

Partially-commutative context-free processes^{*}

Wojciech Czerwiński¹, Sibylle Fröschle², and Sławomir Lasota¹

¹ Institute of Informatics, University of Warsaw

² University of Oldenburg

Abstract. Bisimulation equivalence is decidable in polynomial time for both sequential and commutative normed context-free processes, known as BPA and BPP, respectively. Despite apparent similarity between the two classes, different techniques were used in each case. We provide one polynomial-time algorithm that works in a superclass of both normed BPA and BPP. It is derived in the setting of *partially-commutative context-free processes*, a new process class introduced in the paper. It subsumes both BPA and BPP and seems to be of independent interest.

We investigate the bisimulation equivalence of the *context-free processes*, i.e., the process graphs defined by a context-free grammar in Greibach normal form. In process algebra, there are two essentially different ways of interpreting such grammars, depending on whether the concatenation is understood as the sequential or parallel composition of processes. These two process classes are known as BPA (Basic Process Algebra) and BPP (Basic Parallel Processes) [1].

The bisimulation equivalence is decidable both in BPA and BPP [2, 7]. Under the assumption of *normedness* the polynomial-time algorithms exist [5, 6]. These surprising results were obtained basing on the strong *unique decomposition property* enjoyed by both classes. Despite the apparent similarity of BPA and BPP, the algorithms are fundamentally different; cf. [1] (Chapt. 9, p. 573):

“These algorithms are both based on an exploitation of the decomposition properties enjoyed by normed transition systems; however, despite the apparent similarity of the two problems, different methods appear to be required.”

In [4] a decision procedure was given for normed PA, a superclass of both BPA and BPP. It is however very complicated and has doubly exponential complexity. In [8] a polynomial-time algorithm was proposed for the normed BPA vs. BPP problem. It transforms a BPP process into BPA, if possible, and then refers to a BPA algorithm.

This paper contains a polynomial-time algorithm for a superclass of normed BPA and BPP. The algorithm simultaneously applies to BPA and BPP, thus confirming the similarity of the two classes. Our contributions are as follows.

In Section 1 we introduce a new class of *partially-commutative* context-free processes, called BPC, build on the underlying concept of *(in)dependence* of

^{*} The first and the last author acknowledge a partial support by Polish government grants no. N206 008 32/0810 and N N206 356036.

elementary processes. BPA (no independence) and BPP (full independence) are special cases. Our main motivation was to introduce a common setting for both BPA and BPP; however, the BPC class seems to be of independent interest and may be applied, e.g., as an abstraction in program analysis.

Our first main result is the proof of the unique decomposition property for normed BPC with a transitive dependence relation (Thm 1 in Section 2).

Then in Sections 3–5 we work out our second main result: a polynomial-time algorithm for the bisimulation equivalence in the *feasible* fragment of normed BPC, to be explained below. It clearly subsumes both normed BPA and BPP but allows also for expressing, e.g., a parallel composition of inter-dependent BPA processes. It seems thus suitable for applications, e.g., for modeling of multi-core recursive programs. We sketch the main technical points below.

Recall the classical idea of approximating the bisimulation equivalence from above, by consecutive refinements $R \mapsto R \cap \text{exp}(R)$, where $\text{exp}(R)$ denotes the bisimulation expansion wrt. the relation R . The crucial idea underlying the BPP algorithm [6] was to ensure that the approximant R is always a congruence, and to represent it by a finite *base*; the latter requires a further additional refinement step. Our starting point was an insight of [3] that this latter step yields *the greatest norm-reducing bisimulation* contained in $R \cap \text{exp}(R)$.

The feasibility condition requires the bisimulation expansion to preserve congruences. It appears sufficient for the above scheme to work, after a suitable adaptation, in the general setting of BPC. Roughly speaking, we demonstrate in particular that the BPP algorithm works, surprisingly, for BPA just as well!

Our algorithm efficiently processes both multisets and strings over the set of elementary processes, of pessimistically exponential size. One of technical contributions of this paper is to devise a way of combining the BPP base refinement of [6] with the BPA procedure based on compressed string algorithms [10].

1 Partially commutative context-free processes

BPA and BPP are defined by a context-free grammar in Greibach normal form. The former class is built on sequential composition of processes, while the latter one on parallel composition. Thus BPA processes are elements of the free monoid generated by non-terminal symbols; and BPP processes correspond to the free commutative monoid. Our aim in this section is to define a process class corresponding to the free *partially-commutative* monoid.

A grammar in Greibach normal form consists of a set of non-terminal symbols V , which we call *variables* or *elementary processes*, and a finite set of productions, which we call *rules*, of the form

$$X \xrightarrow{a} \alpha, \tag{1}$$

where $\alpha \in V^*$, $X \in V$, and a is an alphabet letter. Additionally assume a symmetric irreflexive relation $I \subseteq V \times V$, called the *independence relation*. For convenience we will also use the *dependence relation* $D \subseteq V \times V$ defined as $D = (V \times V) \setminus I$. D is thus symmetric and reflexive.

The independence induces an equivalence in \mathbf{V}^* in a standard way: two strings over \mathbf{V} are equivalent, if one can transform one into another by a sequence of transpositions of independent variables. Formally, the equivalence $\sim_{\mathbf{I}} \subseteq \mathbf{V}^* \times \mathbf{V}^*$ is the reflexive-transitive closure of the relation containing all pairs $(wXYv, wYXv)$, for $w, v \in \mathbf{V}^*$, $(X, Y) \in \mathbf{I}$; or equivalently, $\sim_{\mathbf{I}}$ is the smallest congruence in \mathbf{V}^* relating all pairs (XY, YX) where $(X, Y) \in \mathbf{I}$. We work in the monoid $\mathbf{V}_{\mathbf{I}}^{\diamond} = \mathbf{V}^* / \sim_{\mathbf{I}}$ from now on; we call $\mathbf{V}_{\mathbf{I}}^{\diamond}$ the *free partially-commutative monoid* generated by \mathbf{I} . The subscript is usually omitted when \mathbf{I} is clear from the context. Elements of \mathbf{V}^{\diamond} will be called *partially-commutative processes*, or *processes* in short, and usually denoted by Greek letters α, β, \dots . Composition of α and β in \mathbf{V}^{\diamond} is written $\alpha\beta$. Empty process (identity in \mathbf{V}^{\diamond}) will be written as ϵ . Our development is based on the decision to interpret right-hand sides α of productions (1) as elements of \mathbf{V}^{\diamond} , instead of as words or multisets over \mathbf{V} . The induced class of processes we will call BPC (Basic Partially-Commutative algebra).

