Compositional model checking with divergence preserving branching bisimilarity is lively

Sander de Putter, Frédéric Lang, Anton Wijs

► To cite this version:

Sander de Putter, Frédéric Lang, Anton Wijs. Compositional model checking with divergence preserving branching bisimilarity is lively. *Science of Computer Programming*, 2020, 196, pp.102493. 10.1016/j.scico.2020.102493 . hal-02890800

HAL Id: hal-02890800

<https://hal.inria.fr/hal-02890800>

Submitted on 7 Sep 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Compositional Model Checking With Divergence Preserving Branching Bisimilarity is Lively

Sander de Putter^{a,1}, Frédéric Lang^b, Anton Wijs^{a,*}

^a*Eindhoven University of Technology, PO BOX 513, 5600 MB, Eindhoven, The Netherlands*

^b*Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP², LIG, 38000 Grenoble, France*

Abstract

Compositional model checking approaches attempt to limit state space explosion by iteratively combining behaviour of some of the components in the system and reducing the result modulo an appropriate equivalence relation. For an equivalence relation to be applicable, it should be a congruence for parallel composition where synchronisations between the components may be introduced.

An equivalence relation preserving both safety and liveness properties is divergence-preserving branching bisimilarity (DPBB). It has long been generally assumed that DPBB is a congruence for parallel composition. Recently, a congruence format has been proposed that implies that this is the case [1]. In parallel, we were the first to prove this by means of a proof assistant (Coq) for the parallel composition of Labelled Transition Systems (LTSs) with synchronisation on their common alphabet [2]. In the current article, we remove that restriction.

In addition, we show that DPBB is a congruence for LTS networks in which many LTSs are composed in parallel at once with support for multi-party synchronisation. Additionally, we discuss how to safely decompose an existing LTS network into components such that their recomposition is equivalent to the original LTS network.

Finally, to demonstrate the effectiveness of compositional model checking with intermediate DPBB reductions, we discuss the results we obtained after having conducted a number of experiments.

Keywords: Divergence-preserving branching bisimilarity, congruence, parallel composition, synchronisation, compositional model checking

1. Introduction

Model checking [3, 4] is one of the most successful approaches for the analysis and verification of the behaviour of concurrent systems. However, a major issue is the so-called

*Corresponding author

Email addresses: s.m.j.d.putter@tue.nl (Sander de Putter), frederic.lang@inra.fr (Frédéric Lang), a.j.wijs@tue.nl (Anton Wijs)

¹This work is supported by ARTEMIS Joint Undertaking project EMC2 (grant nr. 621429).

²Institute of Engineering Univ. Grenoble Alpes

state space explosion problem: the state space of a concurrent system tends to increase exponentially as the number of parallel processes increases linearly. Often, it is difficult or infeasible to verify realistic large scale concurrent systems. Over time, several methods have been proposed to tackle the state space explosion problem. Prominent approaches are the application of some form of on-the-fly reduction, such as Partial Order Reduction [5] or Symmetry Reduction [6], and compositionally verifying the system, for instance using Compositional Reasoning [7] or Partial Model Checking [8, 9].

The key operations in compositional approaches are the composition and decomposition of systems. First a system is decomposed into two or more components. Then, one or more of these components is manipulated (e.g., reduced). Finally, the components are re-composed. Comparison modulo an appropriate equivalence relation is applied to ensure that the manipulations preserve properties of interest (for instance, expressed in the modal μ -calculus [10]). These manipulations are sound if and only if the equivalence relation is a congruence for the composition expression.

Two prominent equivalences are branching bisimilarity and divergence-preserving branching bisimilarity (DPBB) [11, 12]. Branching bisimilarity preserves safety properties, while DPBB preserves both safety and liveness properties.

In [13] it is proven that DPBB is the coarsest equivalence that is a congruence for parallel composition without synchronisation between the parallel components. However, in general, parallel composition involves some synchronisation mechanism, and compositional reasoning requires equivalences that are a congruence for parallel composition in which parallel components may synchronise their behaviour. Languages to model concurrent systems, such as the process algebras CCS [14], CSP [15], ACP [16], mCRL2 [17], and LOTOS [18], include a parallel composition operator that supports synchronisation. Therefore, in the following, when we refer to parallel composition, we imply, unless stated otherwise, that synchronisation is involved.

It is known that branching bisimilarity is a congruence for parallel composition of Labelled Transition Systems (LTSs). This follows from the fact that parallel composition of LTSs can be expressed as a WB coal language [19]. The authors of [1] have proposed a congruence format for DPBB, from which it follows that DPBB is a congruence for parallel composition of LTSs. We place such a congruence proof in the context of LTS networks and compositional model checking. The current article extends earlier work [2] that was conducted in parallel to [1]. Apart from our earlier work, no results obtained with the use of a proof assistant have been reported.

A popular toolbox that offers a selection of compositional approaches is CADP [20]. CADP offers both *property-independent* approaches (e.g., compositional model generation, smart reduction, and compositional reasoning via behavioural interfaces) and *property-dependent* approaches (e.g., property-dependent reductions [21] and partial model checking [8]). The formal semantics of concurrent systems are described using *networks of LTSs* [22], or *LTS networks* for short. An LTS network consists of n LTSs representing the parallel processes. A set of synchronisation laws $\mathcal{V}$ is used to describe the possible communication, i.e., synchronisation, between the process LTSs. With this synchronisation mechanism, the usual parallel composition operators of standard process algebras, such as ACP, CCS, CSP, mCRL2, and LOTOS, can be encoded.

In this setting, this work considers parallel composition of LTS networks. Given two LTS networks $\mathcal{M}$ and $\mathcal{M}'$ of size n related via a DPBB relation B , another LTS network $\mathcal{N}$ of size m , and a parallel composition operator $\parallel_\sigma$ with a mapping σ that specifies

synchronisation between components, we show there is a DPBB relation C such that

$$\mathcal{M} \ B \ \mathcal{M}' \implies \mathcal{M} \parallel_{\sigma} \mathcal{N} \ C \ \mathcal{M}' \parallel_{\sigma} \mathcal{N}$$

This result subsumes the composition of individual LTSs via composition of LTS networks of size one. Moreover, generalisation to composition of multiple LTS networks can be obtained via a reordering of the processes within LTS networks.

Contributions. In this work, we prove that DPBB is a congruence for parallel composition of LTS networks. From this it follows that DPBB is a congruence for parallel composition of LTSs. Furthermore, we present a method to safely decompose an LTS network in components such that the composition of the components is equivalent to the original LTS network. All proofs, except for the one for Section 7, have been mechanically verified using the Coq proof assistant and are available online.³

This article extends previous work [2] with a number of additional contributions. Associativity and commutativity are desirable properties as they indicate that composition of LTS networks may be done in any order. From these properties it follows that DPBB is a congruence for LTS networks, as defined by Garavel, Lang, and Mateescu [20]. To this end we define a composition operator $\parallel$ for LTS networks that is both associative and commutative. The operator $\parallel$ is built from $\parallel_{\sigma}$ with a mapping σ implementing synchronisation on the common alphabet.

Due to the definition of LTS networks $\parallel$ is not strictly commutative. However, we show that $\parallel$ is commutative with respect to global behaviour. In short, given LTS networks $\mathcal{M}$, $\mathcal{N}$, and $\mathcal{O}$, operator $\parallel$ is associative, i.e.,

$$\mathcal{M} \parallel (\mathcal{N} \parallel \mathcal{O}) = (\mathcal{M} \parallel \mathcal{N}) \parallel \mathcal{O}$$

and commutative with respect to global behaviour, i.e.,

$$\mathcal{G}_{\mathcal{M}} \parallel \mathcal{G}_{\mathcal{N}} = \mathcal{G}_{\mathcal{N}} \parallel \mathcal{G}_{\mathcal{M}}$$

, where $\mathcal{G}_{\mathcal{M}}$ ($\mathcal{G}_{\mathcal{N}}$) is the LTS that results from combining the process LTSs in $\mathcal{M}$ ($\mathcal{N}$) using the synchronisation laws of $\mathcal{M}$ ($\mathcal{N}$). Moreover, we discuss an adaptation of the definition of LTS networks using indexed families for which $\parallel$ is truly commutative.

From associativity and commutativity of $\parallel$ it follows that DPBB is a congruence for LTS networks of which the set of synchronisation laws implements synchronisation on the common alphabet. However, it is actually unnecessary to require this form of synchronisation. As this requirement excludes many LTS networks in practice, we present a proof that does not require synchronisation on a common alphabet. Given two vectors of LTSs Π and P of size n such that for each i ($1 \leq i \leq n$) the i^{th} processes of both vectors are related by a DPBB relation B_i , we show that there is a DPBB relation C that relates the networks $(\Pi, \mathcal{V})$ and $(P, \mathcal{V})$:

$$(\forall i \in 1..n. \Pi_i \ B_i \ P_i) \implies \mathcal{G}_{(\Pi, \mathcal{V})} \ C \ \mathcal{G}_{(P, \mathcal{V})}$$

³http://www.win.tue.nl/mdse/composition/DPBB_is_a_congruence_for_synchronizing_LTSs.zip.

Finally, we discuss the effectiveness of compositionally constructing state spaces with intermediate DPBB reductions in comparison with the classical, non-compositional state space construction. The discussion is based on results we obtained after having conducted a number of experiments using the CADP toolbox. In the current work, we have extended the number of test cases originally presented in [2] showing a higher variety in the effectiveness of the approach. Crouzen and Lang [23] report on experiments comparing the run-time and memory performance of three compositional verification techniques. As opposed to these experiments, our experiments concern the comparison of compositional and classical, non-compositional state space construction.

Structure of the article. Related work is discussed in Section 2. In Section 3, we discuss the notions of LTS, LTS network, so-called LTS network admissibility, and DPBB. Next, the formal composition of LTS networks is presented in Section 4. We prove that DPBB is a congruence for the composition of LTS networks. Section 5 is on the decomposition of an LTS network. Decomposition allows the redefinition of a system as a set of components. Section 6 introduces an instance of the composition operator that is both associative and commutative. From this operator it follows that DPBB is a congruence for LTS networks if the set of synchronisation laws implement synchronisation on the common alphabet. This restriction is lifted in Section 7. In Section 8 we apply the theoretical results to a set of use cases comparing a compositional construction approach with non-compositional state space construction. In Section 9 we present the conclusions and directions for future work.

2. Related Work

Networks of LTSs are introduced in [24]. The authors mention that strong and branching bisimilarity are congruences for the operations supported by LTS networks. Among these operations is the parallel composition with synchronisation on equivalent labels. A congruence proof for branching bisimilarity has been verified in PVS by Van de Pol and a textual proof was written, but both the textual proof and the PVS proof have not been made public [25]. An axiomatisation for a rooted version of divergence-preserving branching bisimulation has been performed in a Master graduation project [26]. However, the considered language does not include parallel composition. In this article, we formally show that DPBB is also a congruence for parallel composition with synchronisations between components. As DPBB is a branching bisimulation relation with an extra case for explicit divergence, the proof we present also formally shows that branching bisimilarity is a congruence for parallel composition with synchronisations between components.

Another approach supporting compositional verification is presented in [22]. Given an LTS network and a component selected from the network, the approach automatically generates an interface LTS from the remainder of the network. This remainder is called the environment. The interface LTS represents the synchronisation possibilities that are offered by the environment. This requires the construction and reduction of the system LTS of the environment. The advantage of this method is that transitions and states that do not contribute to the system LTS can be removed. In our approach only the system LTS of the considered component must be constructed. The environment is left out of scope until the components are composed.

Many process algebras support parallel composition with synchronisation on labels. Often a proof is given showing that some bisimilarity is a congruence for these operators [17, 27, 28, 29]. To generalize the congruence proofs a series of meta-theories has been proposed for algebras with parallel composition [19, 30, 31]. In [31] the *panth* format is proposed. The authors show that strong bisimilarity is a congruence for algebras that adhere to the *panth* format. The focus of the work is on the expressiveness of the format. The author of [19] proposes *WB cool* formats for four bisimilarities: weak bisimilarity, rooted weak bisimilarity, branching bisimilarity, and rooted branching bisimilarity. It is shown that these bisimilarities are congruences for the corresponding formats. In [30] similar formats are proposed for eager bisimilarity and branching bisimilarity. Eager bisimilarity is a kind of weak bisimilarity which is sensitive to divergence. The above mentioned formats do not consider DPBB.

Recently, the authors of [1] have proposed such a format for DPBB. In parallel to that work, we have proven that DPBB is a congruence for the parallel composition of LTSs in which synchronisation is applied on their common alphabet, using the Coq proof assistant [2]. The current article removes that constraint, acknowledging the result obtained in [1]. In addition, we define how to compose and decompose LTS networks, and demonstrate the practical effectiveness of compositional model checking based on DPBB using a larger set of benchmarks than the one used previously [2].

In earlier work, we presented decomposition for LTS transformation systems of LTS networks [32]. The work aims to verify the transformation of a component that may synchronise with other components. The paper proposes to calculate so called detaching laws which are similar to our interface laws. The approach can be modelled with our method. In fact, we show that the derivation of these detaching laws does not amount to a desired decomposition, i.e., the re-composition of the decomposition is *not* equivalent to the original system (see Example 5.1 discussed in Section 5).

A projection of an LTS network given a set of indices is presented in [20]. Their projection operator is similar to the consistent decomposition of LTS networks that we propose. In fact, with a suitable operator for the reordering of LTS networks our decomposition operator is equivalent to their projection operator. The current article contributes to these results that admissibility properties of the LTS network are indeed preserved for such consistent decompositions.

3. Preliminaries

In this section, we introduce the notions of LTS, LTS network, and divergence-preserving branching bisimilarity of LTSs. The potential behaviour of processes is described by means of LTSs. The behaviour of a concurrent system is described by a *network of LTSs* [22], or *LTS network* for short. From an LTS network, a *system LTS* can be derived describing the global behaviour of the network. To compare the behaviour of these systems the notion of divergence-preserving branching bisimilarity (DPBB) is used. DPBB is often used to reduce the state space of system specifications while preserving safety and liveness properties, or to compare the observable behaviour of two systems.

The semantics of a process, or a composition of several processes, can be formally expressed by an LTS as presented in Definition 3.1.

Definition 3.1 (Labelled Transition System). An LTS $\mathcal{G}$ is a tuple $(\mathcal{S}_{\mathcal{G}}, \mathcal{A}_{\mathcal{G}}, \mathcal{T}_{\mathcal{G}}, \mathcal{I}_{\mathcal{G}})$, with

- $\mathcal{S}_{\mathcal{G}}$ a set of states (which we assume to be finite);
- $\mathcal{A}_{\mathcal{G}}$ a set of action labels;
- $\mathcal{T}_{\mathcal{G}} \subseteq \mathcal{S}_{\mathcal{G}} \times \mathcal{A}_{\mathcal{G}} \times \mathcal{S}_{\mathcal{G}}$ a transition relation;
- $\mathcal{I}_{\mathcal{G}} \subseteq \mathcal{S}_{\mathcal{G}}$ a (non-empty) set of initial states.

Action labels in $\mathcal{A}_{\mathcal{G}}$ are denoted by a, b, c , etc. Additionally, there is a special action label τ that represents internal, or hidden, system steps. A transition $(s, a, s') \in \mathcal{T}_{\mathcal{G}}$, or $s \xrightarrow{a}_{\mathcal{G}} s'$ for short, denotes that LTS $\mathcal{G}$ can move from state s to state s' by performing the a -action. The transitive reflexive closure of $\xrightarrow{a}_{\mathcal{G}}$ is denoted as $\xrightarrow{a}_{\mathcal{G}}^*$, and the transitive closure is denoted as $\xrightarrow{a}_{\mathcal{G}}^+$.

LTS Network. An LTS network, presented in Definition 3.2, describes a system consisting of a finite number of concurrent process LTSs and a set of synchronisation laws that define the possible interaction between the processes. We write $1..n$ for the set of integers ranging from 1 to n . A vector $\bar{v}$ of size n contains n elements indexed from 1 to n . For all $i \in 1..n$, $\bar{v}_i$ represents the i^{th} element of vector $\bar{v}$. The *concatenation* of two vectors $\bar{v}$ and $\bar{w}$ of size n and m , respectively, is denoted by $\bar{v} \parallel \bar{w}$. In the context of composition of LTS networks, this concatenation of vectors corresponds to the parallel composition of the behaviour of the two vectors.

Definition 3.2 (LTS network). An LTS network $\mathcal{M}$ of size n is a pair $(\Pi, \mathcal{V})$, where

- Π is a vector of n concurrent LTSs. For each $i \in 1..n$, we write $\Pi_i = (\mathcal{S}_i, \mathcal{A}_i, \mathcal{T}_i, \mathcal{I}_i)$.
- $\mathcal{V}$ is a finite set of synchronisation laws. A *synchronisation law* is a tuple $(\bar{v}, a)$, where $\bar{v}$ is a vector of size n , called the *synchronisation vector*, containing synchronising action labels, and a is an action label representing the result of successful synchronisation. We have $\forall i \in 1..n. \bar{v}_i \in \mathcal{A}_i \cup \{\bullet\}$, where $\bullet$ is a special symbol denoting that Π_i performs no action. The *set of result actions* of a set of synchronisation laws $\mathcal{V}$ is defined as $\mathcal{A}_{\mathcal{V}} = \{a \mid (\bar{v}, a) \in \mathcal{V}\}$.

The explicit behaviour of an LTS network $\mathcal{M}$ is defined by its *system LTS* $\mathcal{G}_{\mathcal{M}}$ which is obtained by combining the processes in Π according to the synchronisation laws in $\mathcal{V}$ as specified by Definition 3.3. The LTS network model subsumes most hiding, renaming, cutting, and parallel composition operators present in process algebras. For instance, hiding can be applied by replacing the a component in a law by τ .

Definition 3.3 (System LTS). Given an LTS network $\mathcal{M} = (\Pi, \mathcal{V})$, its *system LTS* is defined by $\mathcal{G}_{\mathcal{M}} = (\mathcal{S}_{\mathcal{M}}, \mathcal{A}_{\mathcal{M}}, \mathcal{T}_{\mathcal{M}}, \mathcal{I}_{\mathcal{M}})$, with

- $\mathcal{I}_{\mathcal{M}} = \{\langle s_1, \dots, s_n \rangle \mid s_i \in \mathcal{I}_i\}$;

- $\mathcal{T}_{\mathcal{M}}$ and $\mathcal{S}_{\mathcal{M}}$ are the smallest relation and set, respectively, satisfying $\mathcal{I}_{\mathcal{M}} \subseteq \mathcal{S}_{\mathcal{M}}$ and for all $\bar{s} \in \mathcal{S}_{\mathcal{M}}$, $a \in \mathcal{A}_{\mathcal{V}}$, we have $\bar{s} \xrightarrow{a}_{\mathcal{M}} \bar{s}'$ and $\bar{s}' \in \mathcal{S}_{\mathcal{M}}$ iff there exists $(\bar{v}, a) \in \mathcal{V}$ such that for all $i \in 1..n$:

$$\begin{cases} \bar{s}_i = \bar{s}'_i & \text{if } \bar{v}_i = \bullet \\ \bar{s}_i \xrightarrow{\bar{v}_i}_{\Pi_i} \bar{s}'_i & \text{otherwise} \end{cases}$$

- $\mathcal{A}_{\mathcal{M}} = \{a \mid \exists \bar{s}, \bar{s}' \in \mathcal{S}_{\mathcal{M}}. \bar{s} \xrightarrow{a}_{\mathcal{M}} \bar{s}'\}$.

