate, and wrt. $\leq$ on the second one. By Dickson's Lemma there are only finitely many such minimal pairs. The set of all processes $\delta_{k,\gamma}$, for all chosen minimal pairs (γ, a), jointly for all k, has an element which is maximal wrt. size; denote this maximal size by s. The size of any process δ_{k_i,γ_i} in (18) is dominated by s and thus we obtain:

$$\mathrm{size}(\mathrm{nf}(\alpha)) \leq \mathrm{size}(\alpha) \cdot s^n \tag{19}$$

which completes the proof by putting $c = s^n$. Note that c only depends on a process definition, and does not depend on a process α. □

B Proofs missing in Section 5

Consider a sequence $\alpha \xrightarrow{\zeta_1} \ldots \xrightarrow{\zeta_m} \beta$ of decreasing transitions. By a *canonical re-shuffling* we mean any permutation of these transitions that respects the following two conditions: (1) if $X > Y$ then every transition of X occurs before every transition of Y; (2) every generating transition of a variable occurs before every non-generating transition of the same variable. Due to Lemma 3 we have:

Lemma 18. *Any sequence $\alpha \xrightarrow{\zeta_1} \ldots \xrightarrow{\zeta_m} \beta$ of decreasing transitions may be executed after the canonical re-shuffling.*

Proof of Lemma 13. Lemma 13 is formulated for $\approx$ but the major part of the proof either works for weak bisimilarity directly, or may be adapted. The only case that we can not adapt to weak bisimilarity is Case 2.1 below. Importantly, under the assumption of [14] the proof of this subcase is straightforward. Thus we claim that our whole proof covers weak bisimilarity over the restricted subclass studied in [14].

Case 1: $a > 0$**.**

Claim. $\gamma X_k^{a+1} \stackrel{\approx}{\Longrightarrow}_0 \gamma X_k^a \eta$ for some $\eta \in \{X_{k+1} \ldots X_n\}^{\otimes}$.

Proof. Consider the pair $\gamma X_k^{a+1} \approx \gamma X_k^a \delta_{k,\gamma}$ and an arbitrary non-generating decreasing transition $X_k \xrightarrow{\zeta} \omega$; thus $\delta_{k,\gamma}\, \omega \in \{X_{k+1} \ldots X_n\}^{\otimes}$. The transition gives rise to a Spoiler's move

$$\gamma X_k^a \delta_{k,\gamma} \xrightarrow{\zeta} \gamma X_k^{a-1} \delta_{k,\gamma}\, \omega,$$

matched by some sequence of transitions of the form

$$\gamma X_k^{a+1} \Longrightarrow_0 \alpha \xrightarrow{\zeta} \alpha',$$

such that $\alpha \approx \gamma X_k^a \delta_{k,\gamma}$ and $\alpha' \approx \gamma X_k^{a-1} \delta_{k,\gamma}\, \omega$. As γX_k^a is k-unambiguous, by the latter equivalence we deduce that γX_k^{a-1} is the k-prefix of α'. As α'

results from α by a single transition, γX_k^a is necessarily the k-prefix of α. Thus $\alpha = \gamma X_k^a \, \eta$ as required. □

Assume that the sequence of transitions $\gamma \, X_k^{a+1} \Longrightarrow_0 \gamma \, X_k^a \, \eta$ has been canonically re-shuffled, and distinguish the very first process whose k-prefix is γX_k^a:

$$\gamma \, X_k^{a+1} \Longrightarrow_0 \gamma \, X_k^a \, \omega \Longrightarrow_0 \gamma \, X_k^a \, \eta.$$

Due to the re-shuffling, the last transition of the first part involves necessarily X_k, say $X_k \longrightarrow_0 \phi$, and thus the k-prefix of the immediately proceeding process is $\gamma \, X_k^{a+1}$. We obtain:

$$\gamma \, X_k^{a+1} \Longrightarrow_0 \gamma \, X_k^{a+1} \, \delta \longrightarrow_0 \gamma \, X_k^a \, \phi \, \delta \Longrightarrow_0 \gamma \, X_k^a \, \eta.$$

Note that by Lemma 1 we have:

$$\gamma \, X_k^{a+1} \approx \gamma \, X_k^a \, \delta \, \phi. \tag{20}$$

Furthermore, by canonical order we deduce

$$\gamma \, X_k^a \Longrightarrow_0 \gamma \, X_k^a \, \delta$$

and consequently, by Lemma 5 we obtain

$$\gamma \, X_k^a \approx \gamma \, X_k^a \, \delta.$$

This allows us, using substitutivity and (20), to obtain a γ-squeeze of X_k of size at most d:

$$\gamma \, X_k^{a+1} \approx \gamma \, X_k^a \, \delta \, \phi \;\approx \gamma \, X_k^a \phi.$$

Knowing $\phi \in \{X_{k+1} \dots X_n\}^{\otimes}$ and $\text{size}(\phi) \leq d$ we easily deduce the required bound on the weighted size of ϕ:

$$d\text{-size}(\phi) \leq d \cdot d^{n-k-1} = d^{n-k} = d\text{-size}(X_k).$$

Case 2.1: $a = 0$ and X_k has a γ-squeeze δ such that $X_k \Longrightarrow_0 \delta$. This is the only case that we are not able to adapt to weak bisimilarity.