Formally, a *BPC process definition* Δ consists of a finite set \mathbf{V} of variables, a finite alphabet \mathcal{A} , an independence relation $\mathbf{I} \subseteq \mathbf{V} \times \mathbf{V}$, and a finite set of rules (1), where $\alpha \in \mathbf{V}^{\diamond}$, $X \in \mathbf{V}$, and $a \in \mathcal{A}$. The induced transition system has processes as states, and transitions derived according to the following rule:

$$X\beta \xrightarrow{a} \alpha\beta \text{ whenever } (X \xrightarrow{a} \alpha) \in \Delta, \beta \in \mathbf{V}^{\diamond}.$$

Note that when $(X, Y) \in \mathbf{I}$ it may happen that $X\beta = Y\beta'$ in \mathbf{V}^{\diamond} . In such case the rules of both X and Y contribute to the transitions of $X\beta = Y\beta'$. Particular special cases are BPA (\mathbf{I} is empty), and BPP (\mathbf{I} is the identity relation).

Example 1. Let \mathbf{I} contain the pairs (B, C) , (T, C) , (B, U) , (T, U) , and the symmetric ones. In the transition system induced by the rules:

$$\begin{array}{cccc} P \xrightarrow{a} WBCT & W \xrightarrow{a} WBC & T \xrightarrow{t} \epsilon & B \xrightarrow{b} \epsilon \\ & W \xrightarrow{s} U & U \xrightarrow{u} \epsilon & C \xrightarrow{c} \epsilon \end{array}$$

there are, among the others, the following transitions:

$$P \xrightarrow{a^3} W(BC)^3T \xrightarrow{s} U(BC)^3T \sim_{\mathbf{I}} B^3TUC^3 \xrightarrow{b^3} TUC^3 \xrightarrow{tu} C^3 \xrightarrow{c^3} \epsilon.$$

Definition 1. A *bisimulation* is any binary relation R over processes such that $R \subseteq \text{exp}(R)$, where $\text{exp}(R)$, the *bisimulation expansion wrt. R* , contains all pairs (α, β) of processes such that for all $a \in \mathcal{A}$:

1. whenever $\alpha \xrightarrow{a} \alpha'$, there is β' with $\beta \xrightarrow{a} \beta'$ and $(\alpha', \beta') \in R$,
2. the symmetric condition holds,

The *bisimulation equivalence*, written as \sim , is the union of all bisimulations.

An equivalence $\approx \subseteq \mathbf{V}^{\diamond} \times \mathbf{V}^{\diamond}$ is a congruence if it is preserved by composition: $\alpha \approx \alpha'$ and $\beta \approx \beta'$ implies $\alpha\beta \approx \alpha'\beta'$. Bisimulation equivalence is a congruence both in BPA and BPP; however it needs not be so in BPC, as the following simple example shows:

Example 2. Consider $\mathsf{D} = \{(A, B), (B, A)\}$ (plus identity pairs) and the rules below; $AB \not\sim A'B'$, despite that $A \sim A'$ and $B \sim B'$:

$$A \xrightarrow{a} \epsilon \quad A' \xrightarrow{a} \epsilon \quad B \xrightarrow{b} \epsilon \quad B' \xrightarrow{b} \epsilon.$$

2 The unique decomposition for normed processes

Assume that Δ is *normed*, i.e., for every variable $X \in \mathsf{V}$, there is a sequence of transitions $X \xrightarrow{a_1} \alpha_1 \dots \xrightarrow{a_k} \alpha_k = \epsilon$ leading to the empty process ϵ . The length of the shortest such sequence is *the norm* of X , written $|X|$. Norm extends additively to all processes. By the very definition of norm, a transition may decrease norm by at most 1. Those transitions that do decrease norm will be called *norm-reducing* (n-r-transitions, in short). The rules of Δ that induce such transitions will be called norm-reducing as well.

We will need a concept of norm-reducing bisimulation (n-r-bisimulation, in short), i.e., a bisimulation over the transition system restricted to only norm-reducing transitions. The appropriate norm-reducing expansion wrt. R will be written as $\text{n-r-exp}(R)$. Every bisimulation is a n-r-bisimulation (as a norm-reducing transition must be matched in a bisimulation by a norm-reducing one) but the converse does not hold in general.

Proposition 1. *Each n-r-bisimulation, and hence each bisimulation, is norm-preserving, i.e., whenever α and β are related then $|\alpha| = |\beta|$.*

Assume from now on that variables $\mathsf{V} = \{X_1, \dots, X_n\}$ are ordered according to non-decreasing norm: $|X_i| \leq |X_j|$ whenever $i < j$. We write $X_i < X_j$ if $i < j$. Note that $|X_1|$ is necessarily 1, and that norm of a variable is at most exponential wrt. the size of Δ , understood as the sum of lengths of all rules.

We write X^\diamond , for a subset $X \subseteq \mathsf{V}$ of variables, to mean the free partially-commutative monoid generated by X and the independence relation restricted to pairs from X . Clearly, X^\diamond inherits composition and identity from V^\diamond .

Let \equiv be an arbitrary norm-preserving congruence in V^\diamond . Intuitively, an elementary process X_i is *decomposable* if $X_i \equiv \alpha\beta$ for some $\alpha, \beta \neq \epsilon$. Note that $|\alpha|, |\beta| < |X_i|$ then. For technical convenience we prefer to apply a slightly different definition. We say that X_i is *decomposable* wrt. \equiv , if $X_i \equiv \alpha$ for some process $\alpha \in \{X_1, \dots, X_{i-1}\}^\diamond$; otherwise, X_i is called *prime* wrt. \equiv . In particular, X_1 is always prime.

Denote by P the set of primes wrt. \equiv . It is easy to show by induction on norm that for each process α there is some $\gamma \in \mathsf{P}^\diamond$ with $\alpha \equiv \gamma$; in such case γ is called a *prime decomposition* of α . Note that a prime decomposition of X_i is either X_i itself, or it belongs to $\{X_1, \dots, X_{i-1}\}^\diamond$. We say that \equiv has the *unique decomposition property* if each process has precisely one prime decomposition. While the set P of primes depends on the chosen ordering of variables (in case $X_i \equiv X_j$, $i \neq j$), the unique decomposition property does not.