In Figure 1, an example of an LTS network $\mathcal{M} = (\langle \Pi_1, \Pi_2 \rangle, \mathcal{V})$ with four synchronisation laws is shown on the left, and the corresponding system LTS $\mathcal{G}_{\mathcal{M}}$ is shown on the right. Initial states are indicated with an incoming arrow. The states of the system LTS $\mathcal{G}_{\mathcal{M}}$ are constructed by combining the states of Π_1 and Π_2 . In this example, we have $\langle 1, 3 \rangle, \langle 1, 4 \rangle, \langle 2, 3 \rangle \in \mathcal{S}_{\mathcal{M}}$, of which $\langle 1, 3 \rangle$ is the single initial state of $\mathcal{G}_{\mathcal{M}}$.

The transitions of the system LTS in Figure 1 are constructed by combining the transitions of Π_1 and Π_2 according to the set of synchronisation laws $\mathcal{V}$. Law $(\langle c, c \rangle, c)$ specifies that the process LTSs can synchronise on their c -transitions, resulting in c -transitions in the system LTS. Similarly, the process LTSs can synchronise on their d -transitions, resulting in a d -transition in $\mathcal{G}_{\mathcal{M}}$. Furthermore, law $(\langle a, \bullet \rangle, a)$ specifies that process Π_1 can perform an a -transition independently resulting in an a -transition in $\mathcal{G}_{\mathcal{M}}$. Likewise, law $(\langle \bullet, b \rangle, b)$ specifies that the b -transition can be fired independently by process Π_2 . Because Π_1 does not participate in this law, it remains in state $\langle 1 \rangle$ in $\mathcal{G}_{\mathcal{M}}$. The last law states that a - and e -transitions can synchronise, resulting in f -transitions, however, in this example the a - and e -transitions in Π_1 and Π_2 are never able to synchronise since state $\langle 2, 4 \rangle$ is unreachable.

An LTS network is called *admissible* iff the synchronisation laws of the network do not synchronise, rename, or cut τ -transitions [22] as defined in Definition 3.4. The intuition behind this is that internal, i.e., hidden, behaviour should not be restricted by any operation. Partial model checking and compositional construction rely on LTS networks being admissible [20]. Hence, in this article, we also restrict ourselves to admissible LTS networks when presenting our composition and decomposition methods.

Definition 3.4 (LTS network Admissibility). An LTS network $\mathcal{M} = (\Pi, \mathcal{V})$ of length n is called admissible iff the following properties hold:

1. $\forall (\bar{v}, a) \in \mathcal{V}, i \in 1..n. \bar{v}_i = \tau \implies \neg \exists j \neq i. \bar{v}_j \neq \bullet;$ (no synchronisation of τ 's)
2. $\forall (\bar{v}, a) \in \mathcal{V}, i \in 1..n. \bar{v}_i = \tau \implies a = \tau;$ (no renaming of τ 's)
3. $\forall i \in 1..n. \tau \in \mathcal{A}_i \implies \exists (\bar{v}, a) \in \mathcal{V}. \bar{v}_i = \tau.$ (no cutting of τ 's)

Figure 1. An LTS network $\mathcal{M} = (\Pi, \mathcal{V})$ (left) and its system LTS $\mathcal{G}_{\mathcal{M}}$ (right)

Figure 2. Condition 3 of Definition 3.5

Divergence-Preserving Branching Bisimilarity. To compare LTSs, we use DPBB, also called *branching bisimilarity with explicit divergence* [11, 12]. DPBB supports abstraction from actions and preserves both safety and liveness properties. To simplify proofs we use DPBB with the weakest divergence condition (D₄) presented in [12] as presented in Definition 3.5. This definition is equivalent to the standard definition of DPBB [12]. The smallest infinite ordinal is denoted by ω .

Definition 3.5 (Divergence-Preserving Branching bisimulation). A binary relation B between two LTSs $\mathcal{G}_1$ and $\mathcal{G}_2$ is a *divergence-preserving branching bisimulation* iff for all $s \in \mathcal{S}_{\mathcal{G}_1}$ and $t \in \mathcal{S}_{\mathcal{G}_2}$, $s B t$ implies:

1. if $s \xrightarrow{a}_{\mathcal{G}_1} s'$ then
 - (a) either $a = \tau$ with $s' B t$;
 - (b) or $t \xrightarrow{\tau^*}_{\mathcal{G}_2} \hat{t} \xrightarrow{a}_{\mathcal{G}_2} t'$ with $s B \hat{t}$ and $s' B t'$.
2. symmetric to 1.
3. if there is an infinite sequence of states $(s^k)_{k \in \omega}$ such that $s = s^0$, $s^k \xrightarrow{\tau}_{\mathcal{G}_1} s^{k+1}$ and $s^k B t$ for all $k \in \omega$, then there exists a state t' such that $t \xrightarrow{\tau^+}_{\mathcal{G}_2} t'$ and $s^k B t'$ for some $k \in \omega$.
4. symmetric to 3.

Condition 3 (and its symmetric case) is illustrated in Figure 2. For every state t that is related by B to an infinite number of states along an infinite path of τ -transitions, there exists at least one state t' reachable from t via at least one τ -transition that is related by B to one of those states. In other words, the infinite sequence of τ -transitions cannot be simulated in t by zero τ -transitions. Van Glabbeek *et al.* [12] have proven that this condition coincides with the preservation of divergence.

Two states $s \in \mathcal{S}_{\mathcal{G}_1}$ and $t \in \mathcal{S}_{\mathcal{G}_2}$ are *divergence-preserving branching bisimilar*, denoted by $s \xleftrightarrow{b}^{\Delta} t$, iff there is a DPBB relation B such that $s B t$. We say that two LTSs $\mathcal{G}_1$ and $\mathcal{G}_2$ are divergence-preserving branching bisimilar, denoted by $\mathcal{G}_1 \xleftrightarrow{b}^{\Delta} \mathcal{G}_2$, iff $\forall s_1 \in \mathcal{I}_{\mathcal{G}_1}. \exists s_2 \in \mathcal{I}_{\mathcal{G}_2}. s_1 \xleftrightarrow{b}^{\Delta} s_2$ and vice versa.

4. Composition of LTS Networks

This section introduces the compositional construction of LTS networks. Composition of process LTSs results in a system LTS that tends to grow exponentially when more processes are considered.

An LTS network can be seen as being composed of several *components*, each of which consists of a number of individual processes in parallel composition, with *intra-component* synchronisation laws describing how the processes inside a component should synchronise with each other. Furthermore, *inter-component* synchronisation laws define how the components as a whole should synchronise with each other. Compositional construction of a minimal version of the final system LTS may then be performed by first constructing the system LTSs of the different components, then minimising these, and finally combining their behaviour. Example 4.1 presents an example of a network with two components and an inter-component synchronisation law.

Example 4.1 (Component). *Consider an LTS network $\mathcal{M} = (\Pi, \mathcal{V})$ with processes $\Pi = \langle \Pi_1, \Pi_2, \Pi_3 \rangle$ and synchronisation laws $\mathcal{V} = \{(\langle a, \bullet, \bullet \rangle, a), (\langle \bullet, b, b \rangle, b), (\langle c, c, c \rangle, c)\}$. We may split up the network in two components, say $\mathcal{M}_1 = (\langle \Pi_1 \rangle, \mathcal{V}_1)$ and $\mathcal{M}_{\{2,3\}} = (\langle \Pi_2, \Pi_3 \rangle, \mathcal{V}_{\{2,3\}})$. Then, $(\langle c, c, c \rangle, c)$ is an inter-component law describing synchronisation between $\mathcal{M}_1$ and $\mathcal{M}_{\{2,3\}}$. The component $\mathcal{M}_1$ consists of process Π_1 , and the set of intra-component synchronisation laws $\mathcal{V}_1 = \{(\langle a, \bullet, \bullet \rangle, a)\}$ operating solely on Π_1 . Similarly, component $\mathcal{M}_{\{2,3\}}$ consists of Π_2 and Π_3 , and the set of intra-component synchronisation laws $\mathcal{V}_{\{2,3\}} = \{(\langle \bullet, b, b \rangle, b)\}$ operating solely on Π_2 and Π_3 .*

The challenge of compositional construction is to allow manipulation of the components while guaranteeing that the observable behaviour of the system as a whole remains equivalent modulo DPBB. Even though synchronisation laws of a component may be changed, we must somehow preserve synchronisations with the other components. Such a change of synchronisation laws occurs, for instance, when reordering the processes in a component, or renaming actions that are part of inter-component synchronisation laws.

In this article, we limit ourselves to composition of two components: a left and a right component. This simplifies notations and proofs. However, the approach can be generalised to splitting networks given two sets of indices indicating which processes are part of which component, i.e., a projection operator can be used to project distinct parts of a network into components.

In the remainder of this section, first, we formalise LTS networks composition. Then, we show that admissibility is preserved when two admissible networks are composed. Finally, we prove that DPBB is a congruence for composition of LTS networks.

Composing LTS networks. Before defining the composition of two networks, we introduce a mapping indicating how the inter-component laws should be constructed from the interfaces of the two networks. An inter-component law can then be constructed by combining the interface vectors of the components and adding a result action. This is achieved through a given *interface mapping*, presented in Definition 4.1, mapping interface actions to result actions.

Definition 4.1 (Interface Mapping). Consider two LTS networks $\mathcal{M}_\Pi = (\Pi, \mathcal{V})$ and $\mathcal{M}_P = (P, \mathcal{W})$ of size n and m , respectively. An *interface mapping* between $\mathcal{M}_\Pi$ and $\mathcal{M}_P$ is a mapping $\sigma : \mathcal{A}_\mathcal{V} \setminus \{\tau\} \times \mathcal{A}_\mathcal{W} \setminus \{\tau\} \times \mathcal{A}$ describing how the interface actions of $\mathcal{M}_\Pi$ should be combined with interface actions of $\mathcal{M}_P$, and what the action label should be resulting from successful synchronisation. The set $\mathcal{A}$ is the set of actions resulting from successful synchronisation between Π and P . The actions mapped by σ are considered the interface actions.

An interface mapping implicitly defines how inter-component synchronisation laws should be represented in the separate components. These local representatives are called the *interface synchronisation laws*. A mapping between $\mathcal{M}_\Pi = (\Pi, \mathcal{V})$ and $\mathcal{M}_\mathcal{P} = (\mathcal{P}, \mathcal{W})$ implies the following sets of interface synchronisation laws:

$$\mathcal{V}_\sigma = \{(\bar{v}, a) \in \mathcal{V} \mid (a, b, c) \in \sigma\}$$

$$\mathcal{W}_\sigma = \{(\bar{w}, b) \in \mathcal{W} \mid (a, b, c) \in \sigma\}$$

An interface synchronisation law makes a component's potential to synchronise with other components explicit. An interface synchronisation law has a synchronisation vector, called the *interface vector*, that may be part of inter-component laws. The result action of an interface synchronisation law is called an *interface action*. These notions are clarified further in Example 4.2.

Example 4.2 (Interface Vector and Interface Law). *Let $\mathcal{M} = (\langle \Pi_1, \Pi_2, \Pi_3 \rangle, \mathcal{V})$ be a network with inter-component synchronisation law $(\langle a, a, b \rangle, c) \in \mathcal{V}$ and a component $M_{\{1,2\}} = (\langle \Pi_1, \Pi_2 \rangle, \mathcal{V}_{\{1,2\}})$. Then, $\langle a, a \rangle$ is an interface vector of $M_{\{1,2\}}$, and given a corresponding interface action α , the interface law is $(\langle a, a \rangle, \alpha)$.*

Together the interface laws and interface mapping describe the possible synchronisations between two components, i.e., the interface laws and interface mapping describe inter-component synchronisation laws. Given two sets of laws $\mathcal{V}$ and $\mathcal{W}$ and an interface mapping σ , the inter-component synchronisation laws are defined as follows:

$$\sigma(\mathcal{V}, \mathcal{W}) = \{(\bar{v} \parallel \bar{w}, a) \mid (\bar{v}, \alpha) \in \mathcal{V} \wedge (\bar{w}, \beta) \in \mathcal{W} \wedge (\alpha, \beta, a) \in \sigma\}$$

The mapping partitions both $\mathcal{V}$ and $\mathcal{W}$ into two sets of synchronisation laws: the interface and non-interface synchronisation laws.

The application of the interface mapping, i.e., formal composition of two LTS networks, is presented in Definition 4.2. We show that a component may be exchanged with a divergence-preserving branching bisimilar component iff the interface actions are not hidden. In other words, the interfacing with the remainder of the networks is respected when the interface actions remain observable.

Definition 4.2 (Composition of LTS networks). Consider LTS networks $\mathcal{M}_\Pi = (\Pi, \mathcal{V})$ of size n and $\mathcal{M}_\mathcal{P} = (\mathcal{P}, \mathcal{W})$ of size m . Let $\sigma : \mathcal{A}_\mathcal{V} \setminus \{\tau\} \times \mathcal{A}_\mathcal{W} \setminus \{\tau\} \times \mathcal{A}$ be an interface mapping describing the synchronisations between $\mathcal{M}_\Pi$ and $\mathcal{M}_\mathcal{P}$. The *composition* of $\mathcal{M}_\Pi$ and $\mathcal{M}_\mathcal{P}$, denoted by $\mathcal{M}_\Pi \parallel_\sigma \mathcal{M}_\mathcal{P}$, is defined as the LTS network $(\Pi \parallel \mathcal{P}, \mathcal{V} \parallel \mathcal{W})$, where $\mathcal{V} \parallel \mathcal{W} = (\mathcal{V} \setminus \mathcal{V}_\sigma)^\bullet \cup \bullet(\mathcal{W} \setminus \mathcal{W}_\sigma) \cup \sigma(\mathcal{V}, \mathcal{W})$ with $(\mathcal{V} \setminus \mathcal{V}_\sigma)^\bullet = \{(\bar{v} \parallel \bullet^m, a) \mid (\bar{v}, a) \in \mathcal{V} \setminus \mathcal{V}_\sigma\}$ and $\bullet(\mathcal{W} \setminus \mathcal{W}_\sigma) = \{(\bullet^n \parallel \bar{w}, a) \mid (\bar{w}, a) \in \mathcal{W} \setminus \mathcal{W}_\sigma\}$ the sets of synchronisation laws $\mathcal{V} \setminus \mathcal{V}_\sigma$ padded with m $\bullet$'s and $\mathcal{W} \setminus \mathcal{W}_\sigma$ padded with n $\bullet$'s, respectively.

As presented in Proposition 4.1, LTS networks that are composed (according to Definition 4.2) from two admissible networks are admissible as well.

Proposition 4.1. *Let $\mathcal{M}_\Pi = (\Pi, \mathcal{V})$ and $\mathcal{M}_\mathcal{P} = (\mathcal{P}, \mathcal{W})$ be admissible LTS networks of length n and m , respectively. Furthermore, let $\sigma : \mathcal{A}_\mathcal{V} \setminus \{\tau\} \times \mathcal{A}_\mathcal{W} \setminus \{\tau\} \times \mathcal{A}$ be an interface mapping. Then, the network $\mathcal{M} = \mathcal{M}_\Pi \parallel_\sigma \mathcal{M}_\mathcal{P}$, composed according to Definition 4.2, is also admissible.*

Proof. We show that $\mathcal{M}$ satisfies Definition 3.4:

- *No synchronisation and renaming of τ 's.* Let $(\bar{v}, a) \in (\mathcal{V} \setminus \mathcal{V}_\sigma)^\bullet \cup (\mathcal{W} \setminus \mathcal{W}_\sigma) \cup \sigma(\mathcal{V}, \mathcal{W})$ be a synchronisation law with $\bar{v}_i = \tau$ for some $i \in 1..(n+m)$. We distinguish two cases:
 - * $(\bar{v}, a) \in (\mathcal{V} \setminus \mathcal{V}_\sigma)^\bullet \cup (\mathcal{W} \setminus \mathcal{W}_\sigma)$. By construction of $(\mathcal{V} \setminus \mathcal{V}_\sigma)^\bullet$ and $(\mathcal{W} \setminus \mathcal{W}_\sigma)$, and admissibility of $\mathcal{M}_\Pi$ and $\mathcal{M}_P$, we have $\forall j \in 1..n. \bar{v}_j \neq \bullet \implies i = j$, $\forall j \in (n+1)..(n+m). \bar{v}_j \neq \bullet \implies i = j$ and $a = \tau$. Hence, it holds that $\forall j \in 1..(n+m). \bar{v}_j \neq \bullet \implies i = j$ (no synchronisation of τ 's) and $a = \tau$ (no renaming of τ 's).
 - * $(\bar{v}, a) \in \sigma(\mathcal{V}, \mathcal{W})$. By definition of $\sigma(\mathcal{V}, \mathcal{W})$, there are interface laws $(\bar{v}', \alpha') \in \mathcal{V}$ and $(\bar{v}'', \alpha'') \in \mathcal{W}$ such that $(\alpha', \alpha'', a) \in \sigma$. Hence, either $1 \leq i \leq n$ with $\bar{v}'_i = \tau$ or $n < i \leq n+m$ with $\bar{v}''_{i-n} = \tau$. Since $\mathcal{M}_\Pi$ and $\mathcal{M}_P$ are admissible, we must have $\alpha' = \tau$ or $\alpha'' = \tau$, respectively. However, the interface mapping does not allow τ as interface actions, therefore, the proof follows by contradiction.

It follows that $\mathcal{M}$ does not allow synchronisation and renaming of τ 's.

- *No cutting of τ 's.* Let $(\Pi \parallel P)_i$ be a process with $\tau \in \mathcal{A}_{(\Pi \parallel P)_i}$ for some $i \in 1..(n+m)$. We distinguish the two cases $1 \leq i \leq n$ and $n < i \leq n+m$. It follows that $\tau \in \mathcal{A}_{\Pi_i}$ for the former case and $\tau \in \mathcal{A}_{P_{i-n}}$ for the latter case. Since both $\mathcal{M}_\Pi$ and $\mathcal{M}_P$ are admissible and no actions are removed in $(\mathcal{V} \setminus \mathcal{V}_\sigma)^\bullet$ and $(\mathcal{W} \setminus \mathcal{W}_\sigma)$, in both cases there exists a $(\bar{v}, a) \in (\mathcal{V} \setminus \mathcal{V}_\sigma)^\bullet \cup (\mathcal{W} \setminus \mathcal{W}_\sigma) \cup \sigma(\mathcal{V}, \mathcal{W})$ such that $\bar{v}_i = \tau$. Hence, the composite network $\mathcal{M}$ does not allow cutting of τ 's.

