HAL open science

Property-Dependent Reductions for the Modal Mu-Calculus

Radu Mateescu, Anton Wijs

▶ To cite this version:

Radu Mateescu, Anton Wijs. Property-Dependent Reductions for the Modal Mu-Calculus. [Research Report] RR-7690, INRIA. 2011, pp.30. inria-00609585

HAL Id: inria-00609585
https://inria.hal.science/inria-00609585v1

Submitted on 19 Jul 2011

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.

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Property-Dependent Reductions for the Modal Mu-Calculus

Radu Mateescu — Anton Wijs

N° 7690

Juillet 2011

Thème COM

Rapport de recherche

ISSN 0249-6399 ISRN INRIA/RR--7690--FR+ENG

Property-Dependent Reductions for the Modal Mu-Calculus

Radu Mateescu* , Anton Wijs**

Thème COM — Systèmes communicants

Projet Vasy

Rapport de recherche n° 7690 — Juillet 2011 — 30 pages

Abstract: When analyzing the behavior of finite-state concurrent systems by model checking, one way of fighting state explosion is to reduce the model as much as possible whilst preserving the properties under verification. We consider the framework of action-based systems, whose behaviors can be represented by labeled transition systems (LTSs), and whose temporal properties of interest can be formulated in modal μ-calculus (L_μ). First, we determine, for any L_μ formula, the maximal set of actions that can be hidden in the LTS without changing the interpretation of the formula. Then, we define L_μ^{dsbr}, a fragment of L_μ which is compatible with divergence-sensitive branching bisimulation. This enables us to apply the maximal hiding and to reduce the LTS on-the-fly using divergence-sensitive τ-confluence during the verification of any L_μ^{dsbr} formula. The experiments that we performed on various examples of communication protocols and distributed systems show that this reduction approach can significantly improve the performance of on-the-fly verification.

Key-words: divergence-sensitive branching bisimulation, labeled transition system, modal μ-calculus, model checking, on-the-fly verification

A short version of this report is also available as "Property-Dependent Reductions for the Modal Mu-Calculus", *Proceedings of the 18th International SPIN Workshop on Model Checking of Software SPIN'2011, (Cliff Lodge, Snowbird, Utah, USA)*, July 2011.

* Radu.Mateescu@inria.fr

** Technische Universiteit Eindhoven, Postbus 513 5600 MB Eindhoven, The Netherlands.
E-mail: A.J.Wijs@tue.nl

Centre de recherche INRIA Grenoble – Rhône-Alpes
Inovallée, 655, avenue de l'Europe, Montbonnot, 38334 Saint Ismier Cedex (France)
Téléphone : +33 4 76 61 52 00 — Télécopie +33 4 76 61 52 52

Réductions dépendantes de la propriété pour le mu-calcul modal

Résumé : Lorsqu'on analyse le comportement des systèmes concurrents à espace d'états fini par *model checking*, une manière de lutter contre l'explosion d'états est de réduire le modèle le plus possible tout en préservant les propriétés à vérifier. Nous considérons le cadre des systèmes basés sur actions, dont les comportements peuvent être représentés par des systèmes de transitions étiquetées (STEs), et dont les propriétés temporelles d'intérêt peuvent être formulées en μ-calcul modal (L_μ). D'abord, nous déterminons, pour toute formule de L_μ, l'ensemble maximal des actions qui peuvent être masquées dans le STE sans changer l'interprétation de la formule. Ensuite, nous définissons L_μ^{dsbr}, un fragment de L_μ qui est compatible avec la bisimulation de branchement sensible à la divergence. Ceci permet d'appliquer le masquage maximal et de réduire le STE à la volée en utilisant la τ-confluence sensible à la divergence pendant la vérification de la formule L_μ^{dsbr}. Les expériences que nous avons effectuées sur différents exemples de protocoles de communication et de systèmes distribués montrent que cette approche de réduction peut améliorer de manière significative les performances de la vérification à la volée.

Mots-clés : bisimulation de branchement sensible à la divergence, système de transitions étiquetées, μ-calcul modal, vérification basée sur les modèles, vérification à la volée

1 Introduction

Model checking [5] is a technique to systematically verify whether a system specification meets a given temporal property. Although successfully applied in many cases, its usefulness in practice is still hampered by the state explosion phenomenon, which may entail high memory and CPU requirements in order to carry out the verification.