In general \sim , even if it is a congruence, needs not to have the unique decomposition property, as the following example shows:

Example 3. Let $\mathbf{I} = \{(B, C), (C, B)\}$ and the rules be as follows:

$$A \xrightarrow{a} B \quad A' \xrightarrow{a} C \quad B \xrightarrow{b} \epsilon \quad C \xrightarrow{c} \epsilon.$$

Consider two equivalent processes $AC \sim A'B$. As all four variables are prime wrt. \sim , we have thus a process with two different prime decompositions wrt. \sim .

The situation is not as bad as Example 3 suggests. We can prove the unique decomposition property (Thm 1 below) assumed that \mathbf{D} is transitive (hence \mathbf{D} is an equivalence). Abstraction classes of \mathbf{D} we call *threads*. Intuitively, a process (i.e., an abstraction class of $\sim_{\mathbf{I}}$) may be seen as a collection of strings over \mathbf{V} , one for each thread. For convenience, each of the strings will be called thread as well. This concrete representation will be extensively exploited in the sequel.

For $\alpha, \gamma \in \mathbf{V}^{\diamond}$ we say that α *masks* γ if any thread nonempty in γ is also nonempty in α . A binary relation \approx is called:

- *strongly right-cancellative* if whenever $\alpha\gamma \approx \beta\gamma$ then $\alpha \approx \beta$;
- *right-cancellative* if whenever $\alpha\gamma \approx \beta\gamma$ and both α and β mask γ then $\alpha \approx \beta$.

Proposition 2. *If a congruence has the unique decomposition property then it is strongly right-cancellative.*

Proposition 3. *If \approx is strongly right-cancellative then $\text{exp}(\approx)$ and $n\text{-r-exp}(\approx)$ are right-cancellative.*

A counterexample to the unique decomposition property of \equiv , if any, is a pair (α, β) of processes from \mathbf{P}^{\diamond} with $\alpha \equiv \beta$, $\alpha \neq \beta$. If there is a counterexample, there is one of minimal norm; we call it a *minimal \equiv -counterexample*. A congruence \equiv is *weakly right-cancellative* if whenever $\alpha\gamma \equiv \beta\gamma$ and both α and β mask γ and $(\alpha\gamma, \beta\gamma)$ is a minimal \equiv -counterexample then $\alpha \equiv \beta$.

Theorem 1. *Assume \mathbf{D} to be transitive. Then each weakly right-cancellative congruence that is a $n\text{-r-bisimulation}$ has the unique decomposition property.*

It is a generalization of the unique decomposition in BPA and BPP (cf. [5] and [6], resp.), in two respects. Firstly, we consider a wider class of processes. Secondly, we treat every $n\text{-r-bisimulation}$ that is a weak right-cancellative congruence, while in the two cited papers the result was proved only for the bisimulation equivalence \sim (cf. Prop. 10 in Section 3.2, where we consider the unique decomposition property of \sim). The rest of this section is devoted to the proof.

Proof of Thm 1. Fix a weakly right-cancellative congruence \equiv that is a $n\text{-r-bisimulation}$; \equiv is thus norm-preserving. Let $\mathbf{P} \subseteq \mathbf{V}$ denote primes wrt. \equiv , ordered consistently with the ordering \leq of \mathbf{V} . For the sake of contradiction, suppose that the unique decomposition property does not hold, and consider a minimal \equiv -counterexample (α, β) .

We say that a transition of one of α, β is *matched* with a transition of the other if the transitions are equally labelled and the resulting processes are related

by \equiv . Clearly, each norm-reducing transition of one of α, β may be matched with a transition of the other. Due to the minimality of the counterexample (α, β) , *any* prime decompositions of the resulting processes, say α' and β' , are necessarily identical. For convenience assume that each right-hand side of Δ was replaced by a prime decomposition wrt. \equiv . Thus α', β' must be identical.

Let t be the number of threads and let $V = V_1 \cup \dots \cup V_t$ be the partition of V into threads. A process α restricted to the i th thread we denote by $\alpha_i \in V_i^*$. Hence $\alpha = \alpha_1 \dots \alpha_t$ and the order of composing the processes α_i is irrelevant.

A (n-r-)transition of α , or β , is always a transition of the first variable in some α_i , or β_i ; such variables we call *active*. Our considerations will strongly rely on the simple observation: a n-r-transition of an active variable X may 'produce' only variables of strictly smaller norm than X , thus smaller than X wrt. \leq .

In Claims 1–6, to follow, we gradually restrict the possible form of (α, β) .

Claim 1. For each $i \leq t$, one of α_i, β_i is a suffix of the other.

Proof. Suppose that some thread i does not satisfy the requirement, and consider the longest common suffix γ of α_i and β_i . Thus γ is masked in α and β . As \equiv is weakly right-cancellative, γ must be necessarily empty – otherwise we would obtain a smaller counterexample. Knowing that the last letters of α_i and β_i , say P_α, P_β , are different, we perform a case-analysis to obtain a contradiction. The length of a string w is written $\|w\|$.

CASE 1: $\|\alpha_i\| \geq 2, \|\beta_i\| \geq 2$. After performing any pair of matching n-r-transitions, the last letters P_α, P_β will still appear in the resulting processes α', β' , thus necessarily $\alpha' \neq \beta'$ – a contradiction to the minimality of (α, β) .

CASE 2: $\|\alpha_i\| = 1, \|\beta_i\| \geq 2$ (or a symmetric one). Thus $\alpha_i = P_\alpha$. As P_α is prime, some other thread is necessarily nonempty in α . Perform any n-r-transition from that other thread. Irrespective of a matching move in β , the last letters P_α and P_β still appear in the resulting processes – a contradiction.

CASE 3: $\|\alpha_i\| = \|\beta_i\| = 1$. Thus $\alpha_i = P_\alpha, \beta_i = P_\beta$. Similarly as before, some other thread must be nonempty both in α and β . Assume wlog. $|P_\alpha| \geq |P_\beta|$. Perform any n-r-transition in α from a thread different than i . Irrespective of a matching move in β , in the resulting processes α', β' the last letter P_α in α'_i is different from the last letter (if any) in β'_i – a contradiction. \square

Claim 2. For each $i \leq t$, either $\alpha_i = \beta_i$, or $\alpha_i = \epsilon$, or $\beta_i = \epsilon$.