Since the three admissibility properties hold, the composed network $\mathcal{M}$ satisfies Definition 3.4. $\square$

DPBB is a congruence for LTS network composition. Proposition 4.2 shows that DPBB is a congruence for the composition of LTS networks according to Definition 4.2. It is worth noting that an interface mapping does not map τ 's, i.e., synchronisation of τ -actions is not allowed. In particular, this means that interface actions must *not* be hidden when applying verification techniques on a component.

Note that Proposition 4.2 subsumes the composition of single LTSs, via composition of LTS networks of size one with trivial sets of intra-component synchronisation laws.

Proposition 4.2. *Consider LTS networks $\mathcal{M}_\Pi = (\Pi, \mathcal{V})$, $\mathcal{M}_{\Pi'} = (\Pi', \mathcal{V}')$ of size n , and $\mathcal{M}_P = (P, \mathcal{W})$ of size m . Let σ be an interface mapping describing the coupling between the interface actions in $\mathcal{A}_\mathcal{V} \cap \mathcal{A}_{\mathcal{V}'}$ and $\mathcal{A}_\mathcal{W}$. DPBB is a congruence for composition of LTS networks, i.e., it holds that*

$$\mathcal{M}_\Pi \xleftrightarrow[b]{\Delta} \mathcal{M}_{\Pi'} \implies \mathcal{M}_\Pi \parallel_\sigma \mathcal{M}_P \xleftrightarrow[b]{\Delta} \mathcal{M}_{\Pi'} \parallel_\sigma \mathcal{M}_P$$

Proof. Intuitively, we have $\mathcal{M}_\Pi \parallel_\sigma \mathcal{M}_P \xleftrightarrow[b]{\Delta} \mathcal{M}_{\Pi'} \parallel_\sigma \mathcal{M}_P$ because $\mathcal{M}_\Pi \xleftrightarrow[b]{\Delta} \mathcal{M}_{\Pi'}$ and the interface with $\mathcal{M}_P$ is respected. Since $\mathcal{M}_\Pi \xleftrightarrow[b]{\Delta} \mathcal{M}_{\Pi'}$, whenever a transition labelled with an interface action α in $\mathcal{M}_\Pi$ is able to perform a transition together with $\mathcal{M}_P$, then

$\mathcal{M}_{\Pi'}$ is able to simulate the interface α -transition and synchronise with $\mathcal{M}_P$. It follows that the branching structure and divergence is preserved. For the sake of brevity we define the following shorthand notations: $\mathcal{M} = \mathcal{M}_{\Pi} \parallel_{\sigma} \mathcal{M}_P$ and $\mathcal{M}' = \mathcal{M}_{\Pi'} \parallel_{\sigma} \mathcal{M}_P$. We show $\mathcal{M}_{\Pi} \xleftrightarrow[b]{\Delta} \mathcal{M}_{\Pi'} \implies \mathcal{M} \xleftrightarrow[b]{\Delta} \mathcal{M}'$.

Let B be a DPBB relation between $\mathcal{M}_{\Pi}$ and $\mathcal{M}_{\Pi'}$, i.e., $\mathcal{M}_{\Pi} B \mathcal{M}_{\Pi'}$. By definition, we have $\mathcal{M} \xleftrightarrow[b]{\Delta} \mathcal{M}'$ iff there exists a DPBB relation C with $\mathcal{I}_{\mathcal{M}} C \mathcal{I}_{\mathcal{M}'}$. We define C as follows:

$$C = \{(\bar{s} \parallel \bar{r}, \bar{t} \parallel \bar{r}) \mid \bar{s} B \bar{t} \wedge \bar{r} \in \mathcal{S}_{\mathcal{M}_P}\}$$

The component that is subject to change is related via the relation B that relates the states in Π and Π' . The unchanged component of the network is related via the shared state $\bar{r}$, i.e., it relates the states of $\mathcal{M}_P$ to themselves.

To prove the proposition we have to show that C is a DPBB relation. This requires proving that C relates the initial states of $\mathcal{M}$ and $\mathcal{M}'$ and that C satisfies Definition 3.5.

Initial. C relates the initial states of $\mathcal{M}$ and $\mathcal{M}'$, i.e., $\mathcal{I}_{\mathcal{M}} C \mathcal{I}_{\mathcal{M}'}$. We show that $\forall \bar{s} \in \mathcal{I}_{\mathcal{M}}. \exists \bar{t} \in \mathcal{I}_{\mathcal{M}'}. \bar{s} C \bar{t}$, the other case is symmetrical. Take an initial state $\bar{s} \parallel \bar{r} \in \mathcal{I}_{\mathcal{M}}$. Since $\mathcal{I}_{\mathcal{M}_{\Pi}} B \mathcal{I}_{\mathcal{M}_{\Pi'}}$ and $\bar{s} \in \mathcal{I}_{\mathcal{M}_{\Pi}}$, there exists a $\bar{t} \in \mathcal{I}_{\mathcal{M}_{\Pi'}}$ such that $\bar{s} B \bar{t}$. Therefore, we have $\bar{s} \parallel \bar{r} C \bar{t} \parallel \bar{r}$. Since $\bar{s} \parallel \bar{r}$ is an arbitrary state in $\mathcal{I}_{\mathcal{M}}$ the proof holds for all states in $\mathcal{I}_{\mathcal{M}}$. Furthermore, since the other case is symmetrical it follows that $\mathcal{I}_{\mathcal{M}} C \mathcal{I}_{\mathcal{M}'}$.

Def. 3.5, case 1. If $\bar{s} C \bar{t}$ and $\bar{s} \xrightarrow{a}_{\mathcal{M}} \bar{s}'$ then either $a = \tau \wedge \bar{s}' C \bar{t}$, or $\bar{t} \xrightarrow{\tau}_{\mathcal{M}'} \hat{\bar{t}} \xrightarrow{a}_{\mathcal{M}'} \bar{t}' \wedge \bar{s} C \hat{\bar{t}} \wedge \bar{s}' C \bar{t}'$. To better distinguish between the two parts of the networks, we unfold C and reformulate the proof obligation as follows: If $\bar{s} B \bar{t}$ and $\bar{s} \parallel \bar{r} \xrightarrow{a}_{\mathcal{M}} \bar{s}' \parallel \bar{r}'$ then either $a = \tau \wedge \bar{s}' B \bar{t} \wedge \bar{r} = \bar{r}'$, or $\bar{t} \parallel \bar{r} \xrightarrow{\tau}_{\mathcal{M}'} \hat{\bar{t}} \parallel \bar{r} \xrightarrow{a}_{\mathcal{M}'} \bar{t}' \parallel \bar{r}' \wedge \bar{s} B \hat{\bar{t}} \wedge \bar{s}' B \bar{t}'$. Consider synchronisation law $(\bar{v} \parallel \bar{w}, a) \in (\mathcal{V} \setminus \mathcal{V}_{\sigma})^{\bullet} \cup \bullet(\mathcal{W} \setminus \mathcal{W}_{\sigma}) \cup \sigma(\mathcal{V}, \mathcal{W})$ enabling the transition $\bar{s} \parallel \bar{r} \xrightarrow{a}_{\mathcal{M}} \bar{s}' \parallel \bar{r}'$. We distinguish three cases:

1. $(\bar{v} \parallel \bar{u}, a) \in (\mathcal{V} \setminus \mathcal{V}_{\sigma})^{\bullet}$. It follows that $\bar{w} = \bullet^m$, and thus, subsystem $\mathcal{M}_P$ does not participate. Hence, we have $\bar{r} = \bar{r}'$ and $(\bar{v}, a) \in \mathcal{V}$ enables a transition $\bar{s} \xrightarrow{a}_{\mathcal{M}_{\Pi}} \bar{s}'$. Since $\bar{s} B \bar{t}$, by Definition 3.5, we have:
 - * $a = \tau$ with $\bar{s}' B \bar{t}$. Because $\bar{s}' B \bar{t}$ and $\bar{r} = \bar{r}'$, the proof trivially follows.
 - * $\bar{t} \xrightarrow{\tau}_{\mathcal{M}_{\Pi'}} \hat{\bar{t}} \xrightarrow{a}_{\mathcal{M}_{\Pi'}} \bar{t}'$ with $\bar{s} B \hat{\bar{t}}$ and $\bar{s}' B \bar{t}'$. These transitions are enabled by laws in $\mathcal{V}' \setminus \mathcal{V}'_{\sigma}$. The set of derived laws are of the form $(\bar{v}' \parallel \bullet^m, \tau) \in (\mathcal{V}' \setminus \mathcal{V}'_{\sigma})^{\bullet}$ enabling a τ -path from $\bar{t} \parallel \bar{r}$ to $\hat{\bar{t}} \parallel \bar{r}$, and there is a law $(\bar{v}' \parallel \bullet^m, a) \in (\mathcal{V}' \setminus \mathcal{V}'_{\sigma})^{\bullet}$ enabling $\hat{\bar{t}} \parallel \bar{r} \xrightarrow{a}_{\mathcal{M}'} \bar{t}' \parallel \bar{r}$. Take $\bar{r}' := \bar{r}$ and the proof obligation is satisfied.
2. $(\bar{v} \parallel \bar{w}, a) \in \bullet(\mathcal{W} \setminus \mathcal{W}_{\sigma})$. It follows that $\bar{v} = \bullet^n$, and thus, subsystems $\mathcal{M}_{\Pi}$ and $\mathcal{M}_{\Pi'}$ do not participate; we have $\bar{s} = \bar{s}'$ and $\bar{r} \xrightarrow{a}_{\mathcal{M}_P} \bar{r}'$. We take $\bar{t}' := \bar{t}$. Hence, we can conclude $\bar{t} \parallel \bar{r} \xrightarrow{\tau}_{\mathcal{M}'} \hat{\bar{t}} \parallel \bar{r} \xrightarrow{a}_{\mathcal{M}} \bar{t}' \parallel \bar{r}'$, $\bar{s} B \bar{t}$, and $\bar{s}' B \bar{t}'$.
3. $(\bar{v} \parallel \bar{w}, a) \in \sigma(\mathcal{V}, \mathcal{W})$. Both parts of the network participate in the transition $\bar{s} \parallel \bar{r} \xrightarrow{a}_{\mathcal{M}} \bar{s}' \parallel \bar{r}'$. By definition of $\sigma(\mathcal{V}, \mathcal{W})$, there are $(\bar{v}, \alpha) \in \mathcal{V}$, $(\bar{w}, \beta) \in \mathcal{W}$ and $(\alpha, \beta, a) \in \sigma$ such that $(\bar{v}, \alpha)$ enables a transition $\bar{s} \xrightarrow{\alpha}_{\mathcal{M}_{\Pi}} \bar{s}'$ and $(\bar{u}, \beta)$ enables a transition $\bar{r} \xrightarrow{\beta} \bar{r}'$. Since $\bar{s} B \bar{t}$, by Definition 3.5, we have:

- * $\alpha = \tau$ with $\bar{s}' B \bar{t}$. Since $\alpha \in \mathcal{A}_{\mathcal{V}} \setminus \{\tau\}$ we have a contradiction.
- * $\bar{t} \xrightarrow{\tau}_{\mathcal{M}'_{\Pi'}} \hat{\bar{t}} \xrightarrow{\alpha}_{\mathcal{M}'_{\Pi'}} \bar{t}'$ with $\bar{s} B \hat{\bar{t}}$ and $\bar{s}' B \bar{t}'$. Since τ actions are not mapped by the interface mapping we have a set of synchronisation laws of the form $(\bar{v}' \parallel \bullet^m, \tau) \in (\mathcal{V}' \setminus \mathcal{V}'_{\sigma})^{\bullet}$ enabling a τ -path $\bar{t} \parallel \bar{r} \xrightarrow{\tau}_{\mathcal{M}', \hat{\bar{t}}} \bar{r}$. Let $(\bar{v}', \alpha) \in \mathcal{V}'$ be the synchronisation law enabling the α -transition. Since $(\alpha, \beta, a) \in \sigma$, α is an interface action and does not occur in $\mathcal{V}' \setminus \mathcal{V}'_{\sigma}$. It follows that $(\bar{v}', \alpha) \in \mathcal{V}'_{\sigma}$, and consequently $(\bar{v}' \parallel \bar{w}, a) \in \sigma(\mathcal{V}', \mathcal{W})$. Law $(\bar{v}' \parallel \bar{w}, a)$ enables the transition $\hat{\bar{t}} \parallel \bar{r} \xrightarrow{a}_{\mathcal{M}'} \bar{t}' \parallel \bar{r}'$, and the proof follows.

Def. 3.5, case 2. If $\bar{s} C \bar{t}$ and $\bar{t} \xrightarrow{a}_{\mathcal{M}'} \bar{t}'$ then either $a = \tau \wedge \bar{s}' C \bar{t}$, or $\bar{s} \xrightarrow{\tau}_{\mathcal{M}} \hat{\bar{s}} \xrightarrow{a}_{\mathcal{M}} \bar{s}' \wedge \bar{s} C \hat{\bar{t}} \wedge \bar{s}' C \bar{t}'$. This case is symmetric to the previous case.

Def. 3.5, case 3. If $\bar{s} C \bar{t}$ and there is an infinite sequence of states $(\bar{s}^k)_{k \in \omega}$ such that $\bar{s} = \bar{s}^0$, $\bar{s}^k \xrightarrow{\tau}_{\mathcal{M}} \bar{s}^{k+1}$ and $\bar{s}^k C \bar{t}$ for all $k \in \omega$, then there exists a state $\bar{t}'$ such that $\bar{t} \xrightarrow{\tau}_{\mathcal{M}'} \bar{t}'$ and $\bar{s}^k C \bar{t}'$ for some $k \in \omega$. Again we reformulate the proof obligation to better distinguish between the two components: if $\bar{s} \parallel \bar{r} C \bar{t} \parallel \bar{r}$ and there is an infinite sequence of states $(\bar{s}^k \parallel \bar{r}^k)_{k \in \omega}$ such that $\bar{s} \parallel \bar{r} = \bar{s}^0 \parallel \bar{r}^0$, $\bar{s}^k \parallel \bar{r}^k \xrightarrow{\tau}_{\mathcal{M}} \bar{s}^{k+1} \parallel \bar{r}^{k+1}$ and $\bar{s}^k B \bar{t}$ for all $k \in \omega$, then there exists states $\bar{t}'$ and $\bar{r}'$ such that $\bar{t} \parallel \bar{r} \xrightarrow{\tau}_{\mathcal{M}'} \bar{t}' \parallel \bar{r}'$ and $\bar{s}^k B \bar{t}'$ for some $k \in \omega$.

We distinguish two cases:

1. All steps in the τ -sequence are enabled in $\mathcal{M}_{\Pi}$, i.e., $\forall k \in \omega. \bar{s}^k \xrightarrow{\tau}_{\mathcal{M}_{\Pi}} \bar{s}^{k+1}$. Since $\bar{s} B \bar{t}$, by condition 3 of Definition 3.5, it follows that there is a state $\bar{t}'$ with $\bar{t} \xrightarrow{\tau}_{\mathcal{M}'} \bar{t}'$ and $\bar{s}^k B \bar{t}'$ for some $k \in \omega$. Since τ is not an interface action, the synchronization laws enabling $\bar{t} \xrightarrow{\tau}_{\mathcal{M}'} \bar{t}'$ are also present in $\mathcal{M}'$. Hence, we have $\bar{t} \parallel \bar{r} \xrightarrow{\tau}_{\mathcal{M}'} \bar{t}' \parallel \bar{r}$ and $\bar{s}^k B \bar{t}'$ for $k \in \omega$.
2. There is a $k \in \omega$ with $\neg \bar{s}^k \xrightarrow{\tau}_{\mathcal{M}_{\Pi}} \bar{s}^{k+1}$. We do have $\bar{s}^k \parallel \bar{r}^k \xrightarrow{\tau}_{\mathcal{M}} \bar{s}^{k+1} \parallel \bar{r}^{k+1}$ with $\bar{s}^k B \bar{t}$ (see antecedent at the start of the ‘divergence’ case). Since the τ -transition is not enabled in $\mathcal{M}_{\Pi}$ the transition must be enabled by a synchronisation law $(\bar{v} \parallel \bar{w}, \tau) \in \bullet(\mathcal{W} \setminus \mathcal{W}_{\sigma}) \cup \sigma(\mathcal{V}, \mathcal{W})$. We distinguish two cases:
 - * $(\bar{v} \parallel \bar{w}, \tau) \in \bullet(\mathcal{W} \setminus \mathcal{W}_{\sigma})$. The transition $\bar{s}^k \parallel \bar{r}^k \xrightarrow{\tau}_{\mathcal{M}} \bar{s}^{k+1} \parallel \bar{r}^{k+1}$ is enabled by $(\bar{v} \parallel \bar{w}, \tau) \in \bullet(\mathcal{W} \setminus \mathcal{W}_{\sigma})$. Therefore, there is a transition $\bar{r}^k \xrightarrow{\tau}_{\mathcal{M}_{\Pi}} \bar{r}^{k+1}$ enabled by $(\bar{w}, \tau) \in \mathcal{W} \setminus \mathcal{W}_{\sigma}$. Since this transition is part of an infinite τ -sequence, there is a path $\bar{s} \parallel \bar{r} \xrightarrow{\tau}_{\mathcal{M}} \bar{s}^k \parallel \bar{r}^k$. Furthermore, condition 1 of Definition 3.5 holds for C , hence, there is a state $\bar{t}' \in \mathcal{S}_{\mathcal{M}_{\Pi'}}$ and a transition $\bar{t} \parallel \bar{r} \xrightarrow{\tau}_{\mathcal{M}_{\Pi'}} \bar{t}' \parallel \bar{r}^k$ with $\bar{s}^k \parallel \bar{r}^k C \bar{t}' \parallel \bar{r}^k$. Therefore, we have $\bar{t} \parallel \bar{r} \xrightarrow{\tau}_{\mathcal{M}'} \bar{t}' \parallel \bar{r}^{k+1}$. Finally, since $\bar{s}^k \parallel \bar{r}^k C \bar{t}' \parallel \bar{r}^k$, it follows that $\bar{s}^k B \bar{t}'$.
 - * $(\bar{v} \parallel \bar{w}, \tau) \in \sigma(\mathcal{V}, \mathcal{W})$. By definition of $\sigma(\mathcal{V}, \mathcal{W})$, there are two laws $(\bar{v}, \alpha) \in \mathcal{V}$ and $(\bar{u}, \beta) \in \mathcal{W}$ with $(\alpha, \beta, \tau) \in \sigma$. The laws enable transitions $\bar{s}^k \xrightarrow{\alpha}_{\mathcal{M}_{\Pi}} \bar{s}^{k+1}$ and $\bar{r}^k \xrightarrow{\beta}_{\mathcal{M}_{\Pi}} \bar{r}^{k+1}$ respectively. Since $\bar{s}^k B \bar{t}$ and $\alpha \neq \tau$, by Definition 3.5, there are states $\hat{\bar{t}}, \bar{t}' \in \mathcal{S}_{\mathcal{M}_{\Pi'}}$ such that there is a sequence $\bar{t} \xrightarrow{\tau}_{\mathcal{M}_{\Pi'}} \hat{\bar{t}} \xrightarrow{\alpha}_{\mathcal{M}_{\Pi'}} \bar{t}'$ with $\bar{s} B \hat{\bar{t}}$ and $\bar{s}^{k+1} B \bar{t}'$. Let $(\bar{v}', \alpha) \in \mathcal{V}'$ be

the law enabling the α -transition. Since $(\alpha, \beta, \tau) \in \sigma$, and consequently $(\bar{v}' \parallel \bar{w}, \tau) \in \sigma(\mathcal{X}', \mathcal{Y})$. Furthermore, the τ -path from $\bar{t}$ to $\hat{t}$ is enabled by laws of the form $(\bar{v}'', \tau) \in \mathcal{V}' \setminus \mathcal{V}'_\sigma$. Hence, there is a series of transitions $\bar{t} \parallel \bar{r} \xrightarrow{\tau}_{\mathcal{M}'}^* \hat{t} \parallel \bar{r}^k \xrightarrow{\tau}_{\mathcal{M}'} \bar{t}' \parallel \bar{r}^{k+1}$. Finally, recall that $\bar{s}^{k+1} B \bar{t}'$. Hence, also in this case the proof obligation is satisfied.