In the proof of this case we will only use $\longrightarrow_0$ transitions. As γ is unambiguous, the normal form of $\gamma \, X_k$ has the form:

$$\text{nf}(\gamma \, X_k) = \gamma \, [\alpha \cdot \beta]. \tag{21}$$

Consider the sequence of transitions $X_k \Longrightarrow_0 \delta$ after canonical re-shuffling. As some non-generating transition of X_k necessarily occurs, consider the last such transition, say $X_k \longrightarrow \omega$. We obtain:

$$\gamma X_k \Longrightarrow_0 \gamma X_k \beta' \longrightarrow_0 \gamma \omega \beta' \Longrightarrow_0 \gamma \delta,$$

for some β' generated by X_k. As $\gamma X_k \approx \gamma \delta$, by Lemma 1 we obtain $\gamma X_k \approx \gamma \omega \beta'$, and by (21) we obtain

$$\text{nf}(\gamma \omega \beta') = \gamma \, [\alpha \cdot \beta].$$

Recall that the normal form is obtained by squeezing (cf. Definition 9). As squeezing of β' can only yield variables $\{X_{l+1} \dots X_n\}$, we deduce that the normal form of $\gamma\omega$ differs from $\text{nf}(\gamma\omega\beta')$ only on variables $\{X_{l+1} \dots X_n\}$:

$$\text{nf}(\gamma\omega) = \gamma\,[\alpha \cdot \eta]. \tag{22}$$

Now consider the first step in the Bisimulation Game for $\gamma\, X_k \approx \gamma\,[\alpha \cdot \beta]$, starting with the Spoiler's move $\gamma\, X_k \longrightarrow_0 \gamma\,\omega$. Due to (22) we know that l-prefix$(\gamma\,\omega)$ is l-dominating $\gamma\,\alpha$ and thus Lemma 14 applies to give the following claim.

Claim. Duplicator has a response of the following form

$$\gamma\,[\alpha \cdot \beta] \Longrightarrow_0 \gamma\,[\alpha \cdot \beta'] \longrightarrow_0 \gamma\,[\alpha \cdot \beta'']$$

using no transition of γ or α, i.e., induced by the sequence of transitions

$$\beta \Longrightarrow_0 \beta' \longrightarrow_0 \beta''. \tag{23}$$

Let the last transition rule be, say, $Y \longrightarrow_0 \phi$. Thus we may write

$$\beta' = \bar{\beta}\,Y \text{ and } \beta'' = \bar{\beta}\,\phi$$

which allows us to prove the following:

Claim. $\gamma\, X_k\, \phi \approx \gamma\,\omega\, Y$.

Proof. On one side, as $\gamma\, X_k \approx \gamma\,[\alpha \cdot \bar{\beta}\, Y]$, by substitutivity $\gamma\, X_k\, \phi \approx \gamma\,[\alpha \cdot \bar{\beta}\, Y\, \phi]$. On the other hand, for similar reasons $\gamma\,\omega Y \approx \gamma\,[\alpha \cdot \bar{\beta}\,\phi\, Y]$. □

Note that ϕ is generated by X_k as $Y \in \beta'$. Thus $\gamma\, X_k \approx \gamma\, X_k\, \phi$ by Lemma 5. This together with the last claim yields a γ-squeeze of X_k of size at most d:

$$\gamma\, X_k \approx \gamma\,\omega\, Y.$$

As before, we deduce $d\text{-size}(\omega\, Y) \leq d\text{-size}(X_k)$.

Case 2.2: $a = 0$ and X_k has no γ-squeeze δ such that $X_k \Longrightarrow_0 \delta$. This subcase has been treated in Section 5.

As this was the last case, we have thus completed the proof of Lemma 13. □

C Proofs missing in Section 6

Proof of Lemma 16. First note that some $\stackrel{\approx}{\Longrightarrow}_0$-minimal process $\bar{\alpha}$ with $\alpha \stackrel{\approx}{\Longrightarrow}_0 \bar{\alpha}$ surely exists as $\preceq$ is well-founded. Every such process is necessarily $\sqsubseteq$-minimal: indeed, If $\beta \sqsubset \bar{\alpha}$ then $\beta \prec \bar{\alpha}$ and $\bar{\alpha} \Longrightarrow_0 \beta$. Then the size bound follows immediately by Corollary 1. □

Proof of Lemma 17. Note that $\alpha \preceq \beta$ holds by the very definition. We will show that

$$\beta = \alpha\,\delta \tag{24}$$

for some δ generated by α.

For the sake of contradiction assume the shortest sequence of transitions $\alpha \Longrightarrow_0 \beta$ such that $\beta \approx \alpha$ fails to satisfy (24). Consider the last transition, say

$$\alpha\,\delta \longrightarrow_0 \beta,$$

performed necessarily by a variable, say X, that appears in α but not in δ. This last transition has the following form

$$\alpha\,\delta \longrightarrow_0 \alpha'\,\delta,$$

due to a transition $\alpha \longrightarrow_0 \alpha'$. As the last transition is necessarily decreasing and non-generating, $\alpha' \prec \alpha$, and thus $\alpha'\,\delta \prec \alpha\,\delta$. Recall that $\alpha'\,\delta \approx \alpha$. By Lemma 5 we know that those variables in δ that are generated by a variable different than X may be safely removed. Hence

$$\alpha \Longrightarrow_0 \alpha'\,\delta' \approx \alpha$$

where all variables appearing in $\delta' \sqsubseteq \delta$ are generated by X, and thus smaller than X wrt. $\leq$. This implies

$$\alpha'\,\delta' \prec \alpha,$$

a contradiction. □