Proof. By minimality of (α, β) . If α_i , say, is a proper suffix of β_i , then a n-r-transition of α_i may not be matched in β . \square

A thread i is called *identical* if $\alpha_i = \beta_i \neq \epsilon$.

Claim 3. A n-r-transition of one of α, β from an identical thread may be matched only with a transition from the same thread.

Proof. Consider an identical thread i . A n-r-transition of α_i decreases $|\alpha_i|$. By minimality of (α, β) , $|\beta_i|$ must be decreased as well. \square

Claim 4. There is no identical thread.

Proof. Assume thread i is identical. Some other thread j is not as $\alpha \neq \beta$; wlog. assume $|\alpha_j| > |\beta_j|$, using Claim 1. Consider a n-r-transition of the active variable in $\alpha_i = \beta_i$ that maximises the increase of norm on thread j . This transition, performed in α , may not be matched in β , due to Claim 3, so that the norms of α_j and β_j become equal. \square

Claim 5. One of α, β , say α , has only one nonempty thread.

Proof. Consider the greatest (wrt. \leq) active variable and assume wlog. that it appears in α . We claim α has only one nonempty thread. Indeed, if some other thread is nonempty, a n-r-transition of this thread can not be matched in β . \square

Let α_i be the only nonempty thread in α , and let P_i be the active variable in that thread, $\alpha_i = P_i\gamma_i$. The process γ_i is nonempty by primality of P_i .

Claim 6. $|P_i|$ is greater than norm of any variable appearing in γ_i .

Proof. Consider any thread $\beta_j = P_j\gamma_j$ nonempty in β . We know that $|P_i| \geq |P_j|$. As the thread i is empty in β , the norm of P_j must be sufficiently large to "produce" all of γ_i in one n-r-transition, i.e., $|P_j| > |\gamma_i|$. Thus $|P_i| > |\gamma_i|$. \square

Now we easily derive a contradiction. Knowing that P_i has the greatest norm in α , consider the processes $P_i\alpha \equiv P_i\beta$, and an arbitrary sequence of $|P_i|+1$ norm-reducing transitions from $P_i\beta$. We may assume that this sequence does not touch P_i as $|\beta| = |\alpha| > |P_i|$. Let β' be the resulting process, and let α' denote the process obtained by performing some matching transitions from $P_i\alpha = P_iP_i\gamma_i$. The variable P_i may not appear in α' while it clearly appears in β' . Thus $\alpha' \equiv \beta'$, $\alpha' \neq \beta'$ and $|\alpha'| = |\beta'|$ is smaller than $|\alpha| = |\beta|$ – a contradiction to the minimality of the counterexample (α, β) . This completes the proof of the theorem.

3 The algorithm

From now on we only consider normed BPC process definitions Δ with a transitive dependence relation D . Such Δ is called *feasible* if it satisfies the following:

Assumption 1 (Feasibility). Whenever \equiv is a congruence, then the relations $\equiv \cap \text{exp}(\equiv)$ and $\equiv \cap \text{n-r-exp}(\equiv)$ are congruences as well.

Clearly, not all normed BPC process definitions are feasible (cf., e.g., Example 2 and the smallest congruence \equiv such that $A \equiv A'$ and $B \equiv B'$).

Proposition 4. *Every normed BPA or BPP process definition is feasible.*

We prefer to separate description of the algorithm from the implementation details. In this section we provide an outline of the algorithm only. In Section 4 we explain how each step can be implemented. Without further refinement, this would give an exponential-time procedure. Finally, in Section 5 we provide the polynomial-time implementation of crucial subroutines. Altogether, Sections 3–5 contain the proof of our main result:

Theorem 2. *The bisimulation equivalence \sim is decidable in polynomial time for feasible BPC process definitions.*

The algorithm will compute a finite representation of \sim . From now on let Δ be a fixed feasible process definition with variables V and dependence D ; we also fix an ordering of variables $V = \{X_1, \dots, X_n\}$.

3.1 Bases

A *base* will be a finite representation of a congruence having the unique decomposition property. A base $B = (P, E)$ consists of a subset $P \subseteq V$ of variables, and a set E of equations $(X_i = \alpha)$ with $X_i \notin P$, $\alpha \in (P \cap \{X_1, \dots, X_{i-1}\})^\diamond$ and $|X_i| = |\alpha|$. We assume that there is precisely one equation for each $X_i \notin P$.

An equation $(X_i = \alpha) \in E$ is thought to specify a decomposition of a variable $X_i \notin P$ in P^\diamond . Put $d_B(X_i) = \alpha$ if $(X_i = \alpha) \in E$ and $d_B(X_i) = X_i$ if $X_i \in P$. We want d_B to unambiguously extend to all processes as a homomorphism from V^\diamond to P^\diamond . This is only possible when, intuitively, decompositions of independent variables are independent. Formally, we say that a base B is *I-preserving* if whenever $(X_i, X_j) \in I$, and $d_B(X_i) = \alpha$, $d_B(X_j) = \beta$, then $\alpha\beta = \beta\alpha$ in P^\diamond .

The elements of P are, a priori, arbitrarily chosen, and not to be confused with the primes wrt. a given congruence. However, an I-preserving base B naturally induces a congruence $=_B$ on V^\diamond : $\alpha =_B \beta$ iff $d_B(\alpha) = d_B(\beta)$. It is easy to verify that primes wrt. $=_B$ are precisely variables from P and that $=_B$ has the unique decomposition property. Conversely, given a congruence \equiv with the latter property, one easily obtains a base B : take primes wrt. \equiv as P , and the (unique) prime decompositions of decomposable variables as E . B is guaranteed to be I-preserving, by the uniqueness of decomposition of $XY \equiv YX$, for $(X, Y) \in I$. As these two transformations are mutually inverse, we have just shown:

Proposition 5. *A norm-preserving congruence in V^\diamond has unique decomposition property iff it equals $=_B$, for an I-preserving base B .*

This allows us, in particular, to speak of *the base of* a given congruence, if it exhibits the unique decomposition property; and to call elements of P *primes*.