Def. 3.5, case 4. *If $\bar{s} C \bar{t}$ and there is an infinite sequence of states $(\bar{t}^k)_{k \in \omega}$ such that $\bar{t} = \bar{t}^0$, $\bar{t}^k \xrightarrow{\tau}_{\mathcal{M}'} \bar{t}^{k+1}$ and $\bar{s} C \bar{t}^k$ for all $k \in \omega$, then there exists a state $\bar{s}'$ such that $\bar{s} \xrightarrow{\tau}_{\mathcal{M}'}^+ \bar{s}'$ and $\bar{s}' C \bar{t}^k$ for some $k \in \omega$. This case is symmetric to the previous case.* $\square$

5. Decomposition of LTS Networks

In Section 4, we discuss the composition of LTS networks, in which a system is constructed by combining components. However, for compositional model checking approaches, it should also be possible to correctly decompose LTS networks. In this case the inter-component laws are already known. Therefore, we can derive a set of interface laws and an interface mapping specifying how the system is decomposed into components.

Consider the decomposition of an LTS network $\mathcal{M} = (\Pi \parallel P, \mathcal{Z})$ into components $\mathcal{N}$ and $\mathcal{O}$ according to some interface mapping σ . First, the synchronisations laws $\mathcal{Z}$ are split into three disjoint sets: 1) $\mathcal{V}^\bullet$ the laws only applying on the processes in Π ; 2) $\bullet\mathcal{W}$ the laws only applying on the processes in P ; and 3) $\mathcal{X}$ the inter-component laws. Next, given two functions $f, g : \mathcal{X} \rightarrow \mathcal{A} \setminus \{\tau\}$ from inter-component laws to interface actions, the inter-component laws are decomposed into sets $\overleftarrow{\mathcal{X}} = \{(v, f(x)) \mid x \in \mathcal{X} \wedge x = (v \parallel w, a) \in \mathcal{X}\}$ and $\overrightarrow{\mathcal{X}} = \{(w, g(x)) \mid x \in \mathcal{X} \wedge x = (v \parallel w, a) \in \mathcal{X}\}$ of interface laws over Π and P , respectively. Finally, the components are defined as $\mathcal{N} = (\Pi, \mathcal{V}^\bullet \cup \overleftarrow{\mathcal{X}})$ and $\mathcal{O} = (P, \bullet\mathcal{W} \cup \overrightarrow{\mathcal{X}})$.

To be able to apply Proposition 4.2 for compositional state space construction, the composition of the decomposed networks must be equivalent to the original system. If this holds we say a decomposition is *consistent* with respect to $\mathcal{M}$.

Definition 5.1 (Consistent Decomposition). Consider a network $\mathcal{M} = (\Pi \parallel P, \mathcal{V}^\bullet \cup \bullet\mathcal{W} \cup \mathcal{X})$ with $\mathcal{X}$ the set of inter-component laws. Say network $\mathcal{M}$ is decomposed into components $\mathcal{N} = (\Pi, \mathcal{V} \cup \overleftarrow{\mathcal{X}})$ and $\mathcal{O} = (P, \mathcal{W} \cup \overrightarrow{\mathcal{X}})$. Decomposition of $\mathcal{M}$ into components $\mathcal{N}$ and $\mathcal{O}$ is called *consistent* with respect to $\mathcal{M}$ iff $\mathcal{M} = \mathcal{N} \parallel \mathcal{O}$, i.e., we must have $\Sigma = \Pi \parallel P$ and $\mathcal{Z} = ((\mathcal{V} \cup \overleftarrow{\mathcal{X}}) \setminus (\mathcal{V} \cup \overleftarrow{\mathcal{X}})_\sigma) \bullet \bullet((\mathcal{W} \cup \overrightarrow{\mathcal{X}}) \setminus (\mathcal{W} \cup \overrightarrow{\mathcal{X}})_\sigma) \cup \sigma(\mathcal{V} \cup \overleftarrow{\mathcal{X}}, \mathcal{W} \cup \overrightarrow{\mathcal{X}})$.

To show that a decomposition is consistent with the original system it is sufficient to show that the set of inter-component laws of the original system is equivalent to the set of inter-component laws generated by the interface-mapping:

Lemma 5.1. *Consider a network $\mathcal{M} = (\Pi \parallel P, \mathcal{V}^\bullet \cup \bullet\mathcal{W} \cup \mathcal{X})$. A consistent decomposition of $\mathcal{M}$ into components $\mathcal{N} = (\Pi, \mathcal{V} \cup \overleftarrow{\mathcal{X}})$ and $\mathcal{O} = (P, \mathcal{W} \cup \overrightarrow{\mathcal{X}})$ with interface mapping $\sigma = \{(f(\bar{v}, a), g(\bar{v}, a), a) \mid (\bar{v}, a) \in \mathcal{X}\}$ is guaranteed if $\mathcal{X} = \sigma(\overleftarrow{\mathcal{X}}, \overrightarrow{\mathcal{X}})$, $\mathcal{A}_\mathcal{V} \cap \mathcal{A}_{\overleftarrow{\mathcal{X}}} = \emptyset$, and $\mathcal{A}_\mathcal{W} \cap \mathcal{A}_{\overrightarrow{\mathcal{X}}} = \emptyset$.*

Proof. The decomposition of $\mathcal{M} = (\Pi \parallel P, \mathcal{V}^\bullet \cup \bullet \mathcal{W} \cup \mathcal{X})$ is consistent iff $\Pi \parallel P = \Pi \parallel P$ and $\mathcal{Z} = ((\mathcal{V} \cup \overleftarrow{\mathcal{X}}) \setminus (\mathcal{V} \cup \overleftarrow{\mathcal{X}})_\sigma)^\bullet \cup \bullet((\mathcal{W} \cup \overrightarrow{\mathcal{X}}) \setminus (\mathcal{W} \cup \overrightarrow{\mathcal{X}})_\sigma) \cup \sigma(\mathcal{V} \cup \overleftarrow{\mathcal{X}}, \mathcal{W} \cup \overrightarrow{\mathcal{X}})$. The former is trivial. Before we continue with the latter let's number the antecedent propositions of lemma: $\mathcal{X} = \sigma(\overleftarrow{\mathcal{X}}, \overrightarrow{\mathcal{X}})$ (1), $\mathcal{A}_\mathcal{V} \cap \mathcal{A}_{\overleftarrow{\mathcal{X}}} = \emptyset$ (2), and $\mathcal{A}_\mathcal{W} \cap \mathcal{A}_{\overrightarrow{\mathcal{X}}} = \emptyset$ (3). We will show that $\mathcal{V}^\bullet = ((\mathcal{V} \cup \overleftarrow{\mathcal{X}}) \setminus (\mathcal{V} \cup \overleftarrow{\mathcal{X}})_\sigma)^\bullet$, $\bullet \mathcal{W} = \bullet((\mathcal{W} \cup \overrightarrow{\mathcal{X}}) \setminus (\mathcal{W} \cup \overrightarrow{\mathcal{X}})_\sigma)$, and $\mathcal{X} = \sigma(\mathcal{V} \cup \overleftarrow{\mathcal{X}}, \mathcal{W} \cup \overrightarrow{\mathcal{X}})$.

By construction of $\overleftarrow{\mathcal{X}}$ and definition of $\mathcal{V}_\sigma$, we have $\mathcal{A}_{\overleftarrow{\mathcal{X}}} = \mathcal{A}_{(\mathcal{V} \cup \overleftarrow{\mathcal{X}})_\sigma}$ and $\mathcal{A}_{\overrightarrow{\mathcal{X}}} = \mathcal{A}_{(\mathcal{W} \cup \overrightarrow{\mathcal{X}})_\sigma}$ (4). Furthermore, from (2) and (3) it follows that $\mathcal{V}$ and $\mathcal{W}$ are disjoint from $\overleftarrow{\mathcal{X}}$ and $\overrightarrow{\mathcal{X}}$ respectively. Thus, $\mathcal{V}$ and $\mathcal{W}$ are disjoint from $(\mathcal{V} \cup \overleftarrow{\mathcal{X}})_\sigma$ and $(\mathcal{W} \cup \overrightarrow{\mathcal{X}})_\sigma$ (5), respectively, implying that $\overleftarrow{\mathcal{X}} = (\mathcal{V} \cup \overleftarrow{\mathcal{X}})_\sigma$ and $\overrightarrow{\mathcal{X}} = (\mathcal{W} \cup \overrightarrow{\mathcal{X}})_\sigma$ (6). It follows that $\mathcal{V}^\bullet \stackrel{(5,6)}{=} ((\mathcal{V} \cup \overleftarrow{\mathcal{X}}) \setminus \overleftarrow{\mathcal{X}})^\bullet \stackrel{(6)}{=} ((\mathcal{V} \cup \overleftarrow{\mathcal{X}}) \setminus (\mathcal{V} \cup \overleftarrow{\mathcal{X}})_\sigma)^\bullet$ and, symmetrically, $\bullet \mathcal{W} \stackrel{(5,6)}{=} \bullet((\mathcal{W} \cup \overrightarrow{\mathcal{X}}) \setminus (\mathcal{W} \cup \overrightarrow{\mathcal{X}})_\sigma)$.

Recall that $\mathcal{V}$ and $\mathcal{W}$ do not have any result actions in common with $\overleftarrow{\mathcal{X}}$ and $\overrightarrow{\mathcal{X}}$, respectively (2,3), and interface actions in σ are produced by the same functions f and g that are used to produce the result actions of sets $\overleftarrow{\mathcal{X}}$ and $\overrightarrow{\mathcal{X}}$, respectively. These two facts and Definition 4.2 (synchronisation via σ) imply that $\sigma(\mathcal{V} \cup \overleftarrow{\mathcal{X}}, \mathcal{W} \cup \overrightarrow{\mathcal{X}}) \stackrel{(5,6,\sigma)}{=} \sigma(\overleftarrow{\mathcal{X}}, \overrightarrow{\mathcal{X}}) \stackrel{(1)}{=} \mathcal{X}$. Hence, the decomposition of $\mathcal{M}$ is consistent if $\mathcal{X} = \sigma(\overleftarrow{\mathcal{X}}, \overrightarrow{\mathcal{X}})$ (1), $\mathcal{A}_\mathcal{V} \cap \mathcal{A}_{\overleftarrow{\mathcal{X}}} = \emptyset$ (2), and $\mathcal{A}_\mathcal{W} \cap \mathcal{A}_{\overrightarrow{\mathcal{X}}} = \emptyset$ (3). $\square$

Indeed, it is possible to derive an inconsistent decomposition as shown in Example 5.1.

Example 5.1 (Inconsistent Decomposition). *Consider a set of inter-component laws $\mathcal{X} = \{(\langle a, b \rangle, c), (\langle b, a \rangle, c)\}$. To generate interface results actions, consider the functions $f(\bar{v}, a) = g(\bar{v}, a) = \alpha$ with unique result actions α based solely on the result action of the input law, i.e., $\forall(\bar{v}', a') \in \mathcal{X}. a' = a \Rightarrow \alpha = f(\bar{v}', a')$. Partitioning the laws results in the sets of interface laws $\overleftarrow{\mathcal{X}} = \{(\langle a \rangle, \gamma), (\langle b \rangle, \gamma)\}$ and $\overrightarrow{\mathcal{X}} = \{(\langle b \rangle, \gamma), (\langle a \rangle, \gamma)\}$. This system implies the interface mapping $\sigma = \{(\gamma, \gamma, c)\}$. The derived set of inter-component laws is $\sigma(\mathcal{V}, \mathcal{W}) = \{(\langle a, a \rangle, c), (\langle a, b \rangle, c), (\langle b, a \rangle, c), (\langle b, b \rangle, c)\} \neq \mathcal{X}$. Hence, this decomposition is not consistent with the original system.*

However, a consistent decomposition can *always* be derived. Propositions 5.1 and 5.2 give functions f and g that guarantee a consistent decomposition. Consider a synchronisation law $(\bar{v} \parallel \bar{w}, a)$, the idea is to encode this synchronisation law directly in the interface mapping. This way it is explicit which interface law corresponds to which inter-component law.

The intuition behind Propositions 5.1 is to encode the synchronisation laws $(\bar{v} \parallel \bar{w}, a) \in \mathcal{X}$ directly in the interface mapping, i.e., we create unique result actions $\alpha_{\bar{v}}$ and $\alpha_{\bar{w}}$ with $(\alpha_{\bar{v}}, \alpha_{\bar{w}}, a) \in \sigma$. This way it is explicit which interface law corresponds to which inter-component law.

Proposition 5.1. *Consider a network $\mathcal{M} = (\Pi \parallel P, \mathcal{V}^\bullet \cup \bullet \mathcal{W} \cup \mathcal{X})$. We define the functions producing interface result actions as $f(\bar{v} \parallel \bar{w}, a) = \alpha_{\bar{v}}$ and $g(\bar{v} \parallel \bar{w}, a) = \alpha_{\bar{w}}$, where $\alpha_{\bar{v}} \notin \mathcal{A}_\mathcal{V} \cup \{\tau\}$ and $\alpha_{\bar{w}} \notin \mathcal{A}_\mathcal{W} \cup \{\tau\}$ are unique interface result actions identified by the corresponding interface law, that is, $\forall(\bar{v}', a) \in \mathcal{V} \cup \overleftarrow{\mathcal{X}}. a = \alpha_{\bar{v}} \implies \bar{v}' = \bar{v}$ and $\forall(\bar{w}', a) \in \mathcal{W} \cup \overrightarrow{\mathcal{X}}. a = \alpha_{\bar{w}} \implies \bar{w}' = \bar{w}$. The decomposition of $\mathcal{M}$ into $\mathcal{M}_\Pi = (\Pi, \mathcal{V} \cup \overleftarrow{\mathcal{X}})$ and $\mathcal{M}_P = (P, \mathcal{W} \cup \overrightarrow{\mathcal{X}})$ given by f and g is consistent.*

Proof. Functions f and g imply interface mapping $\sigma = \{(\alpha_{\bar{v}}, \alpha_{\bar{w}}, a) \mid (\bar{v} \parallel \bar{w}, a) \in \mathcal{X}\}$, and sets of interface laws $\overleftarrow{\mathcal{X}} = \{(\bar{v}, \alpha_{\bar{v}}) \mid (\bar{v} \parallel \bar{w}, a) \in \mathcal{X}\}$ and $\overrightarrow{\mathcal{X}} = \{(\bar{w}, \alpha_{\bar{w}}) \mid (\bar{v} \parallel \bar{w}, a) \in \mathcal{X}\}$.

By Lemma 5.1, we have to show:

- $\mathcal{X} = \sigma(\overleftarrow{\mathcal{X}}, \overrightarrow{\mathcal{X}})$: By (1) definition of $\sigma(\mathcal{V}, \mathcal{W}_\sigma)$, and (2) construction of $\overleftarrow{\mathcal{X}}, \overrightarrow{\mathcal{X}}$, and σ , it follows that

$$\begin{aligned} \sigma(\overleftarrow{\mathcal{X}}, \overrightarrow{\mathcal{X}}) &\stackrel{(1)}{=} \{(\bar{v} \parallel \bar{w}, a) \mid (\bar{v}, \alpha_{\bar{v}}) \in \mathcal{V}_\sigma \wedge (\bar{w}, \alpha_{\bar{w}}) \in \mathcal{W}_\sigma \wedge (\alpha_{\bar{v}}, \alpha_{\bar{w}}, a) \in \sigma\} \\ &\stackrel{(2)}{=} \{(\bar{v} \parallel \bar{w}, a) \mid (\bar{v} \parallel \bar{w}, a) \in \mathcal{X}\} = \mathcal{X} \end{aligned}$$

- $\mathcal{A}_\mathcal{V} \cap \mathcal{A}_{\overleftarrow{\mathcal{X}}} = \emptyset$: Since $\alpha_{\bar{v}} \notin \mathcal{A}_\mathcal{V} \cup \{\tau\}$ the proof follows.
- $\mathcal{A}_\mathcal{W} \cap \mathcal{A}_{\overrightarrow{\mathcal{X}}} = \emptyset$: Since $\alpha_{\bar{w}} \notin \mathcal{A}_\mathcal{W} \cup \{\tau\}$ the proof follows.

□

Example 5.2. Consider an admissible network with the following set of rules:

$$\{(\langle a, \bullet, a \rangle, a), (\langle a, a, \bullet \rangle, a), (\langle b, b, b \rangle, \tau), (\langle c, \bullet, c \rangle, c)\}$$

This set of rules can be decomposed along the lines of Proposition 5.1 as follows:

$$\begin{aligned} \mathcal{V} &= \{(\langle a, a \rangle, a)\} \\ \mathcal{W} &= \{\} \\ \overleftarrow{\mathcal{X}} &= \{(\langle a, \bullet \rangle, \alpha_{\langle a, \bullet \rangle}), (\langle b, b \rangle, \alpha_{\langle b, b \rangle}), (\langle c, \bullet \rangle, \alpha_{\langle c, \bullet \rangle})\} \\ \overrightarrow{\mathcal{X}} &= \{(\langle a \rangle, \alpha_{\langle a \rangle}), (\langle b \rangle, \alpha_{\langle b \rangle}), (\langle c \rangle, \alpha_{\langle c \rangle})\} \\ \sigma &= \{(\alpha_{\langle a, \bullet \rangle}, \alpha_{\langle a \rangle}, a), (\alpha_{\langle b, b \rangle}, \alpha_{\langle b \rangle}, \tau), (\alpha_{\langle c, \bullet \rangle}, \alpha_{\langle c \rangle}, c)\} \end{aligned}$$

Proposition 5.2 proposes an alternative decomposition that is implemented in CADP's smart reduction [15]. The idea is (1) to generate only interface mapping rules of the form (a, a, b) , so that components always synchronise through a common label a , while (2) keeping a equal to b and thus avoiding α labels whenever possible. Rules in this simple form make the decomposition more understandable by users.