One way to improve the performance of model checking is to check the property at a higher level of abstraction; by abstracting parts of the system behavior away from the specification, its corresponding state space will be smaller, thereby easier to check. This can either be done globally, i.e., before verifying the property, or on-the-fly, i.e., during verification. However, one needs to be careful not to abstract away any details crucial for the outcome of the check, i.e., relevant for the property. This is known as *action abstraction* in action-based formalisms, where state spaces are represented by *Labeled Transition Systems* (LTSs), specifications are written using some flavor of *process algebra* [2], and temporal properties are described using a temporal logic such as the *μ-calculus* (L_μ) [18, 28]. Abstracted behavior is then represented by some predefined action, denoted τ in process algebras. In the past, the main focus in this area has been on devising μ-calculus variants adequate with specific relations, such as μACTL\X [9], which is adequate w.r.t. divergence-sensitive branching bisimulation [15, 14], or weak μ-calculus [28], which is adequate w.r.t. weak bisimulation [25]. For such fragments, the minimization of an LTS modulo the specific relation preserves the truth value of all formulas written in the adequate μ-calculus. Other works focused on devising reductions targeted to specific formulas, such as those written in the *selective μ-calculus* [3]. For each selective μ-calculus formula, it is possible to hide all actions not occurring in the formula, and subsequently minimize the LTS modulo $\tau^*.a$ bisimulation [10] before verifying the formula.

In this report, we propose two enhancements with respect to existing work. Firstly, starting from an arbitrary L_μ formula, we determine automatically the maximal set of actions which can be hidden in an LTS without affecting the outcome of the verification of that formula. This yields the maximum potential for reduction, and therefore for improving the performance of model checking. After hiding, the LTS can be minimized modulo strong bisimulation without disturbing the truth value of the formula. This method is not intrusive, in the sense that it does not force the user to write formulas in a certain way. Secondly, we identify a fragment of L_μ, called L_μ^{dsbr}, which is compatible with divergence-sensitive branching bisimulation. We show that this fragment subsumes μACTL\X, the modalities of selective μ-calculus, and the weak μ-calculus. Compared to these μ-calculi, which require that action formulas contain only names of visible actions, our L_μ^{dsbr} fragment also allows the presence of the invisible action τ, therefore providing additional flexibility in the specification of properties.

The reduction approach for L_μ^{dsbr} is now supported within the CADP[1] verification toolbox [13]. The model checking of a L_μ^{dsbr} formula can be optimized generally in two ways: globally, by generating the LTS, then hiding the maximal set of actions according to the formula, and minimizing the LTS modulo strong or divergence-sensitive branching bisimulation before checking the formula; and on-the-fly, by applying maximal hiding and reduction modulo divergence-

[1] `http://www.inrialpes.fr/vasy/cadp`

sensitive τ-confluence simultaneously with the verification. The experiments we carried out on several examples of protocols and distributed systems show that these optimizations can lead to significant performance improvements.

Section 2 defines the formalisms and relations considered in this report. Section 3 studies the maximal hiding of actions in an LTS w.r.t. a given L_μ formula. Section 4 introduces the L_μ^{dsbr} fragment, shows its compatibility with divergence-sensitive branching bisimulation, and compares its expressiveness with other logics. Section 5 describes and illustrates experimentally the model checking optimizations obtained by applying maximal hiding and reductions for L_μ^{dsbr} formulas. Section 6 gives concluding remarks and directions of future work.

2 Background

Labeled transition system. We consider as interpretation model the classical LTS, which underlies process algebras and related action-based description languages. An LTS is a tuple $\langle S, A, T, s_0 \rangle$, where S is the set of states, A is the set of actions (including the invisible action τ), $T \subseteq S \times A \times S$ is the transition relation, and $s_0 \in S$ is the initial state. The visible actions in $A \setminus \{\tau\}$ are noted a and the actions in A are noted b. A transition $\langle s_1, b, s_2 \rangle \in T$ (also noted $s_1 \xrightarrow{b} s_2$) means that the system can move from state s_1 to state s_2 by performing action b. The reflexive transitive closure of $\xrightarrow{\tau}$ is denoted by $\Longrightarrow$. A finite path is denoted by $s_0 \xrightarrow{b_0 \cdots b_{k-1}} s_k$, which is a finite sequence $s_0, s_1, \ldots, s_k$, such that there exist actions $b_0, \ldots, b_{k-1}$ with $\forall 0 \leq i < k . s_i \xrightarrow{b_i} s_{i+1}$. We assume below the existence of an LTS $M = \langle S, A, T, s_0 \rangle$ on which temporal formulas will be interpreted.