3.2 Outline of the algorithm

For an equivalence \equiv over processes, let $\text{gnrb}(\equiv)$ denote the greatest n-r-bisimulation that is contained in \equiv , defined as the union of all n-r-bisimulations contained in \equiv . It admits the following fix-point characterization:

Proposition 6. $(\alpha, \beta) \in \text{gnrb}(\equiv)$ iff $\alpha \equiv \beta$ and $(\alpha, \beta) \in n\text{-r-exp}(\text{gnrb}(\equiv))$.

Proposition 7. (i) if \equiv is a congruence then $\text{gnrb}(\equiv)$ is a congruence as well. (ii) if \equiv is right-cancellative then $\text{gnrb}(\equiv)$ is weakly right-cancellative.

Proof. (i) $\text{gnrb}(\equiv)$ is the intersection of the descending chain of relations $\equiv_1 := \equiv \cap \text{n-r-exp}(\equiv)$, $\equiv_2 := \equiv_1 \cap \text{n-r-exp}(\equiv_1)$, \dots . Due to Assumption 1 each \equiv_i is a congruence, hence $\text{gnrb}(\equiv)$ is a congruence too.

(ii) Consider a minimal $\text{gnrb}(\equiv)$ -counterexample $(\alpha\gamma, \beta\gamma)$ such that γ is masked both by α and β . Hence each n-r-transition of α (β) is matched by a transition of β (α); the resulting processes have the same prime decompositions, due to minimality of $(\alpha\gamma, \beta\gamma)$, and thus are related by $\text{gnrb}(\equiv)$. This proves that $(\alpha, \beta) \in \text{n-r-exp}(\text{gnrb}(\equiv))$. Due to right-cancellativity of \equiv we have also $\alpha \equiv \beta$. Now by the *if* implication of Prop. 6 we deduce $(\alpha, \beta) \in \text{gnrb}(\equiv)$. \square

Here is the overall idea. We start with the initial congruence $=_{\mathbf{B}}$ that relates processes of equal norm, and then perform the fixpoint computation by refining $=_{\mathbf{B}}$ until it finally stabilizes. The initial approximant has the unique decomposition property. To ensure that all consecutive approximants also have the property, we apply the refinement step: $=_{\mathbf{B}} \mapsto \text{gnrb}(=_{\mathbf{B}} \cap \text{exp}(=_{\mathbf{B}}))$. By Assumption 1, $=_{\mathbf{B}} \cap \text{exp}(=_{\mathbf{B}})$ is a congruence, and by Prop. 2 and 3 it is right-cancellative. Thus Prop. 7 applies to $\text{gnrb}(=_{\mathbf{B}} \cap \text{exp}(=_{\mathbf{B}}))$ and in consequence of Thm. 1 we get:

Proposition 8. *$\text{gnrb}(=_{\mathbf{B}} \cap \text{exp}(=_{\mathbf{B}}))$ is a congruence with the unique decomposition property.*

Outline of the algorithm:

- (1) Compute the base \mathbf{B} of 'norm equality'.
- (2) If $=_{\mathbf{B}}$ is a bisimulation then halt and return \mathbf{B} .
- (3) Otherwise, compute the base of the congruence $\text{gnrb}(=_{\mathbf{B}} \cap \text{exp}(=_{\mathbf{B}}))$.
- (4) Assign this new base to \mathbf{B} and go to step (2).

This scheme is a generalization of the BPP algorithm [6]. As our setting is more general, the implementation details, to be given in the next sections, will be necessarily more complex than in [6]. However, termination and correctness may be proved without inspecting the details of implementation:

Proposition 9 (termination). *The number of iterations is smaller than n .*

Proof. In each iteration the current relation $=_{\mathbf{B}}$ gets strictly finer (if \mathbf{B} did not change in one iteration, then $=_{\mathbf{B}}$ would necessarily be a bisimulation). Therefore all prime variables stay prime, and at least one non-prime variable becomes prime. To prove this suppose the contrary. Consider the smallest X_i wrt. \leq such that its prime decomposition changes during the iteration. X_i has thus two different prime decompositions (wrt. the 'old' relation $=_{\mathbf{B}}$), a contradiction. \square

Proposition 10 (correctness). *The algorithm computes the base of \sim .*

Proof. The invariant $\sim \subseteq =_{\mathbf{B}}$ is preserved by each iteration. The opposite inclusion $=_{\mathbf{B}} \subseteq \sim$ follows when $=_{\mathbf{B}}$ is a bisimulation. \square

The unique decomposition property of \sim is thus only a corollary, as we did not have to prove it prior to the design of the algorithm! A crucial discovery is that the unique decomposition must only hold for the relations $\text{gnrb}(=_{\mathbf{B}} \cap \text{exp}(=_{\mathbf{B}}))$ and that these relations play a prominent role in the algorithm (cf. [3]).

4 Implementation

Step (1) is easy: recalling that $|X_1| = 1$, initialize \mathbf{B} by $\mathbf{P} := \{X_1\}$, $\mathbf{E} := \{X_i = X_1^{|X_i|} : i = 2 \dots n\}$. On the other hand implementations of steps (2) and (3) require some preparation. We start with a concrete characterization of I-preserving bases $\mathbf{B} = (\mathbf{P}, \mathbf{E})$. Distinguish *monic threads* as those containing precisely one prime variable. \mathbf{B} is called *pure* if for each decomposition $(X_i = \alpha) \in \mathbf{E}$, α contains only variables from the thread of X_i and from (other) monic threads.

Proposition 11. *A base \mathbf{B} is I-preserving if and only if it is pure.*

Proof. The *if* implication is immediate: if \mathbf{B} is pure and the decompositions $\alpha = d_{\mathbf{B}}(X_i)$, $\beta = d_{\mathbf{B}}(X_j)$ of independent variables X_i, X_j both contain a prime from some thread, then the thread is necessarily monic. Thus $\alpha\beta = \beta\alpha$. This includes the case when any of X_i, X_j is prime. The *only if* implication is shown as follows. Consider a decomposition $(X_i = \alpha) \in \mathbf{E}$ and any prime X_j , appearing in α , from a thread different than that of X_i . As \mathbf{B} is I-preserving, $X_j\alpha = \alpha X_j$. Hence α , restricted to the thread of X_j , must be a monomial X_j^k . As X_j was chosen arbitrary, we deduce that this thread must be a monic one. \square

Let $\mathbf{B} = (\mathbf{P}, \mathbf{E})$ be a pure base. Two variables $X_i, X_j \notin \mathbf{P}$ are *compatible* if either they are independent, or $(X_i, X_j) \in \mathbf{D}$, $(X_i = \alpha)$, $(X_j = \beta) \in \mathbf{E}$ and α and β contain primes from the same threads. That is, α contains a prime from a thread iff β contains a prime from that thread. Note that it must be the same prime only in case of a (necessarily monic) thread different from the thread of X_i and X_j . \mathbf{B} is compatible if all pairs of non-prime variables are compatible.