Proposition 5.2. Consider a network $\mathcal{M} = (\Pi \parallel P, \mathcal{V}^\bullet \cup \bullet\mathcal{W} \cup \mathcal{X})$. We define the functions producing interface result actions as

$$f(\bar{v}, a) = g(\bar{v}, a) = \begin{cases} a & \text{if } \text{visible_unique}(a, \mathcal{V}) \\ \alpha_{(\bar{v}, a)} & \text{otherwise} \end{cases}$$

where each $\alpha_{(\bar{v}, a)} \notin \mathcal{A}_\mathcal{V} \cup \mathcal{A}_\mathcal{W} \cup \{\tau\}$ is a unique interface result action identified by the corresponding inter component law, that is, $\forall(\bar{v}', a) \in \overleftarrow{\mathcal{X}} \cup \overrightarrow{\mathcal{X}}. a = \alpha_{\bar{v}} \implies \bar{v}' = \bar{v}$, and where $\text{visible_unique}(a, \mathcal{V})$ is defined by the following predicate:

$$a \neq \tau \wedge \forall(\bar{v}, a), (\bar{v}', a) \in \mathcal{V}. \bar{v} = \bar{v}'.$$

The decomposition of $\mathcal{M}$ into $\mathcal{M}_\Pi = (\Pi, \mathcal{V} \cup \overleftarrow{\mathcal{X}})$ and $\mathcal{M}_P = (P, \mathcal{W} \cup \overrightarrow{\mathcal{X}})$ given by f and g is consistent.

The proof of Proposition 5.2 is similar to the proof of Proposition 5.1. The most relevant difference is the presence of $(a, a, a) \in \sigma$ if $(\bar{v} \parallel \bar{w}, a) \in \mathcal{X}$ and a is unique in $\mathcal{V}^\bullet \cup \bullet \mathcal{W} \cup \mathcal{X}$ ($\text{visible_unique}(a, \mathcal{V}^\bullet \cup \bullet \mathcal{W} \cup \mathcal{X})$ holds). In this case we must also have $\mathcal{A}_{\mathcal{V}} \cap \mathcal{A}_{\overleftarrow{\mathcal{X}}} = \emptyset$ and $\mathcal{A}_{\mathcal{W}} \cap \mathcal{A}_{\overrightarrow{\mathcal{X}}} = \emptyset$, otherwise a contradiction with $\text{visible_unique}(a, \mathcal{V}^\bullet \cup \bullet \mathcal{W} \cup \mathcal{X})$ can be derived.

The decomposition of Proposition 5.2 implies the interface mapping

$$\begin{aligned} \sigma = & \{(\alpha_{\bar{v} \parallel \bar{w}}, \alpha_{\bar{v} \parallel \bar{w}}, a) \mid (\bar{v} \parallel \bar{w}, a) \in \mathcal{X} \wedge \neg \text{visible_unique}(a, \mathcal{V}^\bullet \cup \bullet \mathcal{W} \cup \mathcal{X})\} \\ & \cup \{(a, a, a) \mid (\bar{v} \parallel \bar{w}, a) \in \mathcal{X} \wedge \text{visible_unique}(a, \mathcal{V}^\bullet \cup \bullet \mathcal{W} \cup \mathcal{X})\} \end{aligned}$$

and sets of interface laws

$$\begin{aligned} \mathcal{V}_\sigma = & \{(\bar{v}, \alpha_{\bar{v} \parallel \bar{w}}) \mid (\bar{v} \parallel \bar{w}, a) \in \mathcal{X} \wedge \neg \text{visible_unique}(a, \mathcal{V}^\bullet \cup \bullet \mathcal{W} \cup \mathcal{X})\} \\ & \cup \{(\bar{v}, a) \mid (\bar{v} \parallel \bar{w}, a) \in \mathcal{X} \wedge \text{visible_unique}(a, \mathcal{V}^\bullet \cup \bullet \mathcal{W} \cup \mathcal{X})\} \\ \mathcal{W}_\sigma = & \{(\bar{w}, \alpha_{\bar{v} \parallel \bar{w}}) \mid (\bar{v} \parallel \bar{w}, a) \in \mathcal{X} \wedge \neg \text{visible_unique}(a, \mathcal{V}^\bullet \cup \bullet \mathcal{W} \cup \mathcal{X})\} \\ & \cup \{(\bar{w}, a) \mid (\bar{v} \parallel \bar{w}, a) \in \mathcal{X} \wedge \text{visible_unique}(a, \mathcal{V}^\bullet \cup \bullet \mathcal{W} \cup \mathcal{X})\} \end{aligned}$$

Example 5.3. Consider again the admissible network of Example 5.2. Its set of rules can be decomposed along the lines of Proposition 5.2 as follows, using the same definition of $\mathcal{V}$ and $\mathcal{W}$:

$$\begin{aligned} \overleftarrow{\mathcal{X}} = & \{(\langle a, \bullet \rangle, \alpha_{\langle a, \bullet, a \rangle}), (\langle b, b \rangle, \alpha_{\langle b, b, b \rangle}), (\langle c, \bullet \rangle, c)\} \\ \overrightarrow{\mathcal{X}} = & \{(\langle a \rangle, \alpha_{\langle a, \bullet, a \rangle}), (\langle b \rangle, \alpha_{\langle b, b, b \rangle}), (\langle c \rangle, c)\} \\ \sigma = & \{(\alpha_{\langle a, \bullet, a \rangle}, \alpha_{\langle a, \bullet, a \rangle}, a), (\alpha_{\langle b, b, b \rangle}, \alpha_{\langle b, b, b \rangle}, \tau), (c, c, c)\} \end{aligned}$$

Preservation of Admissibility. Proposition 5.3 shows that LTS networks resulting from the consistent decomposition of an admissible LTS network are also admissible. Hence, consistent decomposition is compatible with the compositional verification approaches presented in [20].

Proposition 5.3. Consider an admissible LTS network $\mathcal{M} = (\Pi \parallel P, \mathcal{V}^\bullet \cup \bullet \mathcal{W} \cup \mathcal{X})$ of length $n + m$. If the decomposition is consistent, then the decomposed networks $\mathcal{M}_\Pi = (\Pi, \mathcal{V} \cup \overleftarrow{\mathcal{X}})$ and $\mathcal{M}_P = (P, \mathcal{W} \cup \overrightarrow{\mathcal{X}})$ are also admissible.

Proof. We show that $\mathcal{M}_\Pi$ and $\mathcal{M}_P$ satisfy Definition 3.4:

No synchronisation and renaming of τ 's. Let $(\bar{v}, a) \in \mathcal{V} \cup \overleftarrow{\mathcal{X}}$ be a synchronisation law such that $\bar{v}_i = \tau$ for some $i \in 1..n$. We distinguish two cases:

- $(\bar{v}, a) \in \overleftarrow{\mathcal{X}}$. Since $(\bar{v}, a)$ is an interface law and the decomposition is consistent, its result action a may not be τ . However, since $\mathcal{M}$ is admissible, no renaming of τ 's is allowed. By contradiction it follows that $(\bar{v}, a) \notin \overleftarrow{\mathcal{X}}$ completing this case.
- $(\bar{v}, a) \in \mathcal{V}$. By construction, there exists a law $(\bar{v} \parallel \bullet^m, a) \in \mathcal{V}^\bullet$. Since $\mathcal{V}^\bullet \subseteq \mathcal{V}^\bullet \cup \bullet \mathcal{W} \cup \mathcal{X}$, by admissibility of $\mathcal{M}$, we have $\forall j \in 1..n. \bar{v}_j \neq \bullet \implies i = j$ (no synchronisation of τ 's) and $a = \tau$ (no renaming of τ 's).

Hence, $\mathcal{M}_\Pi$ does not synchronise or rename τ 's. The proof for $\mathcal{M}_P$ is similar.

No cutting of τ 's. Let Π_i be a process with $i \in 1..n$ such that $\tau \in \mathcal{A}_{\Pi_i}$. Since $\mathcal{M}$ is admissible there exists a law $(\bar{v} \parallel \bar{w}, a) \in \mathcal{V}^\bullet \cup \bullet\mathcal{W} \cup \mathcal{X}$ such that $(\bar{v} \parallel \bar{u})_i = \tau$. We distinguish three cases:

- $(\bar{v} \parallel \bar{w}, a) \in \mathcal{V}^\bullet$. Since $(\bar{v} \parallel \bar{w})_i = \tau$ and $i \leq n$ it follows that $\bar{v}_i = \tau$. By construction of $\mathcal{V}^\bullet$, there is a $(\bar{v}, a) \in V$ with $\bar{v}_i = \tau$.
- $(\bar{v} \parallel \bar{w}, a) \in \bullet\mathcal{W}$. In this case we must have $i > n$ which contradicts our assumption: $i \in 1..n$. The proof follows by contradiction.
- $(\bar{v} \parallel \bar{w}, a) \in \mathcal{X}$. Then, $(\bar{v} \parallel \bar{w}, a)$ is an inter-component law with at least one participating process for each component. Hence, there exists a $j \in (n+1)..m$ such that $(\bar{v} \parallel \bar{w})_j \neq \bullet$. Moreover, since $\mathcal{M}$ is admissible, no synchronisation of τ 's are allowed. Therefore, since $(\bar{v} \parallel \bar{w})_j \neq \bullet$, we must have $j = i$. However, this would mean $j \in 1..n$, contradicting $j \in (n+1)..m$. By contradiction the proof follows.

We conclude that $\mathcal{M}_\Pi$ does not cut τ 's. The proof for $\mathcal{M}_P$ is symmetrical.

All three admissibility properties hold for $\mathcal{M}_\Pi$ and $\mathcal{M}_P$. Hence, the networks resulting from the decomposition satisfy Definition 3.4. $\square$

6. Associative and Commutative LTS Network Composition

In this section we create an instance of the composition operator that is commutative and associative. This operator uses an interface mapping that synchronises on a common alphabet for interface actions. It is desirable for the composition operator to be both associative and commutative as then the composition order is irrelevant with respect to the resulting system. This allows users to select any composition order; in some situation a particular order may be better than others [23].

Definition 6.1 (Composition with Synchronisation on a Common Alphabet).

Consider LTS networks $\mathcal{M}_\Pi = (\Pi, \mathcal{V})$ of size n and $\mathcal{M}_P = (P, \mathcal{W})$ of size m . Take the interface mapping $\sigma = \{(a, a, a) \mid a \in (\mathcal{A}_\mathcal{V} \cap \mathcal{A}_\mathcal{W}) \setminus \{\tau\}\}$. The *composition on a common alphabet* of $\mathcal{M}_\Pi$ and $\mathcal{M}_P$ is defined as $\mathcal{M}_\Pi \parallel \mathcal{M}_P = \mathcal{M}_\Pi \parallel_\sigma \mathcal{M}_P$.

Associativity. The intuition behind the associativity of LTS network composition is that vector concatenation is associative and synchronisation on the common alphabet is insensitive to the order of composition. Thus, the concatenation of process vectors and synchronisation vectors enjoy the associativity property. The challenge, however, is to show that the $\bullet$ and $\sigma(\dots, \dots)$ operations support the mathematical properties needed for associativity of the composition of sets of synchronisation laws.

Given two networks $\mathcal{M}_\Pi = (\Pi, \mathcal{V})$ and $\mathcal{M}_P = (P, \mathcal{W})$ the composition of two sets of synchronisation laws $\mathcal{V} \parallel \mathcal{W}$ consists of the union of three sets: two describing independent behaviour $(\mathcal{V} \setminus \mathcal{V}_\sigma)^\bullet$ and $\bullet(\mathcal{W} \setminus \mathcal{W}_\sigma)$, and one describing synchronising behaviour $\sigma(\mathcal{V}, \mathcal{W})$ (Definition 4.2). When three networks are composed one (inner) composition is performed before the other. The outer composition applies the operators $\bullet$ and $\sigma(\dots, \dots)$ on a union of sets. We show how these operators distribute over set union.

Lemma 6.1. Consider sets of synchronisation laws $\mathcal{V}$ and $\mathcal{W}$ with synchronisation vectors of the same size. Padding of n $\bullet$'s distributes over set union:

$$\bullet(\mathcal{V} \cup \mathcal{W}) = \bullet\mathcal{V} \cup \bullet\mathcal{W}, \text{ and } (\mathcal{V} \cup \mathcal{W})^\bullet = \mathcal{V}^\bullet \cup \mathcal{W}^\bullet.$$

Proof. The proof of $(\mathcal{V} \cup \mathcal{W})^\bullet = \mathcal{V}^\bullet \cup \mathcal{W}^\bullet$ is analog to the proof of $\bullet(\mathcal{V} \cup \mathcal{W}) = \bullet\mathcal{V} \cup \bullet\mathcal{W}$. We only prove $\bullet(\mathcal{V} \cup \mathcal{W}) = \bullet\mathcal{V} \cup \bullet\mathcal{W}$ here. The proof follows from (1) application of the definition of $\bullet$ and (2) splitting of the set on $\mathcal{V} \cup \mathcal{W}$:

$$\begin{aligned} \bullet(\mathcal{V} \cup \mathcal{W}) &\stackrel{(1)}{=} \{(\bullet^n \parallel \bar{v}, a) \mid (\bar{v}, a) \in \mathcal{V} \cup \mathcal{W}\} \\ &\stackrel{(2)}{=} \{(\bullet^n \parallel \bar{v}, a) \mid (\bar{v}, a) \in \mathcal{V}\} \cup \{(\bullet^n \parallel \bar{v}, a) \mid (\bar{v}, a) \in \mathcal{W}\} \\ &\stackrel{(1)}{=} \bullet\mathcal{V} \cup \bullet\mathcal{W} \end{aligned}$$

□

Lemma 6.2. Consider an interface mapping σ and sets of synchronisation laws $\mathcal{V}$, $\mathcal{W}$, and $\mathcal{X}$. The $\sigma(\dots, \dots)$ operation distributes over set union as follows:

$$\sigma(\mathcal{V}, \mathcal{W} \cup \mathcal{X}) = \sigma(\mathcal{V}, \mathcal{W}) \cup \sigma(\mathcal{V}, \mathcal{X})$$

Proof. The proof follows from (1) application of the definition of $\sigma(\dots, \dots)$ and (2) splitting of the set on $\mathcal{W} \cup \mathcal{X}$:

$$\begin{aligned} \sigma(\mathcal{V}, \mathcal{W} \cup \mathcal{X}) &\stackrel{(1)}{=} \{(\bar{v} \parallel \bar{w}, a) \mid (\bar{v}, \alpha) \in \mathcal{V} \wedge (\bar{w}, \beta) \in \mathcal{W} \cup \mathcal{X} \wedge (\alpha, \beta, a) \in \sigma\} \\ &\stackrel{(2)}{=} \{(\bar{v} \parallel \bar{w}, a) \mid (\bar{v}, \alpha) \in \mathcal{V} \wedge (\bar{w}, \beta) \in \mathcal{W} \wedge (\alpha, \beta, a) \in \sigma\} \cup \\ &\quad \{(\bar{v} \parallel \bar{w}, a) \mid (\bar{v}, \alpha) \in \mathcal{V} \wedge (\bar{w}, \beta) \in \mathcal{X} \wedge (\alpha, \beta, a) \in \sigma\} \\ &\stackrel{(1)}{=} \sigma(\mathcal{V}, \mathcal{W}) \cup \sigma(\mathcal{V}, \mathcal{X}) \end{aligned}$$

□

Now we prove that the composition of LTS networks with synchronisation on the common alphabet is associative.

Proposition 6.1. For all LTS networks $\mathcal{M}_\Pi = (\Pi, \mathcal{V})$, $\mathcal{M}_\mathcal{P} = (\mathcal{P}, \mathcal{W})$, and $\mathcal{M}_\Sigma = (\Sigma, \mathcal{X})$ of sizes n , m , and o respectively. The composition of LTS networks following Definition 6.1 is associative, i.e., it holds that $(\mathcal{M}_\Pi \parallel \mathcal{M}_\mathcal{P}) \parallel \mathcal{M}_\Sigma = \mathcal{M}_\Pi \parallel (\mathcal{M}_\mathcal{P} \parallel \mathcal{M}_\Sigma)$.

Proof. If the networks $(\mathcal{M}_\Pi \parallel \mathcal{M}_\mathcal{P}) \parallel \mathcal{M}_\Sigma$ and $\mathcal{M}_\Pi \parallel (\mathcal{M}_\mathcal{P} \parallel \mathcal{M}_\Sigma)$ are equivalent, then this means that their process vectors are equivalent and their sets of synchronisation laws are equivalent.

The process vectors are equivalent due to associativity of the vector concatenation operator $\parallel$:

$$\Pi \parallel (\mathcal{P} \parallel \Sigma) = (\Pi \parallel \mathcal{P}) \parallel \Sigma$$

Next, we show $\mathcal{V} \parallel (\mathcal{W} \parallel \mathcal{X}) = (\mathcal{V} \parallel \mathcal{W}) \parallel \mathcal{X}$. First, given a set of laws $\mathcal{V}$, we will introduce an alternative notation for $\mathcal{V} \setminus \mathcal{Y}_\sigma$ in the context of composition, $\mathcal{V} \parallel \mathcal{Z}$, of $\mathcal{V}$ with a set of laws $\mathcal{Z}$. We will write $\mathcal{V} \setminus \mathcal{Z}_\mathcal{A}$ to emphasize the relevance of the alphabet of $\mathcal{V}$ that is in common with that of $\mathcal{Z}$. We define $\mathcal{V} \setminus \mathcal{Z}_\mathcal{A} = \{(\bar{y}, a) \in \mathcal{V} \mid a \notin \mathcal{A}_\mathcal{Z}\}$. The set $\mathcal{V} \setminus \mathcal{Z}_\mathcal{A}$ is equivalent to $\mathcal{V} \setminus \mathcal{Y}_\sigma$:

$$\mathcal{V} \setminus \mathcal{Z}_\mathcal{A} = \{(\bar{y}, a) \in \mathcal{V} \mid a \notin \mathcal{A}_\mathcal{Z}\} = \mathcal{V} \setminus \{(\bar{y}, a) \in \mathcal{V} \mid (a, a, a) \in \sigma\} = \mathcal{V} \setminus \mathcal{Y}_\sigma$$

The set $\mathcal{Z} \setminus \mathcal{Y}_\mathcal{A}$ is defined similarly.