Modal μ-calculus. The variant of L_μ that we consider here consists of action formulas (noted α) and state formulas (noted φ), which characterize subsets of LTS actions and states, respectively. The syntax and semantics of these formulas are defined in Figure 1. Action formulas are built over the set of actions by using Boolean connectors in a way similar to ACTL (Action-based CTL) [26], which is a slight extension w.r.t. the original definition of L_μ [18]. Derived action operators can be defined as usual: $\mathsf{true} = \neg\mathsf{false}$, $\alpha_1 \wedge \alpha_2 = \neg(\neg\alpha_1 \vee \neg\alpha_2)$, etc. State formulas are built from Boolean connectors, the possibility modality ($\langle\ \rangle$), and the minimal fixed point operator (μ) defined over propositional variables X belonging to a set $\mathcal{X}$. Derived state operators can be defined as usual: $\mathsf{true} = \neg\mathsf{false}$, $\varphi_1 \wedge \varphi_2 = \neg(\neg\varphi_1 \vee \neg\varphi_2)$, $[\alpha]\,\varphi = \neg\,\langle\alpha\rangle\,\neg\varphi$ is the necessity modality, and $\nu X.\varphi = \neg\mu X.\neg\varphi[\neg X/X]$ is the maximal fixed point operator ($\varphi[\neg X/X]$ stands for φ in which all free occurrences of X have been negated).

The interpretation $[\![\alpha]\!]_A$ of an action formula on the set of actions of an LTS denotes the subset of actions satisfying α. An action b satisfies a formula α (also noted $b \models_A \alpha$) if and only if $b \in [\![\alpha]\!]_A$. A transition $s_1 \xrightarrow{b} s_2$ such that $b \models_A \alpha$ is called an α-transition. A propositional context $\rho : \mathcal{X} \to 2^S$ is a partial function mapping propositional variables to subsets of states. The notation $\rho \oslash [U/X]$ stands for a propositional context identical to ρ except for variable X, which is mapped to the state subset U. The interpretation $[\![\varphi]\!]_M\,\rho$ of a state formula on an LTS M and a propositional context ρ (which assigns a set of states

Action formulas:

$$
\begin{array}{lll}
\alpha ::= b & & [\![b]\!]_A = \{b\} \\
\quad | \ \mathsf{false} & & [\![\mathsf{false}]\!]_A = \emptyset \\
\quad | \ \neg\alpha_1 & & [\![\neg\alpha_1]\!]_A = A \setminus [\![\alpha_1]\!]_A \\
\quad | \ \alpha_1 \vee \alpha_2 & & [\![\alpha_1 \vee \alpha_2]\!]_A = [\![\alpha_1]\!]_A \cup [\![\alpha_2]\!]_A
\end{array}
$$

State formulas:

$$
\begin{array}{lll}
\varphi ::= \mathsf{false} & & [\![\mathsf{false}]\!]_M\, \rho = \emptyset \\
\quad | \ \neg\varphi_1 & & [\![\neg\varphi_1]\!]_M\, \rho = S \setminus [\![\varphi_1]\!]_M\, \rho \\
\quad | \ \varphi_1 \vee \varphi_2 & & [\![\varphi_1 \vee \varphi_2]\!]_M\, \rho = [\![\varphi_1]\!]_M\, \rho \cup [\![\varphi_2]\!]_M\, \rho \\
\quad | \ \langle\alpha\rangle\, \varphi_1 & & [\![\langle\alpha\rangle\, \varphi_1]\!]_M\, \rho = \{s \in S \mid \exists s \xrightarrow{b} s' \in T.b \in [\![\alpha]\!]_A \wedge s' \in [\![\varphi_1]\!]_M\, \rho\} \\
\quad | \ X & & [\![X]\!]_M\, \rho = \rho(X) \\
\quad | \ \mu X.\varphi_1 & & [\![\mu X.\varphi_1]\!]_M\, \rho = \bigcap\{U \subseteq S \mid [\![\varphi_1]\!]_M\, (\rho \oslash [U/X]) \subseteq U\}
\end{array}
$$

Figure 1: Syntax and semantics of L_μ

to each propositional variable occurring free in φ) denotes the subset of states satisfying φ in that context. The Boolean connectors are interpreted as usual in terms of set operations. The possibility modality $\langle\alpha\rangle\, \varphi_1$ (resp. the necessity modality $[\alpha]\, \varphi_1$) denotes the states for which some (resp. all) of their outgoing transitions labeled by actions satisfying α lead to states satisfying φ_1. The minimal fixed point operator $\mu X.\varphi_1$ (resp. the maximal fixed point operator $\nu X.\varphi_1$) denotes the least (resp. greatest) solution of the equation $X = \varphi_1$ interpreted over the complete lattice $\langle 2^S, \emptyset, S, \cap, \cup, \subseteq\rangle$. A state s satisfies a closed formula φ (also noted $s \models_M \varphi$) if and only if $s \in [\![\varphi]\!]_M$ (the propositional context ρ can be omitted since φ does not contain free variables).