Proposition 12. *Let \mathbf{B} be pure. If $=_{\mathbf{B}}$ is a n - r -bisimulation then \mathbf{B} is compatible.*

Proof. Assume $(X_i, X_j) \in \mathbf{D}$, $(X_i = \alpha)$ and $(X_j = \beta) \in \mathbf{E}$. For the sake of contradiction, suppose that some thread t is nonempty in α but empty in β : $\alpha_t \neq \epsilon$, $\beta_t = \epsilon$. We know that $(X_j X_i, \beta\alpha) \in n\text{-}r\text{-exp}(=_{\mathbf{B}})$. Hence a norm-reducing transition of $(\beta\alpha)_t = \alpha_t$ is matched by a transition of X_j , so that the decompositions of the two resulting processes are equal. In the decomposition of the left process the norm of thread t is at least $|\alpha_t|$, as X_i was not involved in the transition. On the right side, the norm decreased due to the fired transition, and is thus smaller than $|\alpha_t|$ – a contradiction. \square

Step (2) may be implemented using the fact stated below. It is essentially an adaptation of the property of *Causal base* [1] to the setting of partially-commutative processes, but the proof requires more care than in previously studied settings.

Proposition 13. *Let $\mathbf{B} = (\mathbf{P}, \mathbf{E})$ be pure and compatible. Then $=_{\mathbf{B}}$ is a bisimulation if and only if $(X, \alpha) \in \text{exp}(=_{\mathbf{B}})$ for each $(X = \alpha) \in \mathbf{E}$.*

Proof. We only need to consider the *if* direction. For any pair α, β of processes such that $\alpha =_{\mathbf{B}} \beta$, we should show that $(\alpha, \beta) \in \text{exp}(=_{\mathbf{B}})$. Let $\gamma := d_{\mathbf{B}}(\alpha) = d_{\mathbf{B}}(\beta)$ be the prime decomposition of α and β . It is sufficient to prove that $(\alpha, \gamma) \in \text{exp}(=_{\mathbf{B}})$, as $\text{exp}(=_{\mathbf{B}})$ is symmetric and transitive. We will analyse the possible transitions of α and γ , knowing that all decompositions in \mathbf{E} are pure.

First consider the possible transitions of α . Let X_i be an active variable in α , i.e., $\alpha = X_i \alpha'$ for some α' , and let $\delta := d_{\mathbf{B}}(X_i)$. Then γ may be also split into $\gamma = \delta \gamma'$, where $\gamma' = d_{\mathbf{B}}(\alpha')$. Thus, any transition of α may be matched by a transition of γ , as we know, by assumption, that $(X_i, \delta) \in \text{exp}(=_{\mathbf{B}})$.

Now we consider the possible transitions of γ . Let a prime X_j be active in γ . Choose a variable X_i such that X_j appears in the decomposition $\delta = d_{\mathbf{B}}(X_i)$. Due to compatibility of \mathbf{B} we may assume that the chosen X_i is active in α , i.e., $\alpha = X_i \alpha'$, for some α' . Similarly as above we have $\gamma = \delta \gamma'$. A transition of X_j is necessarily a transition of δ , hence may be matched by a transition of X_i , by the assumption that $(X_i, \delta) \in \text{exp}(=_{\mathbf{B}})$. \square

Proposition 14. *The base \mathbf{B} is pure and compatible in each iteration.*

Proof. Initially \mathbf{B} is pure and compatible. After each iteration \mathbf{B} , being the base of $\text{gnrb}(=_{\mathbf{B}} \cap \text{exp}(=_{\mathbf{B}}))$, is pure by Prop. 11, 8 and 5, and compatible by Prop. 12. \square

Therefore in step (2), the algorithm only checks the condition of Prop. 13.

Implementation of step (3). We compute the base $\mathbf{B}' = (\mathbf{P}', \mathbf{E}')$ of the greatest n-r-bisimulation contained in $=_{\mathbf{B}} \cap \text{exp}(=_{\mathbf{B}})$. As only norm-reducing transitions are concerned, the base is obtained in a sequence of consecutive extensions, by inspecting the variables according to their ordering, as outlined below.

In the following let \equiv denote the relation $=_{\mathbf{B}} \cap \text{exp}(=_{\mathbf{B}})$. The algorithm below is an implementation of the fix-point characterization of $\text{gnrb}(\equiv)$ (cf. Prop. 6).

Implementation of step (3):

Start with the set $\mathbf{P}' = \{X_1\}$ of primes and the empty set \mathbf{E}' of decompositions. Then for $i := 2, \dots, n$ do the following:

Check if there is some $\alpha \in \mathbf{P}'^{\diamond}$ such that

(a) $(X_i, \alpha) \in \text{n-r-exp}(=_{\mathbf{B}'})$, and (b) $X_i \equiv \alpha$.

If one is found, add $(X_i = \alpha)$ to \mathbf{E}' . Otherwise, add X_i to \mathbf{P}' and thus declare X_i prime in \mathbf{B}' .

Before explaining how searching for a decomposition α of X_i is implemented, we consider the correctness issue.

Proposition 15 (correctness of step (3)). *The base B' computed in step (3) coincides with the base of $gnrb(\equiv)$.*

Now we return to the implementation of step (3). Seeking $\alpha \in (P')^\diamond$ appropriate for the decomposition of X_i is performed by an exhaustive check of all 'candidates' computed according to the procedure described below. The computation implements a necessary condition for (a) to hold: if $(X_i, \alpha) \in n\text{-r-exp}(=_{B'})$ then α is necessarily among the candidates.