The associativity proof proceeds as follows. Following Definition 4.2, we break the set $\mathcal{V} \parallel (\mathcal{W} \parallel \mathcal{X})$ down to seven partitions, we will show that there is a one-to-one mapping of these partitions to one of the seven partitions of $(\mathcal{V} \parallel \mathcal{W}) \parallel \mathcal{X}$. Each of the rewrite equations consists of four steps:

- (1) unfold the outer definition of $\bullet$ or $\sigma(\dots, \dots)$
- (2) unfold the inner definition of $\bullet$ or $\sigma(\dots, \dots)$
- (3) apply associativity of vector concatenation and the inner definition of $\bullet$ or $\sigma(\dots, \dots)$
- (4) apply the outer definition of $\bullet$ or $\sigma(\dots, \dots)$

Furthermore, in cases 2a and 3c the following property of composition of sets of laws is applied in steps (2) and (3) respectively: $a \notin \mathcal{A}_{\mathcal{W} \parallel \mathcal{X}} = a \notin \mathcal{A}_\mathcal{W} \cup \mathcal{A}_\mathcal{X} = a \notin \mathcal{A}_\mathcal{W} \wedge a \notin \mathcal{A}_\mathcal{X}$

The partitioning and partition mapping proceed as follows.

1. $\sigma(\mathcal{V}, \mathcal{W} \parallel \mathcal{X})$ is partitioned, according to Lemma 6.2, into:

- (a) $\sigma(\mathcal{V}, \sigma(\mathcal{W}, \mathcal{X}))$, the set of laws specifying synchronisations involving all networks.

$$\begin{aligned} & \sigma(\mathcal{V}, \sigma(\mathcal{W}, \mathcal{X})) \\ & \stackrel{(1)}{=} \{\bar{v} \parallel (\bar{w} \parallel \bar{x}) \mid (\bar{v}, a) \in \mathcal{V} \wedge (\bar{w} \parallel \bar{x}, a) \in \sigma(\mathcal{W}, \mathcal{X})\} \\ & \stackrel{(2)}{=} \{\bar{v} \parallel (\bar{w} \parallel \bar{x}) \mid (\bar{v}, a) \in \mathcal{V} \wedge (\bar{w}, a) \in \mathcal{W} \wedge (\bar{x}, a) \in \mathcal{X}\} \\ & \stackrel{(3)}{=} \{(\bar{v} \parallel \bar{w}) \parallel \bar{x} \mid (\bar{v} \parallel \bar{w}, a) \in \sigma(\mathcal{V}, \mathcal{W}) \wedge (\bar{x}, a) \in \mathcal{X}\} \\ & \stackrel{(4)}{=} \sigma(\sigma(\mathcal{V}, \mathcal{W}), \mathcal{X}) \end{aligned}$$

- (b) $\sigma(\mathcal{V}, (\mathcal{W} \setminus \mathcal{X}_\mathcal{A})^\bullet)$, the set of laws synchronising only $\mathcal{M}_\Pi$ and $\mathcal{M}_\mathcal{P}$.

$$\begin{aligned} & \sigma(\mathcal{V}, (\mathcal{W} \setminus \mathcal{X}_\mathcal{A})^\bullet) \\ & \stackrel{(1)}{=} \{\bar{v} \parallel (\bar{w} \parallel \bullet^o) \mid (\bar{v}, a) \in \mathcal{V} \wedge (\bar{w} \parallel \bullet^o, a) \in (\mathcal{W} \setminus \mathcal{X}_\mathcal{A})^\bullet\} \\ & \stackrel{(2)}{=} \{\bar{v} \parallel (\bar{w} \parallel \bullet^o) \mid (\bar{v}, a) \in \mathcal{V} \wedge (\bar{w}, a) \in \mathcal{W} \wedge a \notin \mathcal{A}_\mathcal{X}\} \\ & \stackrel{(3)}{=} \{(\bar{v} \parallel \bar{w}) \parallel \bullet^o \mid (\bar{v} \parallel \bar{w}, a) \in \sigma(\mathcal{V}, \mathcal{W}) \wedge a \notin \mathcal{A}_\mathcal{X}\} \\ & \stackrel{(4)}{=} (\sigma(\mathcal{V}, \mathcal{W}) \setminus \mathcal{X}_\mathcal{A})^\bullet \end{aligned}$$

(c) $\sigma(\mathcal{V}, \bullet(\mathcal{X} \setminus \mathcal{W}_{\mathcal{A}}))$, the set of laws synchronising only $\mathcal{M}_{\Pi}$ and $\mathcal{M}_{\Sigma}$.

$$\begin{aligned}
& \sigma(\mathcal{V}, \bullet(\mathcal{X} \setminus \mathcal{W}_{\mathcal{A}})) \\
& \stackrel{(1)}{=} \{ \bar{v} \parallel (\bullet^m \parallel \bar{x}) \mid (\bar{v}, a) \in \mathcal{V} \wedge (\bullet^m \parallel \bar{x}, a) \in \mathcal{X} \setminus \mathcal{W}_{\mathcal{A}} \} \\
& \stackrel{(2)}{=} \{ \bar{v} \parallel (\bullet^m \parallel \bar{x}) \mid (\bar{v}, a) \in \mathcal{V} \wedge (\bar{x}, a) \in \mathcal{X} \wedge a \notin \mathcal{A}_{\mathcal{W}} \} \\
& \stackrel{(3)}{=} \{ (\bar{v} \parallel \bullet^m) \parallel \bar{x} \mid (\bar{v} \parallel \bullet^m) \in (\mathcal{V} \setminus \mathcal{W}_{\mathcal{A}})^{\bullet} \wedge (\bar{x}, a) \in \mathcal{X} \} \\
& \stackrel{(4)}{=} \sigma((\mathcal{V} \setminus \mathcal{W}_{\mathcal{A}})^{\bullet}, \mathcal{X})
\end{aligned}$$

2. $(\mathcal{V} \setminus (\mathcal{W} \parallel \mathcal{X})_{\mathcal{A}})^{\bullet}$ requires no partitioning:

(a) $(\mathcal{V} \setminus (\mathcal{W} \parallel \mathcal{X})_{\mathcal{A}})^{\bullet}$, the set of laws specifying the independent behaviour of $\mathcal{M}_{\Pi}$.

$$\begin{aligned}
& (\mathcal{V} \setminus (\mathcal{W} \parallel \mathcal{X})_{\mathcal{A}})^{\bullet} \\
& \stackrel{(1)}{=} \{ \bar{v} \parallel \bullet^{m+o} \mid (\bar{v}, a) \in \mathcal{V} \wedge a \notin \mathcal{A}_{\mathcal{W} \parallel \mathcal{X}} \} \\
& \stackrel{(2)}{=} \{ \bar{v} \parallel (\bullet^m \parallel \bullet^o) \mid (\bar{v}, a) \in \mathcal{V} \wedge a \notin \mathcal{A}_{\mathcal{W}} \wedge a \notin \mathcal{A}_{\mathcal{X}} \} \\
& \stackrel{(3)}{=} \{ (\bar{v} \parallel \bullet^m) \parallel \bullet^o \mid (\bar{v} \parallel \bullet^m, a) \in (\mathcal{V} \setminus \mathcal{W}_{\mathcal{A}})^{\bullet} \wedge a \notin \mathcal{A}_{\mathcal{X}} \} \\
& \stackrel{(4)}{=} ((\mathcal{V} \setminus \mathcal{W}_{\mathcal{A}})^{\bullet} \setminus \mathcal{X}_{\mathcal{A}})^{\bullet}
\end{aligned}$$

3. $\bullet((\mathcal{W} \parallel \mathcal{X}) \setminus \mathcal{V}_{\mathcal{A}})$ is partitioned, applying Lemma 6.1, into:

(a) $\bullet(\sigma(\mathcal{W}, \mathcal{X}) \setminus \mathcal{V}_{\mathcal{A}})$, the set of laws synchronising only $\mathcal{M}_{\Pi}$ and $\mathcal{M}_{\Sigma}$.

$$\begin{aligned}
& \bullet(\sigma(\mathcal{W}, \mathcal{X}) \setminus \mathcal{V}_{\mathcal{A}}) \\
& \stackrel{(1)}{=} \{ (\bullet^n \parallel (\bar{w} \parallel \bar{x}), a) \mid (\bar{w} \parallel \bar{x}, a) \in \sigma(\mathcal{W}, \mathcal{X}) \wedge a \notin \mathcal{A}_{\mathcal{V}} \} \\
& \stackrel{(2)}{=} \{ (\bullet^n \parallel (\bar{w} \parallel \bar{x}), a) \mid (\bar{w}, a) \in \mathcal{W} \wedge (\bar{x}, a) \in \mathcal{X} \wedge a \notin \mathcal{A}_{\mathcal{V}} \} \\
& \stackrel{(3)}{=} \{ ((\bullet^n \parallel \bar{w}) \parallel \bar{x}, a) \mid (\bullet^n \parallel \bar{w}, a) \in \bullet(\mathcal{W} \setminus \mathcal{V}_{\mathcal{A}}) \wedge (\bar{x}, a) \in \mathcal{X} \} \\
& \stackrel{(4)}{=} \sigma(\bullet(\mathcal{W} \setminus \mathcal{V}_{\mathcal{A}}), \mathcal{X})
\end{aligned}$$

(b) $\bullet((\mathcal{W} \setminus \mathcal{X}_{\mathcal{A}})^{\bullet} \setminus \mathcal{V}_{\mathcal{A}})$, the set of laws regarding independent behaviour of $\mathcal{M}_{\Pi}$.

$$\begin{aligned}
& \bullet((\mathcal{W} \setminus \mathcal{X}_{\mathcal{A}})^{\bullet} \setminus \mathcal{V}_{\mathcal{A}}) \\
& \stackrel{(1)}{=} \{ (\bullet^n \parallel (\bar{w} \parallel \bullet^o), a) \mid (\bar{w} \parallel \bullet^o, a) \in (\mathcal{W} \setminus \mathcal{X}_{\mathcal{A}})^{\bullet} \wedge a \notin \mathcal{A}_{\mathcal{V}} \} \\
& \stackrel{(2)}{=} \{ (\bullet^n \parallel (\bar{w} \parallel \bullet^o), a) \mid (\bar{w}, a) \in \mathcal{W} \wedge a \notin \mathcal{A}_{\mathcal{V}} \wedge a \notin \mathcal{A}_{\mathcal{X}} \} \\
& \stackrel{(3)}{=} \{ ((\bullet^n \parallel \bar{w}) \parallel \bullet^o, a) \mid (\bullet^n \parallel \bar{w}, a) \in \bullet(\mathcal{W} \setminus \mathcal{V}_{\mathcal{A}}) \wedge a \notin \mathcal{A}_{\mathcal{X}} \} \\
& \stackrel{(4)}{=} (\bullet(\mathcal{W} \setminus \mathcal{V}_{\mathcal{A}}) \setminus \mathcal{X}_{\mathcal{A}})^{\bullet}
\end{aligned}$$

$$\begin{aligned}
(c) \quad & \bullet(\bullet(\mathcal{X} \setminus \mathcal{W}_A) \setminus \mathcal{V}_A), \text{ the set of laws specifying independent behaviour of } \mathcal{M}_\Sigma. \\
& \bullet(\bullet(\mathcal{X} \setminus \mathcal{W}_A) \setminus \mathcal{V}_A) \\
& \stackrel{(1)}{=} \{(\bullet^n \parallel (\bullet^m \parallel \bar{x}), a) \mid (\bullet^m \parallel \bar{x}, a) \in \bullet(\mathcal{X} \setminus \mathcal{W}_A) \wedge a \notin \mathcal{A}_\mathcal{V}\} \\
& \stackrel{(2)}{=} \{(\bullet^n \parallel (\bullet^m \parallel \bar{x}), a) \mid (\bar{x}, a) \in \mathcal{X} \wedge a \notin \mathcal{A}_\mathcal{W} \wedge a \notin \mathcal{A}_\mathcal{V}\} \\
& \stackrel{(3)}{=} \{(\bullet^{n+m} \parallel \bar{x}, a) \mid (\bar{x}, a) \in \mathcal{X} \wedge a \notin \mathcal{A}_{\mathcal{V} \parallel \mathcal{W}}\} \\
& \stackrel{(4)}{=} \bullet(\mathcal{X} \setminus (\mathcal{V} \parallel \mathcal{W})_A)
\end{aligned}$$

These equations constitute a one-to-one mapping between the partitioning of $\mathcal{V} \parallel (\mathcal{W} \parallel \mathcal{X})$ and that of $(\mathcal{V} \parallel \mathcal{W}) \parallel \mathcal{X}$. Therefore, we have $\mathcal{V} \parallel (\mathcal{W} \parallel \mathcal{X}) = (\mathcal{V} \parallel \mathcal{W}) \parallel \mathcal{X}$.

Since both $\Pi \parallel (\mathcal{P} \parallel \Sigma) = (\Pi \parallel \mathcal{P}) \parallel \Sigma$ and $\mathcal{V} \parallel (\mathcal{W} \parallel \mathcal{X}) = (\mathcal{V} \parallel \mathcal{W}) \parallel \mathcal{X}$ it follows that $(\mathcal{M}_\Pi \parallel \mathcal{M}_\mathcal{P}) \parallel \mathcal{M}_\Sigma = \mathcal{M}_\Pi \parallel (\mathcal{M}_\mathcal{P} \parallel \mathcal{M}_\Sigma)$. $\square$

Commutativity. It is clear that composition of LTS networks is not commutative as is indicated by Example 6.1.

Example 6.1. Let $\mathcal{M}_\Pi = (\Pi, \mathcal{V})$ and $\mathcal{M}_\mathcal{P} = (\mathcal{P}, \mathcal{W})$ be two LTS networks. Furthermore, consider compositions $\mathcal{M}_1 = \mathcal{M}_\Pi \parallel \mathcal{M}_\mathcal{P}$ and $\mathcal{M}_2 = \mathcal{M}_\mathcal{P} \parallel \mathcal{M}_\Pi$. The network $\mathcal{M}_1$ has process vector $\Pi \parallel \mathcal{P}$ while $\mathcal{M}_2$ has process vector $\mathcal{P} \parallel \Pi$. Unless $\mathcal{M}_\Pi = \mathcal{M}_\mathcal{P}$, $\mathcal{M}_1$ and $\mathcal{M}_2$ are strictly not equivalent. Similarly, the synchronisation laws of both composite networks are in a different order.

Network composition as defined in Definition 6.1 is, however, commutative with respect to the system LTS of the composition. That is, for the global behaviour of the composition of the networks, it does not matter in which order the networks are composed. We first prove that this network composition is commutative with respect to (strong) bisimulation [4] in Proposition 6.2. Afterwards, we will propose an adaption of the definition of LTS network, fixing the ordering issue by replacing vectors with indexed families gaining a commutative operator for composition of LTS networks with synchronisation on the common alphabet.

Proposition 6.2. Let $\mathcal{M}_\Pi = (\Pi, \mathcal{V})$ and $\mathcal{M}_\mathcal{P} = (\mathcal{P}, \mathcal{W})$ be LTS networks of sizes n and m respectively. Composition of LTS networks according to Definition 6.1 is commutative with respect to (strong) bisimulation, i.e., it holds that $\mathcal{G}_{\mathcal{M}_\Pi \parallel \mathcal{M}_\mathcal{P}} \stackrel{\leftrightarrow}{=} \mathcal{G}_{\mathcal{M}_\mathcal{P} \parallel \mathcal{M}_\Pi}$.

Proof. Take the relation $C = \{(\bar{s} \parallel \bar{t}, \bar{t} \parallel \bar{s}) \mid \bar{s} \in \mathcal{G}_{\mathcal{M}_\Pi} \wedge \bar{t} \in \mathcal{G}_{\mathcal{M}_\mathcal{P}}\}$. The relation C is a (strong) bisimulation relation.

- C relates the initial states of $\mathcal{M}_\Pi$ and $\mathcal{M}_\mathcal{P}$. Since every state $\bar{s} \parallel \bar{t} \in \mathcal{I}_{\mathcal{M}_\Pi \parallel \mathcal{M}_\mathcal{P}}$ is related by C to state $\bar{t} \parallel \bar{s} \in \mathcal{I}_{\mathcal{M}_\mathcal{P} \parallel \mathcal{M}_\Pi}$ and vice versa.
- If $\bar{s} \parallel \bar{t} \xrightarrow{C} \bar{t} \parallel \bar{s}$ and $\bar{s} \parallel \bar{t} \xrightarrow{a}_{\mathcal{M}_\Pi \parallel \mathcal{M}_\mathcal{P}} \bar{s}' \parallel \bar{t}'$ then $\bar{t} \parallel \bar{s} \xrightarrow{a}_{\mathcal{M}_\mathcal{P} \parallel \mathcal{M}_\Pi} \bar{t}'' \parallel \bar{s}'' \wedge \bar{s}' \parallel \bar{t}' \xrightarrow{C} \bar{t}'' \parallel \bar{s}''$. Let $(\bar{v} \parallel \bar{w}, a) \in \mathcal{V} \parallel \mathcal{W}$ be the law enabling the transition $\bar{s} \parallel \bar{t} \xrightarrow{a}_{\mathcal{M}_\Pi \parallel \mathcal{M}_\mathcal{P}} \bar{s}' \parallel \bar{t}'$. It follows that there is a law $(\bar{w} \parallel \bar{v}, a) \in \mathcal{W} \parallel \mathcal{V}$ that enables the transition $\bar{t} \parallel \bar{s} \xrightarrow{a}_{\mathcal{M}_\mathcal{P} \parallel \mathcal{M}_\Pi} \bar{t}' \parallel \bar{s}'$. As $\bar{s}' \parallel \bar{t}' \xrightarrow{C} \bar{t}'' \parallel \bar{s}''$, the proof follows by taking $\bar{t}'$ for $\bar{t}''$, and $\bar{s}'$ for $\bar{s}''$.
- If $\bar{s} \parallel \bar{t} \xrightarrow{C} \bar{t} \parallel \bar{s}$ and $\bar{t} \parallel \bar{s} \xrightarrow{a}_{\mathcal{M}_\mathcal{P} \parallel \mathcal{M}_\Pi} \bar{t}' \parallel \bar{s}'$ then $\bar{s} \parallel \bar{t} \xrightarrow{a}_{\mathcal{M}_\Pi \parallel \mathcal{M}_\mathcal{P}} \bar{s}'' \parallel \bar{t}'' \wedge \bar{t}' \parallel \bar{s}' \xrightarrow{C} \bar{s}'' \parallel \bar{t}''$. This case is symmetric to the previous case.

$\square$

Commutativity of composition of LTS networks. To avoid the issues discussed in Example 6.1 an alternative definition of LTS network can be designed. Both process vectors and synchronisation vectors may be replaced by indexed families. An *indexed family* consists of a sets of objects (the process LTSs or synchronisation vectors), an index set, and a surjective function mapping elements from the index set to elements of the set of objects. When the index sets of two networks are disjoint, then the union of sets can be applied, where we previously would use vector concatenation, to compose the collections of process LTSs and synchronisation laws. The union of two indexed families is commutative, as such, commutativity of composition of LTS networks with indexed families is also commutative.