Propositional Dynamic Logic with Looping. In addition to plain L_μ operators, we will use the modalities of PDL-Δ (*Propositional Dynamic Logic with Looping*) [29], which characterize finite (resp. infinite) sequences of transitions whose concatenated actions form words belonging to regular (resp. ω-regular) languages. The syntax and semantics of PDL-Δ (defined by translation to L_μ) are given in Figure 2. Regular formulas (noted β) are built from action formulas and the testing (?), concatenation (.), choice (|), and transitive reflexive closure ($*$) operators. Apart from Boolean connectors, state formulas are built from the possibility modality ($\langle\ \rangle$) and the infinite looping operator ($\langle\ \rangle\, @$), both containing regular formulas. Derived state operators are defined as follows: $[\beta]\, \varphi = \neg\, \langle\beta\rangle\, \neg\varphi$ is the necessity modality, and $[\beta]\, \dashv = \neg\, \langle\beta\rangle\, @$ is the saturation operator.

A transition sequence satisfies a formula β if the word obtained by concatenating all actions of the sequence belongs to the regular language defined by β. The testing operator makes it possible to specify state formulas that must hold in the intermediate states of a transition sequence. The possibility modality $\langle\beta\rangle\, \varphi_1$ (resp. the necessity modality $[\beta]\, \varphi_1$) denotes the states for which some (resp. all) of their outgoing transition sequences satisfying β lead to states satisfying φ_1. The infinite looping operator $\langle\beta\rangle\, @$ (resp. the saturation operator

$$\beta ::= \alpha \mid \varphi? \mid \beta_1.\beta_2 \mid \beta_1|\beta_2 \mid \beta_1^*$$
$$\varphi ::= \mathsf{false} \mid \neg\varphi_1 \mid \varphi_1 \vee \varphi_2 \mid \langle\beta\rangle\,\varphi_1 \mid \langle\beta\rangle\,@$$

$$\begin{aligned}
\langle\varphi'?\rangle\,\varphi &= \varphi' \wedge \varphi \\
\langle\beta_1.\beta_2\rangle\,\varphi &= \langle\beta_1\rangle\,\langle\beta_2\rangle\,\varphi \\
\langle\beta_1|\beta_2\rangle\,\varphi &= \langle\beta_1\rangle\,\varphi \vee \langle\beta_2\rangle\,\varphi \\
\langle\beta^*\rangle\,\varphi &= \mu X.(\varphi \vee \langle\beta\rangle\,X) \\
\langle\beta\rangle\,@ &= \nu X.\,\langle\beta\rangle\,X
\end{aligned}$$

Figure 2: Syntax and semantics of PDL-Δ

$[\beta]\,\dashv$) denotes the states having some (resp. no) outgoing transition sequence consisting of an infinite concatenation of sub-sequences satisfying β.

The operators of PDL-Δ can be freely mixed with those of L_μ, and in practice they allow a much more concise and intuitive description of properties. The variant of L_μ extended with PDL-Δ operators, noted L_μ^{reg}, has been considered and efficiently implemented in [21] (in fact, the syntax used for PDL-Δ operators in Fig. 2 is that of L_μ^{reg} and not the original one). In the remainder of the report, we will use L_μ^{reg} whenever possible for specifying properties.

Divergence-sensitive branching bisimulation. As equivalence relation between LTSs, we consider divergence-sensitive branching bisimulation [15, 14], which preserves branching-time properties such as inevitable reachability and also the existence of divergences (τ-cycles), while still making possible substantial reductions of LTSs. This relation is finer than plain branching bisimulation and weak bisimulation [25] (none of which preserves divergences), therefore being a good candidate for comparing the behaviour of concurrent systems.