Computing candidates α :

Fix an arbitrarily chosen norm-reducing rule $X_i \xrightarrow{a} \beta$ (hence $\beta \in \{X_1, \dots, X_{i-1}\}^\diamond$) and let $\beta' := d_{B'}(\beta)$ be a decomposition of β wrt. B' (hence $\beta' \in (P')^\diamond$). For any $j < i$ such that $X_j \in P'$, for any norm reducing rule $X_j \xrightarrow{a} \gamma$, do the following: let $\gamma' := d_{B'}(\gamma)$; if $\beta' = \gamma'\gamma''$, for some γ'' , then let $\alpha := X_j\gamma'' \in (P')^\diamond$ be a candidate.

We will write $\gamma \leq_{B'} \beta$ to mean that $d_{B'}(\gamma)$ is a prefix of $d_{B'}(\beta)$.

5 Polynomial-time implementation

The algorithm performs various manipulations on processes. The most important are the following 'subroutines', invoked a polynomial number of times:

- (i) Given $\alpha, \beta \in V^\diamond$ and B , check if $\alpha =_B \beta$.
- (ii) Given $\alpha, \beta \in V^\diamond$ and B , check if $\alpha \leq_B \beta$. If so, compute γ such that $\alpha\gamma =_B \beta$.

Recall that the processes involved in the algorithm are tuples of strings over prime variables, one string for each thread, of pessimistically exponential length and norm. We need thus to consider two inter-related issues:

- a succinct representation of processes in polynomial space; and
- polynomial-time implementations of all manipulations, including subroutines (i) and (ii), that preserve the succinct representation of manipulated data.

The special case of BPP is straightforward: the commutative processes are essentially multisets, may be thus succinctly represented by storing exponents in binary, and effectively manipulated in polynomial time.

In the general case of BPC, and even in BPA, we need a more elaborate approach. To get a polynomial-time implementation, we need to use a method of 'compressed' representation of strings. Moreover, all the operations performed will have to be implemented on compressed representations, without 'decompressing' to the full exponential-size strings. After preparatory Section 5.1, in Section 5.2 we explain how to implement steps (1)–(3) in polynomial time.

5.1 Compression by an acyclic morphism

Let \mathbf{A} be a finite alphabet and $\mathbf{S} = \{z_1, \dots, z_m\}$ a finite set of non-terminal symbols. An *acyclic morphism* is a mapping $h : \mathbf{S} \rightarrow (\mathbf{S} \cup \mathbf{A})^*$ such that

$$h(z_i) \in (\mathbf{A} \cup \{z_1, \dots, z_{i-1}\})^*.$$

We assume thus a numbering of symbols such that in string $h(z_i)$, only z_j with smaller index $j < i$ are allowed. Due to this acyclicity requirement, h induces a monoid morphism $h^* : \mathbf{S}^* \rightarrow \mathbf{A}^*$, as the limit of compositions $h, h^2 = h \circ h, \dots$. Formally, $h^*(z_i) = h^k(z_i)$, for the smallest k with $h^k(z_i) \in \mathbf{A}^*$. Then the extension of h^* to all strings in \mathbf{S}^* is as usual. Therefore each symbol z_i represents a nonempty string over \mathbf{A} . Its length $\|h^*(z_i)\|$ may be exponentially larger than the size of h , defined as the sum of lengths of all strings $h(z_i)$.

Action of h^* on a symbol z may be presented by a finite tree, that we call the *derivation tree* of z . The root is labeled by z . If a node is labeled by some z' , then the number of its children is equal to the length of $h(z')$. Their labels are consecutive letters from $h(z')$ and their ordering is consistent with the ordering of letters in $h(z')$. Nodes labeled by an alphabet letter are leaves. By acyclicity of h the tree is necessarily finite; the labels of its leaves store $h^*(z)$.

Lemma 1 ([9]). *Given an acyclic morphism h and two symbols $z, z' \in \mathbf{S}$, one may answer in polynomial-time (wrt. the size of h) the following questions:*

- is $h^*(z) = h^*(z')$?
- is $h^*(z)$ a prefix of $h^*(z')$?

A relevant parameter of a symbol z , wrt. an acyclic morphism h , is its *depth*, written $\text{depth}_h(z)$, and defined as the longest path in the derivation tree of z . A depth of h , $\text{depth}(h)$, is the greatest depth of a symbol.

An acyclic morphism h is *binary* if $\|h(z_i)\| \leq 2$, for all $z_i \in \mathbf{S}$. Any acyclic morphism h may be transformed to the equivalent binary one: replace each $h(z_i)$ of length greater than 2 with a balanced binary tree, using $\|h(z_i)\| - 2$ auxiliary symbols. Note that the depth d of h may increase to $d \log d$.

Lemma 2. *Given a binary acyclic morphism h , a symbol $z \in \mathbf{S}$, and $k < \|h^*(z)\|$, one may compute in polynomial-time an acyclic morphism h' extending h , such that one of new symbols of h' represents the suffix of $h^*(z)$ of length k , and $\text{size}(h') \leq \text{size}(h) + \mathcal{O}(\text{depth}_h(z))$ and $\text{depth}(h') \leq \text{depth}(h)$.*

Proof. By inspecting the path in the derivation tree of z leading to the “cutting point”, that is, to the first letter of the suffix of length k . For each symbol y appearing on this path, we will add its copy \tilde{y} to \mathbf{S} . Our intention is that \tilde{y} represents a suitable suffix of $h^*(y)$. This is achieved as follows. Assume that $h(y) = y_1 y_2$. If the path to the cutting point traverses y and y_1 , we define h for \tilde{y} as: $h(\tilde{y}) = \tilde{y}_1 y_2$. Otherwise, if the path traverses y_2 , we put $h(\tilde{y}) = \tilde{y}_2$.

The total overhead in increase of size of h is constant. Hence by repeating this procedure along the whole path from the root, labelled by z , to the cutting point, the size of h will increase by $\mathcal{O}(\text{depth}_h(z))$. Clearly, \tilde{z} represents the required suffix of $h^*(z)$. The new acyclic morphism is an extension of h , in the sense that value of h^* is preserved for all symbols that were previously in \mathbf{S} . \square

5.2 Representation of a base by an acyclic morphism

We focus on the case of BPA first, $V^\diamond = V^*$. Extension to BPC, being straightforward, is discussed at the end of this section. The complexity considerations are wrt. the size N of Δ , i.e., the sum of lengths of all rules. The subroutines (i) and (ii) may be implemented in polynomial time due to Lemma 1 and 2.