7. Congruence for LTS networks

In this section we prove that DPBB is a congruence for LTS networks as defined in Definition 7.1 [20].

Definition 7.1 (Congruence for LTS networks). An LTS equivalence R is a *congruence for LTS networks* if and only if for all networks $(\Pi, \mathcal{V})$ of some size n , and all vectors of LTSs P (also of size n) it holds that

$$(\forall i \in 1..n. \Pi_i R P_i) \implies \mathcal{G}_{(\Pi, \mathcal{V})} R \mathcal{G}_{(P, \mathcal{V})}$$

That DPBB is a congruence for LTS networks follows from associativity and commutativity rules presented in Section 6 for LTS networks in which the synchronisation laws implement synchronisation on the common alphabet of the LTSs.

However, it is unnecessary to require that the set of synchronisation laws implements synchronisation on the common alphabet. As this requirement excludes many LTS networks in practice, we discuss next an alternative proof, presented for Proposition 7.1, that does not require synchronisation on a common alphabet.

Proposition 7.1. *Consider two vectors of LTSs Π and P , and a set of synchronisation laws $\mathcal{V}$. Furthermore, assume that τ -transitions are not renamed, cut, or synchronised. It holds that*

$$(\forall i \in 1..n. \Pi_i \xleftrightarrow{b} P_i) \implies \mathcal{G}_{(\Pi, \mathcal{V})} \xleftrightarrow{b} \mathcal{G}_{(P, \mathcal{V})}$$

Proof. Given two vectors of Labelled Transitions Systems (LTSs) Π and P such that for all $i \in 1..n$ there is a DPBB relation B_i with $\Pi_i \xleftrightarrow{b} P_i$, we define the bisimulation relation C as follows:

$$C = \{(\bar{s}, \bar{t}) \mid \bar{s} \in \mathcal{S}_{(\Pi, \mathcal{V})} \wedge \bar{t} \in \mathcal{S}_{(P, \mathcal{V})} \wedge \forall i \in 1..n. \bar{s}_i B_i \bar{t}_i\}$$

We prove that C is a DPBB relation as defined in Definition 3.5. We will use $Ac(\bar{v}) = \{i \mid i \in 1..n \wedge \bar{v}_i \neq \bullet\}$ as a shorthand for the set of indices of processes participating in a synchronisation law $(\bar{v}, a)$; e.g., $Ac(\langle c, b, \bullet \rangle) = \{1, 2\}$.

- C relates the initial states of $\mathcal{M}_\Pi$ and $\mathcal{M}_P$. Consider states $\bar{s} \in \mathcal{I}_{(\Pi, \nu)}$. For each $i \in 1..n$ there is a state $q_i \in \mathcal{I}_{P_i}$ such that $\bar{s}_i B_i t_i$. Let $\bar{t}$ be the state build from these q_i such that $\bar{t}_i = q_i$ for all $i \in 1..n$. Then, $\bar{t} \in \mathcal{I}_{(P, \nu)}$ and $\bar{s} C \bar{t}$. The symmetric case follows similarly.
- If $\bar{s} C \bar{t}$ and $\bar{s} \xrightarrow{a}_{(\Pi, \nu)} \bar{s}'$ then either $a = \tau \wedge \bar{s}' C \bar{t}$, or $\bar{t} \xrightarrow{\tau^*}_{(P, \nu)} \hat{\bar{t}} \xrightarrow{a}_{(P, \nu)} \bar{t}' \wedge \bar{s} C \hat{\bar{t}} \wedge \bar{s}' C \bar{t}'$. Consider a law $(\bar{v}, a) \in \mathcal{V}$ enabling transition $\bar{s} \xrightarrow{a}_{(\Pi, \nu)} \bar{s}'$. We distinguish two cases:

1. There is a τ -action in synchronisation vector $\bar{v}$, i.e., $\exists i \in 1..n. \bar{v}_i = \tau$. Therefore, there is a transition $\bar{s}_i \xrightarrow{\tau}_i \bar{s}'_i$. Since τ -transitions do not synchronise it follows that it is the only action in the synchronisation vector, i.e., $\{i\} = Ac(\bar{v})$. Hence, $a = \tau$ as renaming τ -transitions is not allowed. Furthermore, by Definition 3.3, for all $j \in 1..n \setminus \{i\}$ it holds that $\bar{s}_j = \bar{s}'_j$. As we also have $\bar{s}_j B_j \bar{t}_j$, it follows that $\bar{s}'_j B_j \bar{t}_j$.

Because $\bar{s}_i B_i \bar{t}_i$ and $\bar{s}_i \xrightarrow{\tau}_i \bar{s}'_i$, by Definition 3.5, two cases can occur:

- * $a = \tau$ with $\bar{s}'_i B_i \bar{t}_i$. Hence, for all $j \in 1..n$ we have $\bar{s}_j B_j \bar{t}_j$. By definition of C , it follows that $\bar{s}' C \bar{t}$.
- * $\bar{t}_i \xrightarrow{\tau^*}_i \hat{\bar{t}} \xrightarrow{a}_i \bar{t}'$ with $\bar{s}_i B_i \hat{\bar{t}}$ and $\bar{s}'_i B_i \bar{t}'$. Since no τ -transitions are cut, there also exists a path $\bar{t} \xrightarrow{\tau^*}_{(P, \nu)} \hat{\bar{t}} \xrightarrow{a}_{(P, \nu)} \bar{t}'$ with $\hat{\bar{t}}_i = \hat{\bar{t}}$, $\bar{t}'_i = \bar{t}'$, and for all $j \in 1..n \setminus \{i\}$ we have $\bar{t}_j = \hat{\bar{t}}_j = \bar{t}'_j$. Therefore, from $\bar{s}_i B_i \hat{\bar{t}}$, $\bar{s}'_i B_i \bar{t}'$, and $\forall j \in 1..n \setminus \{i\}. \bar{s}_j B_j \bar{t}_j$ we deduce that $\bar{s} C \hat{\bar{t}}$ and $\bar{s}' C \bar{t}'$.

2. There is *no* τ -action in synchronisation vector $\bar{v}$, i.e., $\forall i \in 1..n. \bar{v}_i \neq \tau$. By Definition 3.3, for all $j \in 1..n \setminus Ac(\bar{v})$ we have $\bar{s}'_j = \bar{s}_j$. Thus, since $\bar{s}_j B_j \bar{t}_j$ it follows that $\bar{s}'_j B_j \bar{t}_j$. Furthermore, we have for all $j \in Ac(\bar{v})$ a transition $\bar{s}_i \xrightarrow{\bar{v}_i}_i \bar{s}'_i$. Hence, as $\bar{v}_j \neq \tau$ for all those $j \in Ac(\bar{v})$, there exists a path $\bar{t}_j \xrightarrow{\tau^*}_j \hat{\bar{t}}_j \xrightarrow{\bar{v}_j}_j \bar{t}'_j$ with $\bar{s}_j B_j \hat{\bar{t}}_j$ and $\bar{s}'_j B_j \bar{t}'_j$ (by Definition 3.5). From Definition 3.3 it follows that there also is a path $\bar{t} \xrightarrow{\tau^*}_{(P, \nu)} \hat{\bar{t}} \xrightarrow{a}_{(P, \nu)} \bar{t}'$ where for all $j \in 1..n \setminus Ac(\bar{v})$ $\hat{\bar{t}}_j$ and $\bar{t}'_j$ are defined by $\bar{t}'_j = \hat{\bar{t}}_j = \bar{t}_j$. Hence, from $\forall i \in 1..n. \bar{s}_i B_i \bar{t}_i$ and $\forall j \in Ac(\bar{v}). \bar{s}_j B_j \hat{\bar{t}}_j$, and $\forall k \in Ac(\bar{v}). \bar{s}'_k B_k \bar{t}'_k$ we deduce that $\bar{s} C \hat{\bar{t}}$ and $\bar{s}' C \bar{t}'$.

- If $\bar{s} C \bar{t}$ and $\bar{t} \xrightarrow{a}_{(P, \nu)} \bar{t}'$ then either $a = \tau \wedge \bar{s}' C \bar{t}$, or $\bar{s} \xrightarrow{\tau^*}_{(\Pi, \nu)} \hat{\bar{s}} \xrightarrow{a}_{(\Pi, \nu)} \bar{s}' \wedge \bar{s} C \hat{\bar{t}} \wedge \bar{s}' C \bar{t}'$. This case is symmetric to the previous case.
- If $\bar{s} C \bar{t}$ and there is an infinite sequence of states $(\bar{s}^k)_{k \in \omega}$ such that $\bar{s} = \bar{s}^0$, $\bar{s}^k \xrightarrow{\tau}_{(\Pi, \nu)} \bar{s}^{k+1}$ and $\bar{s}^k C \bar{t}$ for all $k \in \omega$, then there exists a state $\bar{t}'$ such that $\bar{t} \xrightarrow{\tau^+}_{(P, \nu)} \bar{t}'$ and $\bar{s}^k C \bar{t}'$ for some $k \in \omega$. For all $k \in \omega$ let $\bar{v}^k$ be the synchronisation law enabling transition $\bar{s}^k \xrightarrow{\tau} \bar{s}^{k+1}$.

We distinguish two cases:

- * There is a $k \in \omega$ such that $\bar{s}^k \xrightarrow{\tau} \bar{s}^{k+1}$ is the result of the synchronisation of multiple processes in Π , i.e., $\exists k \in \omega, i \in 1..n. \{i\} \subset Ac(\bar{v}^k)$. In the τ -sequence $\bar{s}^\ell C \bar{t}$ for all $\ell \in \omega$, hence, we have $\bar{s}^k C \bar{t}$. Furthermore, since C is a

bisimulation relation it follows that there are states $\hat{t}, \bar{t}' \in \mathcal{S}_{(P, \nu)}$ with a τ -path $\bar{t} \xrightarrow{(\tau, \nu)}^* \hat{t} \xrightarrow{(\tau, \nu)} \bar{t}'$ such that $\bar{s}^{k+1} B \bar{t}'$. Thus, $\bar{t} \xrightarrow{(\tau, \nu)}^+ \bar{t}'$ and for $k+1 \in \omega$ it holds that $\bar{s}^{k+1} C \bar{t}'$ completing the case.

- * The τ -sequence only consists of τ -transitions performed independently by the processes in Π , i.e., $\forall k \in \omega. \forall i \in 1..n. \{i\} \not\subseteq \text{Ac}(\bar{v}^k)$. Since all of the τ -transitions are performed independently, there has to be at least one process of which its infinite τ -sequence is embedded in the global infinite τ -sequence, otherwise the given τ -sequence starting from s would not be infinite. Suppose the i^{th} process has such a sequence, then the sequence starts from state $\bar{s}_i$. The independent infinite sequence is embedded in that of the network, hence, for all $k \in \omega$ it holds that $\bar{s}_i^k B_i \bar{t}_i$. Since $\bar{s}_i B_i \bar{t}_i$, by Definition 3.5, there is a state $t' \in \mathcal{S}_{P_i}$ with $\bar{t}_i \xrightarrow{(\tau, \nu)}^+ t'$ and some $\ell \in \omega$ such that $\bar{s}_i^\ell B_i t'$. We construct state $\bar{t}'$ such that for all $j \in 1..n$ if $j = i$, then $\bar{t}'_j = t'$, and otherwise $\bar{t}'_j = \bar{t}_j$. As local τ -transitions are not cut nor renamed, it follows that $\bar{t} \xrightarrow{(\tau, \nu)}^+ \bar{t}'$. Moreover, since $\bar{s}^k C \bar{t}$ for all $k \in \omega$, by definition of C , we have $\bar{s}_j^\ell B_j \bar{t}_j$ for all $j \in 1..n$. Finally, because $\bar{s}_i^\ell B_i \bar{t}_i$ and for all $j \in 1..n \setminus \{i\}$ it holds that $\bar{s}_j^\ell B_j \bar{t}_j$, by construction of $\bar{t}'$ it follows that $\bar{s}^\ell C \bar{t}'$.

- If $\bar{s} C \bar{t}$ and there is an infinite sequence of states $(\bar{t}^k)_{k \in \omega}$ such that $\bar{t} = \bar{t}^0$, $\bar{t}^k \xrightarrow{(\tau, \nu)} \bar{t}^{k+1}$ and $\bar{s} C \bar{t}^k$ for all $k \in \omega$, then there exists a state $\bar{s}'$ such that $\bar{s} \xrightarrow{(\tau, \nu)}^+ \bar{s}'$ and $\bar{s}' C \bar{t}^k$ for some $k \in \omega$. This case is symmetric to the previous case.

□

8. Application

In order to compare compositional approaches with the classical, non-compositional approach, we have employed CADP to minimise a set of test cases modulo DPBB.

Each test case consists of a model that is minimised with respect to a given liveness property. To achieve the best minimisation we applied maximal hiding [21] in all approaches. Intuitively, maximal hiding hides all actions except for the interface actions and actions relevant for the given liveness property. In general, one can also hide fewer actions, but this will decrease the impact of DPBB reduction, both in compositional and non-compositional model checking.

As *composition strategy* we have used the *smart reduction* approach described in [23]. In CADP, the *classical approach*, where the full state space is constructed at once and no intermediate minimisations are applied, is the *root reduction* strategy.

We have measured the *running time* and the *maximum number of states and transitions* generated by the two methods.

Experimental setup. To facilitate replication we briefly discuss the methods used for our experiments.

For compositional approaches, the running time and largest state space considered depend heavily on the composition order, i.e., the order in which the components are combined. The smart reduction approach uses a heuristic to determine the order in which

to compose processes. In [23], it has been experimentally established that this heuristic frequently works very well. After each composition step the result is minimised.

We use the following expression from the scripting language SVL of CADP to invoke the smart reduction modulo DPBB approach:

```
smart total divbranching reduction of (<m>)
```

where <m> is the test model.

In the classical approach the state space of the entire system is generated before minimisation is applied. This approach is invoked as follows

```
root total divbranching reduction of (<m>)
```

where <m> is the test model.

The experiments were run on the DAS-5 cluster [33] machines. They have an INTEL HASWELL E5-2630-v3 2.4 GHz CPU, 64 GB memory, and run CENTOS LINUX 7.2. The running time of the two approaches was measured as the wall clock time (i.e., the real elapsed time) using the Unix time command:

```
/usr/bin/time -f "%e" svl <file>
```

The argument `-f "%e"` specifies that the time written as output should follow format `"%e"` where `%e` indicates the wall clock time. The `svl <file>` argument invokes the SVL-engine with script file `<file>`. The command measures the wall clock time of the execution of the SVL-script.

The maximum number of states and transitions that were generated were extracted from the SVL-log files after execution of the script.

The set of test cases. To prevent source bias case studies were selected from four different source. In total 19 case studies were selected: four MCRL2 [17] models distributed with its tool set, nine CADP models, three from the BEEM database [34], and three from Example Repository for Finite State Verification Tools [35].

The models stemming from the MCRL2 tool set distribution are the following:

1. The *1394* model, created by Luttik [36], specifies the 1394 or firewire protocol. *Property:* every PAreq with parameter immediate is eventually followed by a matching PAcon with parameter won.
2. The *1394'* model is the 1394 model scaled up with extra internal transitions. This model is our own adaptation of the 1394 model and is therefore not distributed with the tool set. *Property:* same as the 1394 model.
3. The *ACS* model describes the ACS Manager that is part of the ALMA project of the European Southern Observatory. The ACS Manager is part of a system controlling a large collection of radio telescopes. The model consists of a manager and some containers and components and was created by Ploeger [37]. *Property:* every time container MT1 is locked, eventually it is freed again.
4. *Wafer Stepper* models a wafer stepper used in the manufacturing of integrated circuits. *Property:* always, eventually, all wafers in the system will be exposed.

The CADP models consist of:

1. *Cache* models a directory-based cache coherency protocol for a multi-processor architecture. The model was developed by Kahlouche et al.[38]. *Property*: there is no live-lock.
2. The *DES* model describes an implementation of the data encryption standard, which allows to cipher and decipher 64-bit vectors using a 64-bit key vector [39]. *Property*: the DES can always deliver outputs.
3. *HAVi-LE* describes the asynchronous Leader Election protocol used in the HAVi (Home Audio-Video) standard, involving three device control managers. The model is fully described by Romijn [40]. *Property*: always eventually a leader is selected.
4. *HAVi-LE'* is an adaptation of the HAVi-LE model containing transitions denoting logging events. Since the model is our own adaptation it is not distributed with CADP. *Property*: same as the HAVi-LE model.
5. *Le Lann* models a distributed leader election algorithm for unidirectional ring networks. The CADP model was developed by Garavel and Mounier [41]. *Property*: process P0 is infinitely many times in the critical section.
6. *ODP* is a model of an open distributed processing trader [42]. *Property*: work is always executed eventually.
7. The *Erat. Sieve* model computes prime numbers implementing a distributed Eratosthenes Sieve; the model describes a pipeline of units, of which each unit blocks input numbers that are multiples of a given number. The model consists of four units. *Property*: if the number two is generated, then it is eventually reported as a prime number.
8. *Erat. Sieve'* is a variant of Erat. Sieve consisting of seven units. *Property*: same as the Erat. Sieve model.
9. The *Transit* model describes a transit-node, it models an abstraction of a routing component of a communication network. The model was developed by Mounier [43]. *Property*: every time a message is receive, it is eventually either sent out the node or buffered as faulty.

The BEEM models are:

1. The *Peterson* model describes Peterson's mutual exclusion algorithm [44] for seven processes. *Property*: every time process P0 waits for access to the critical section, it will eventually enter it.
2. *Anderson* models Anderson's queue lock mutual exclusion algorithm [45] for three processes. *Property*: Every time a process waits for access to the critical section, it will eventually enter it.
3. *Anderson'* is a variant of the Anderson model considering four processes competing for a lock. *Property*: same as the Anderson model.