In each iteration of the algorithm, a base $B = (P, E)$ will be represented succinctly by a binary acyclic morphism h . For each $X_i \notin P$, the right-hand side of its decomposition $(X_i = \alpha_i) \in E$ will be represented by a designated symbol $x_i \in S$. Thus $d_B(X_i) = \alpha_i = h^*(x_i)$. The set P of primes will be the alphabet.

In the initial step (1), cf. Section 4, it is easy to construct such h of size $\mathcal{O}(N)$ and depth $\mathcal{O}(N \log N)$. Implementation of step (2) will be similar to checking the conditions (a) and (b) in step (3) (cf. Sect. 4), to be described now.

Given a 'compressed' representation h of B , we now show how to construct a representation h' of B' in each execution of step (3) of the algorithm; h' will be of size $\mathcal{O}(N^2 \log N)$ and depth $\mathcal{O}(N \log N)$. The alphabet A will be P' .

We construct h' by consecutive extensions, according to step (3) of the algorithm. Initially, h' is empty and $A = \{X_1\}$. For the extension step, suppose that each non-prime $X_j \notin P'$, $j < i$, has already a designated symbol x_j in h' that represents α_j , where $(X_j = \alpha_j) \in E'$. The most delicate point is the size of the representation of a 'candidate' α in step (3), as the decomposition α_i of X_i , if any, is finally found among the candidates.

Let $X_i \xrightarrow{a} \beta$ be a chosen norm-reducing rule. Replace occurrences of non-prime X_j 's in β by the corresponding x_j 's. Extend h' by $h'(y_i) = \beta$, for a fresh auxiliary symbol y_i , and transform h' to the binary form (we say that we *encode* the rule in h'). This increases the size of h' by $\mathcal{O}(N)$ and its depth by at most $\log N$. To compute a representation of a candidate α , we extend h' similarly as above, to encode a rule $X_j \xrightarrow{a} \gamma$, by $h'(y_j) := \gamma$. Then we apply Lemma 1 to check whether $h^*(y_j)$ is a prefix of $h^*(y_i)$, and if so, apply Lemma 2 to y_i , so that the newly added symbol \tilde{y}_i represents the required suffix of $h^*(x_i)$. This increases the size of h' by $\mathcal{O}(\text{depth}(h'))$ and keeps its previous depth. Finally, we compose X_j , represented by x_j , with \tilde{y}_i : $h(x_i) = x_j \tilde{y}_i$, according to step (3) of the algorithm. The cumulative increase of size and depth, after at most N repetitions of the above procedure, fits the required bounds, $\mathcal{O}(N^2 \log N)$ and $\mathcal{O}(N \log N)$, on the size and depth of h' , respectively.

To test the conditions (a) and (b) we invoke the subroutine (i) for the successors of X_i and the candidate α . This involves encoding the rules of X_i and X_j in h' , in the similar way as above. The condition (b) refers to B , so we need to merge h' with h . As the total number of candidates is polynomial, it follows that the whole algorithm runs in polynomial time.

Remark 1. The input process definition may be well given in a compressed form, i.e., by an acyclic morphism representing the right-hand sides of all rules.

Implementation for BPC. The only difference is that there may be more threads than one. Hence instead of single symbol x_i , representing decomposition

of X_i , we need to have a tuple of symbols, x_i^1, \dots, x_i^t , where t is the number of threads, to represent the content of each thread separately. The overall idea is that the algorithm should work thread-wise, i.e., process separately the strings appearing in each thread. E.g., the subroutines (i) and (ii) may be implemented analogously as for BPA, by referring to Lemma 1 and 2 for each thread.

6 Conclusions

We have provided an evidence that the bisimulation equivalence in both normed BPA and BPP can be solved by essentially the same polynomial-time algorithm.

The algorithm works correctly in a feasible fragment of normed BPC. An interesting open question remains whether the procedure may be extended to work for all of BPC with transitive D. This would probably require a quite different method, as the core ingredient of the approach of this paper is that the bisimulation equivalence is a congruence, which is not the case in general. Another open question is whether our setting can solve the normed BPA vs. BPP problem (disjoint union of normed BPA and BPP needs not be feasible, cf. Example 2).

It also remains to investigate possible applications of the BPC framework as an abstraction of programs, e.g., of multi-core computations.

Concerning the expressibility, normed BPC class and normed PA seem to be incomparable, even with respect to the trace equivalence (e.g., the process in Example 1 is not expressible in normed PA).

References

1. O. Burkart, D. Caucal, F. Moller, and B. Steffen. Verification of infinite structures. In *Handbook of Process Algebra*, pages 545–623. Elsevier, 2001.
2. S. Christensen, H. Hüttel, and C. Stirling. Bisimulation equivalence is decidable for all context-free processes. *Inf. Comput.*, 121(2):143–148, 1995.
3. S. Fröschle and S. Lasota. Normed processes, unique decomposition, and complexity of bisimulation equivalences. In *Proc. INFINITY'06*, ENTCS. To appear.
4. Y. Hirshfeld and M. Jerrum. Bisimulation equivalence is decidable for normed process algebra. In *ICALP'99*, volume 1644 of *LNCS*, pages 412–421, 1999.
5. Y. Hirshfeld, M. Jerrum, and F. Moller. A polynomial algorithm for deciding bisimilarity of normed context-free processes. *Theor. Comput. Sci.*, 158(1-2):143–159, 1996.
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*, 6:251–259, 1996.
7. P. Jančar. Bisimilarity of Basic Parallel Processes is PSPACE-complete. In *Proc. LICS'03*, pages 218–227, 2003.
8. P. Jančar, M. Kot, and Z. Sawa. Normed BPA vs. normed BPP revisited. In *Proc. CONCUR'08*, volume 5201 of *LNCS*, pages 434–446. Springer-Verlag, 2008.
9. M. Karpiński, W. Rytter, and A. Shinohara. An efficient pattern-matching algorithm for strings with short descriptions. *Nord. J. Comput.*, 4(2):172–186, 1997.
10. S. Lasota and W. Rytter. Faster algorithm for bisimulation equivalence of normed context-free processes. In *Proc. MFCS'06*, volume 4162 of *LNCS*, pages 646–657. Springer-Verlag, 2006.