Table 1. Experiments: smart reduction vs. root reduction

Test case	Running time (sec.)		Max. #states		Max. #transitions		Reduced #states	Reduced #transitions
	smart	root	smart	root	smart	root		
1394	14.41	8.25	102,983	198,692	187,714	355,338	1	1
1394'	47.51	460.53	2,832,074	36,855,184	5,578,078	96,553,318	1	1
ACS	70.87	11.22	1,854	4,764	4,760	14,760	29	61
Anderson	26.56	15.42	153,664	384,104	2,118,368	5,892,964	1	1
Anderson'	373.56	1852.42	15,116,544	56,250,000	268,738,560	1,188,000,000	1	1
Cache	20.55	7.84	616	616	4,631	4,631	1	1
Chiron	22.76	13.66	317,115	481,140	2,563,650	3,456,675	216	1286
Chiron'	1,171.06	1,236.06	49,076,280	56,293,380	467,536,860	513,857,520	2376	17,974
DES	54.61	948.66	1,404	64,498,297	3,510	518,438,860	1	1
Erat. Sieve	63.00	8.43	1,156,781	234	2,891,692	406	3	2
Erat. Sieve'	—	10.64	—	865	—	2,012	3	2
Gas Station	325.10	362.31	11,042,816	11,436,032	84,254,720	87,105,536	432	5616
HAVi-LE	114.27	493.01	970,772	15,688,570	5,803,552	80,686,289	131,873	644,695
HAVi-LE'	93.08	5,255.56	453,124	190,208,728	2,534,371	876,008,628	159,318	849,227
Le Lann	96.35	5,599.15	12,083	160,025,986	701,916	944,322,648	83,502	501,573
ODP	32.90	9.97	10,397	91,394	87,936	641,226	432	2,268
Peterson	63.04	—	9	—	139	—	9	22
Transit	25.50	59.69	22,928	3,763,192	132,712	39,925,524	636	3,188
Wafer Stepper	74.18	57.54	962,122	3,772,753	4,537,240	16,977,692	905,955	4,095,389

From the Example Repository for Finite State Verification Tools we selected the following models:

1. The *Chiron* model describes a user interface development system with two clients. The system consists of the Chiron server, managing generic aspects of a user interface, and artists (the clients). This server is responsible for notifying artists when a user interface event occurs, while the clients listen for notifications from the server. The formal model was developed by Avrunin et al. [46]. *Property*: If an artist is registered for event e_1 , then it will eventually be notified for this event.
2. *Chiron'* is a adapted version of the Chiron model where another client is added. There are a total of three clients. *Property*: same as the Chiron model.
3. The *Gas Station* problem [47] simulates a self-serve gas station. The gas station consists of two pumps, an operator, and three customers. *Property*: A charge is made eventually after a customer has started pumping.

Measurement Results. The results of our experiments are shown in Table 1. The *Test case* column indicates the test case model corresponding to the measurements.

The *smart* and *root* sub-columns denote the measurement for the smart reduction and root reduction approaches, respectively.

In the *Running time (sec.)* column the running time until completion of the experiment is shown in seconds. Indicated in bold are the shortest running times comparing the *smart* and *root* sub-columns. The maximum running time of an experiment was set to 80 hours, after which the experiment was discontinued (indicated with —).

The columns *Max. #states* and *Max. #transitions* show the largest number of states and transitions, respectively, generated during the experiment. Of both methods the best result is indicated in bold.

The number of states and transitions after minimisation are shown in the *Reduced #states* and *Reduced #transitions* columns, respectively.

Discussion. In terms of running time smart reduction performs best for ten out of nineteen models, whereas root reduction performs best in eight of the models. For *Wafer stepper* and *Transit* smart reduction is only a few seconds to a minute faster than root reduction. The gain is a few minutes for *1394*, *DES*, and *HAVi-LE*. For *HAVi-LE*, *Le Lann*, and *Peterson* smart reduction is several hours faster. In general, the smart reduction approach performs better for large models where the state space can be reduced significantly before composition.

Root reduction performs best in relatively small models; *1394*, *ACS*, *Cache*, *Lamport*, and *ODP*. However, for these cases the difference in running times is negligible. Only the *Erat. Sieve* model is minimised significantly faster by root reduction. For smaller models the overhead of the smart reduction heuristic is too high to obtain any benefits from the nominated ordering.

In terms of state space, smart reduction is the clear winner in all cases, with the exception of the *Erat. Sieve* models. Indeed, smart reduction performs particularly badly for the *Erat. Sieve* model. The model consists of a pipeline where data is being pushed from one end to another. While the data domain considered by the nodes in the pipeline consists of 32 elements, in the minimised state space only one element remains. As synchronising actions may not be hidden in the local process LTSs the incremental composition and minimisation leads to a state space that is several orders of magnitude larger than the final state space.

The efficiency of smart reduction seems to increase when it is more successful in keeping compositions small; e.g., *Transit* and *Wafer stepper* have a similar number of states before reduction, however, when minimizing the *Transit* model the smart reduction approach is able to reduce the number of states much more significantly than in the *Wafer stepper* model (see the *Max. #states smart* column).

Lessons learned. In summary, the following lessons can be learned from this experiment:

- The overhead of applying its heuristics makes smart reduction less efficient in terms of running time when applied on small models.
- In general, smart reduction produces significantly smaller state spaces, especially for larger models, compared to other approaches.
- The data domain can have a significant impact on the effectiveness of smart reduction. In particular, this is the case when data can only be eliminated at a late stage of the compositional minimisation (as demonstrated by *Erat. Sieve*).
- The efficiency of smart reduction seems to increase when it is more successful in keeping compositions small.

Threads to validity. The following threads to validity must be considered when interpreting the results:

- Only one tool has been involved to conduct the experiment, hence the results may be implementation specific. This only affects measured running times, as the state spaces produced during compositional minimisation is deterministic with respect to the composition order.

- A more controlled experiment needs to be considered in order to extrapolate the results of this experiment. In particular, the effect of action hiding, number of processes, and chosen composition order should be controlled to determine correlation of these aspects with running time and size of the state space.
- The study only considers the DPBB equivalence as minimisation relation. Results may vary depending on the chosen equivalence relation. The DPBB equivalence is the strongest equivalence relation offered by CADP that still allows abstraction. Thus, the expectation is that other relations show equal or better performance improvements.
- The number and variety of case studies is still too limited for generalizing claims on the effectiveness of smart reduction in comparison with root reduction in the field.

9. Conclusions

In this article we have shown that DPBB is a congruence for parallel composition of LTS networks where there is synchronisation on given label combinations. Therefore, DPBB may be used to reduce components in the compositional verification of LTS networks. It had already been shown that compositional verification of LTS networks is adequate for safety properties, as this follows from the fact that branching bisimilarity is a congruence for the parallel composition of synchronising LTS networks [19]. As DPBB preserves both safety and liveness properties, compositional verification can be used to verify liveness properties as well.

Furthermore, we have discussed how to safely decompose an LTS network in the case where verification has to start from the system as a whole. Both the composition and consistent decomposition of LTS networks preserve the admissibility property of LTS networks. Hence, the composition operator remains compatible with the compositional verification approaches for LTS networks described by [20].

We have shown that parallel composition of LTS networks with synchronisation on the common result action alphabet is associative and commutative. From this it follows that DPBB is also a congruence for LTS networks as defined by Garavel, Lang, and Mateescu [20] if the set of synchronisation laws implements synchronisation on the common alphabet. We have shown, however, that the requirement to synchronise on the common alphabet is unnecessarily restrictive. This has been shown in a direct proof of DPBB being a congruence for LTS networks.

All proofs in this work, except for the one in Section 7, have been mechanically verified using the Coq proof assistant ⁴ and are available online. ⁵

Although our work focuses on the composition of LTS networks, the results are also applicable on composition of individual LTSs. Our parallel composition operator subsumes the usual parallel composition operators of standard process algebra languages such as CCS [14], CSP [48], MCRL2 [17], and LOTOS [18].

⁴<https://coq.inria.fr>.

⁵http://www.win.tue.nl/mdse/composition/DPBB_is_a_congruence_for_synchronizing_LTSs.zip.

Finally, we have run a set of experiments to compare compositional and traditional DPBB reduction. The compositional approach applies CADP’s smart reduction employing a heuristic to determine an efficient compositional reduction order. The traditional reduction generates the complete state space before applying reduction. The compositional approach performed better in the medium to large models where the intermediate state space can be kept small.

Future Work. This work has been inspired by an approach for the compositional verification of transformations of LTS networks [49, 32, 50, 51, 52]. We would like to apply the results of this article to the improved transformation verification algorithm [49], thus guaranteeing its correctness for the compositional verification of transformations of LTS networks.

In future experiments, we would like to involve recent advancements in the computation of branching bisimulation, and therefore also DPBB, both sequentially [53, 54] and in parallel on graphics processors [55], and we may consider other equivalence relations, such as simulation equivalence [56]. It will be interesting to measure the effect of applying these algorithms to compositionally solve a model checking problem.

Finally, we can consider various extensions of the LTS network formalism. For instance, by encoding timing in the LTSs, it is possible to reason about timed system behaviour. Combining approaches such as [57, 58] with our results would allow to compositionally reason about timed behaviour. Other possibilities are to distinguish *must* and *may* transitions, and explicitly involve data variables, for instance as is suggested in [59]. We plan to investigate this further.

References

- [1] W. Fokkink, R. J. Glabbeek, S. P. Luttik, Divide and congruence III: Stability & divergence, in: CONCUR, Vol. 85 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017, pp. 15:1–15:16.
- [2] S. de Putter, A. J. Wijs, Compositional model checking is lively, in: FACS, Vol. 10487 of LNCS, Springer, 2017, pp. 117–136.
- [3] E. M. Clarke, O. Grumberg, D. Peled, Model Checking, The MIT Press, 1999.
- [4] C. Baier, J.-P. Katoen, Principles of model checking, MIT Press, 2008.
- [5] D. Peled, Ten Years of Partial Order Reduction, in: CAV, Vol. 1427, Springer, Berlin, Heidelberg, 1998, pp. 17–28.
- [6] E. M. Clarke, E. A. Emerson, S. Jha, A. P. Sistla, Symmetry reductions in model checking, in: CAV, Springer, Berlin, Heidelberg, 1998, pp. 147–158.
- [7] E. M. Clarke, D. E. Long, K. L. McMillan, Compositional model checking, in: LICS, IEEE Computer Society Press, 1989, pp. 353–362.
- [8] H. R. Andersen, Partial Model Checking, in: LICS, IEEE Computer Society Press, 1995, pp. 398–407.
- [9] H. R. Andersen, Partial Model Checking of Modal Equations: A Survey, STTT 2 (3) (1999) 242–259.
- [10] D. Kozen, Results on the Propositional μ -calculus, Theoretical Computer Science 27 (1983) 333–354.
- [11] R. J. van Glabbeek, W. P. Weijland, Branching Time and Abstraction in Bisimulation Semantics, Journal of the ACM 43 (3) (1996) 555–600.
- [12] R. J. van Glabbeek, S. P. Luttik, N. Trčka, Branching bisimilarity with explicit divergence, Fundam. Inf. 93 (4) (2009) 371–392.
- [13] R. J. van Glabbeek, S. P. Luttik, N. Trčka, Computation Tree Logic with Deadlock Detection, LMCS 5 (4) (2009) 1–24.
- [14] R. Milner, Communication and Concurrency, Prentice-Hall, 1989.
- [15] C. A. R. Hoare, Communicating sequential processes, Prentice Hall, 1985.
- [16] J. A. Bergstra, J. W. Klop, Process algebra for synchronous communication, Information and Control 60 (1/3) (1984) 109–137.

- [17] S. Cranen, J. F. Groote, J. J. A. Keiren, F. P. M. Stappers, E. P. de Vink, W. Wesselink, T. Willemse, An Overview of the mCRL2 Toolset and Its Recent Advances, in: TACAS, Vol. 7795 of LNCS, Springer, 2013, pp. 199–213.
- [18] ISO/IEC, LOTOS — A Formal Description Technique Based on the Temporal Ordering of Observational Behaviour. International Standard 8807, International Organization for Standardization — Information Processing Systems — Open Systems Interconnection (1989).
- [19] B. Bloom, Structural operational semantics for weak bisimulations, Theoretical Computer Science 146 (1) (1995) 25 – 68.
- [20] H. Garavel, F. Lang, R. Mateescu, Compositional Verification of Asynchronous Concurrent Systems using CADP, Acta Informatica 52 (4-5) (2015) 337–392.
- [21] R. Mateescu, A. J. Wijs, Property-Dependent Reductions Adequate with Divergence-Sensitive Branching Bisimilarity, Science of Computer Programming 96 (3) (2014) 354–376.
- [22] F. Lang, Refined Interfaces for Compositional Verification, in: FORTE, Vol. 4229 of LNCS, Springer, 2006, pp. 159–174.
- [23] P. Crouzen, F. Lang, Smart Reduction, in: FASE, Vol. 6603 of LNCS, Springer, 2011, pp. 111–126.
- [24] F. Lang, Exp.Open 2.0: A Flexible Tool Integrating Partial Order, Compositional, and On-The-Fly Verification Methods, in: IFM, Vol. 3771 of LNCS, Springer, 2005, pp. 70–88.
- [25] F. Lang, Unpublished PVS proof (by Jaco van de Pol) and textual proof showing that branching bisimulation is a congruence for Networks of LTSs. This proof does not consider DPBB. Personal communication (2016).
- [26] L. Spaninks, An Axiomatisation for Rooted Branching Bisimulation with Explicit Divergence, Master’s thesis, Eindhoven University of Technology (2013).
- [27] J.-P. Krimm, L. Mounier, Compositional state space generation from LOTOS programs, in: TACAS, Vol. 1217, Springer, Berlin, Heidelberg, 1997, pp. 239–258.
- [28] F. Maraninchi, Operational and compositional semantics of synchronous automaton compositions, in: CONCUR, Vol. 630 of LNCS, Springer, 1992, pp. 550–564.
- [29] M. Mazzara, I. Lanese, Towards a unifying theory for web services composition, in: WS-FM’06, Vol. 4184 of LNCS, Springer, 2006, pp. 257–272.
- [30] I. Ulidowski, I. Phillips, Ordered SOS Process Languages for Branching and Eager Bisimulations, Information and Computation 178 (1) (2002) 180–213.
- [31] C. Verhoef, A congruence theorem for structured operational semantics with predicates and negative premises, in: CONCUR, Vol. 836, Springer, Berlin, Heidelberg, 1994, pp. 433–448.
- [32] A. J. Wijs, Define, Verify, Refine: Correct Composition and Transformation of Concurrent System Semantics, in: FACS, Vol. 8348 of LNCS, Springer, 2013, pp. 348–368.
- [33] H. Bal, D. Epema, C. de Laat, R. van Nieuwpoort, J. Romein, F. Seinstra, C. Snoek, H. Wijshoff, A Medium-Scale Distributed System for Computer Science Research: Infrastructure for the Long Term, IEEE Computer 49 (5) (2016) 54–63.
- [34] R. Pelánek, BEEM: Benchmarks for Explicit Model Checkers, in: SPIN’07, Vol. 4595 of LNCS, Springer, 2007, pp. 263–267.
- [35] Laboratory for Advanced Software Engineering Research, Example repository for finite state verification tools, <http://laser.cs.umass.edu/verification-examples/>, last Accessed 20 Dec. 2017 (8 Jan. 2003).
- [36] S. P. Luttik, Description and Formal Specification of the Link Layer of P1394, Tech. Rep. SEN-R9706, CWI (1997).
- [37] B. Ploeger, Analysis of ACS Using mCRL2, Tech. Rep. 09-11, Eindhoven University of Technology (2009).
- [38] H. Kahlouche, C. Viho, M. Zendri, An industrial experiment in automatic generation of executable test suites for a cache coherency protocol, in: Proceedings of the IFIP TC6 11th International Workshop on Testing Communicating Systems, IWTCS, Kluwer, B.V., 1998, pp. 211–226.
- [39] National Institute of Standards and Technology, Data Encryption Standard (DES). Federal Information Processing Standards 46-3 (1999).
- [40] J. Romijn, Model Checking a HAVi Leader Election Protocol, Tech. Rep. SEN-R9915, CWI (1999).
- [41] H. Garavel, L. Mounier, Specification and Verification of various Distributed Leader Election Algorithm for Unidirectional Ring Networks, Research Report RR-2986, INRIA (1996).
- [42] H. Garavel, M. Sighireanu, A Graphical Parallel Composition Operator for Process Algebras, in: FORTE/PSTV’99, Vol. 156 of IFIP Conference Proceedings, Kluwer, 1999, pp. 185–202.
- [43] L. Mounier, A LOTOS specification of a "transit-node", Tech. Rep. SPECTRE 94-8, VERIMAG (3 1994).
- [44] G. L. Peterson, Myths about the mutual exclusion problem, Inf. Process. Lett. 12 (1981) 115–116.

- [45] J. H. Anderson, Y.-J. Kim, T. Herman, Shared-memory mutual exclusion: major research trends since 1986, *Distrib. Comput.* 16 (2-3) (2003) 75–110.
- [46] G. S. Avrunin, J. C. Corbett, M. B. Dwyer, C. S. Pasareanu, S. F. Siegel, Comparing finite-state verification techniques for concurrent software, Technical Report UM-CS-1999-069, University of Massachusetts (1999).
- [47] D. Heimbald, D. Luckham, Debugging Ada tasking programs, *IEEE Software* 2 (2) (1985) 47–57.
- [48] A. Roscoe, *The Theory and Practice of Concurrency*, Prentice-Hall, 1998.
- [49] S. de Putter, A. J. Wijs, Verifying a Verifier: On the Formal Correctness of an LTS Transformation Verification Technique, in: *FASE*, Vol. 9633 of LNCS, Springer, 2016, pp. 383–400.
- [50] A. J. Wijs, Confluence Detection for Transformations of Labelled Transition Systems, in: *Proceedings of the 2nd Graphs as Models Workshop (GaM 2015)*, Vol. 181 of EPTCS, Open Publishing Association, 2015, pp. 1–15.
- [51] A. J. Wijs, L. J. P. Engelen, Efficient Property Preservation Checking of Model Refinements, in: *TACAS*, Vol. 7795 of LNCS, Springer, 2013, pp. 565–579.
- [52] A. J. Wijs, L. J. P. Engelen, *REFINER*: Towards Formal Verification of Model Transformations, in: *NFM*, Vol. 8430 of LNCS, Springer, 2014, pp. 258–263.
- [53] J. F. Groote, A. J. Wijs, An $O(m \log n)$ Algorithm for Stuttering Equivalence and Branching Bisimulation, in: *TACAS*, Vol. 9636 of LNCS, Springer, 2016, pp. 607–624.
- [54] J. F. Groote, D. N. Jansen, J. J. A. Keiren, A. J. Wijs, An $O(m \log n)$ Algorithm for Computing Stuttering Equivalence and Branching Bisimulation, *ACM Transactions on Computational Logic* 18 (2) (2017) 13:1–13:34.
- [55] A. J. Wijs, GPU Accelerated Strong and Branching Bisimilarity Checking, in: *TACAS*, Vol. 9035 of LNCS, Springer, 2015, pp. 368–383.
- [56] G. Cécé, Foundation for a series of efficient simulation algorithms, in: *LICS*, IEEE, 2017, pp. 1–12.
- [57] A. J. Wijs, Achieving Discrete Relative Timing with Untimed Process Algebra, in: *Proc. 12th Conference on Engineering of Complex Computer Systems (ICECCS 2007)*, IEEE Computer Society Press, 2007, pp. 35–44.
- [58] A. J. Wijs, W. J. Fokkink, From χ_t to μCRL : Combining Performance and Functional Analysis, in: *Proc. 10th Conference on Engineering of Complex Computer Systems (ICECCS 2005)*, IEEE Computer Society Press, 2005, pp. 184–193.
- [59] S. S. Bauer, K. G. Larsen, A. Legay, U. Nyman, A. Wąsowski, A model specification theory for components with data, in: *FACS*, Vol. 7253 of LNCS, Springer, 2011, pp. 61–78.