Definition 1 (Divergence-Sensitive Branching Bisimulation [15]) *A binary relation R on the set of states S is a* divergence-sensitive branching bisimulation *if R is symmetric and $s\ R\ t$ implies that*

- *if $s \xrightarrow{b} s'$ then*
  - *either $b = \tau$ with $s'\ R\ t$;*
  - *or $t \Longrightarrow \hat{t} \xrightarrow{b} t'$ with $s\ R\ \hat{t}$ and $s'\ R\ t'$.*
- *if for all $k \geq 0$ and $s = s_0$, $s_k \xrightarrow{\tau} s_{k+1}$ then for all $\ell \geq 0$ and $t = t_0$, $t_\ell \xrightarrow{\tau} t_{\ell+1}$ and $s_k\ R\ t_\ell$ for all k, ℓ.*

Two states s and t are divergence-sensitive branching bisimilar*, noted $s \approx_{br}^{ds} t$, if there is a divergence-sensitive branching bisimulation R with $s\ R\ t$.*

When expressing certain properties (e.g., inevitable reachability), it is necessary to characterize deadlock states in the LTS, i.e., states from which the execution cannot progress anymore. From the $\approx_{br}^{ds}$ point of view, deadlock states are precisely those states leading eventually to sink states (i.e., states without successors) after a finite number of τ-transitions. These states can be characterized by the PDL-Δ formula below:

$$deadlock = [\mathsf{true}^*.\neg\tau]\,\mathsf{false} \wedge [\tau]\,\dashv$$

where the box modality forbids the reachability of visible actions and the saturation operator forbids the presence of divergences.

3 Maximal Hiding

When checking a state formula φ over an LTS, some actions of the LTS can be hidden (i.e., renamed into τ) without disturbing the interpretation of φ.

Definition 2 (Hiding Set) *Let α be an action formula interpreted over a set of actions A. The* hiding set *of α w.r.t. A is defined as follows:*

$$h_A(\alpha) = \begin{cases} [\![\alpha]\!]_A & \text{if } \tau \models \alpha \\ A \setminus [\![\alpha]\!]_A & \text{if } \tau \not\models \alpha \end{cases}$$

The hiding set of a state formula φ w.r.t. A, noted $h_A(\varphi)$, is defined as the intersection of $h_A(\alpha)$ for all action subformulas α of φ.

Definition 3 (Hiding) *Let A be a set of actions and $B \subseteq A$. The* hiding *of an action $b \in A$ w.r.t. B is defined as follows:*

$$hide_B(b) = \begin{cases} b & \text{if } b \notin B \\ \tau & \text{if } b \in B \end{cases}$$

The hiding of an LTS *$M = \langle S, A, T, s_0 \rangle$ w.r.t. B is defined as follows:*

$$hide_B(\langle S, A, T, s_0 \rangle) = \left\langle S, (A \setminus B) \cup \{\tau\}, \{s_1 \overset{hide_B(b)}{\rightarrow} s_2 \mid s_1 \overset{b}{\rightarrow} s_2 \in T\}, s_0 \right\rangle.$$

The following lemma states that hiding an action b w.r.t. the hiding set of an action formula α does not disturb the satisfaction of α by b.

Lemma 1 *Let α be an action formula interpreted over a set of actions A. Then, the hiding set $h_A(\alpha)$ is the maximal set $B \subseteq A$ such that:*

$$b \models_A \alpha \Leftrightarrow hide_B(b) \models_A \alpha$$

for any action $b \in A$.

Proof. We show first that $h_A(\alpha)$ satisfies the statement in the lemma. Let $b \in h_A(\alpha)$. By Definition 3, this means $hide_{h_A(\alpha)}(b) = \tau$. Two cases are possible. If $\tau \models \alpha$, then $h_A(\alpha) = [\![\alpha]\!]_A$ by Definition 2, and therefore $b \models_A \alpha$. If $\tau \not\models \alpha$, then $h_A(\alpha) = A \setminus [\![\alpha]\!]_A$ by Definition 2, and therefore $b \not\models_A \alpha$.

To show the maximality of $h_A(\alpha)$, suppose there exists $b \in A \setminus h_A(\alpha)$ such that $b \models_A \alpha \Leftrightarrow \tau \models_A \alpha$. Two cases are possible, both leading to a contradiction. If $\tau \models \alpha$, then $h_A(\alpha) = [\![\alpha]\!]_A$ by Definition 2, and since $b \notin h_A(\alpha)$, this means $b \not\models \alpha$. If $\tau \not\models \alpha$, then $h_A(\alpha) = A \setminus [\![\alpha]\!]_A$ by Definition 2, and since $b \notin h_A(\alpha)$, this means $b \models \alpha$. □

To enable LTS reductions prior to (or simultaneously with) the verification of a state formula φ, it is desirable to hide as many actions as possible in the LTS, i.e., all actions in $h_A(\varphi)$. The following proposition ensures that this hiding preserves the interpretation of φ.

Proposition 1 (Maximal Hiding) *Let $M = \langle S, A, T, s_0 \rangle$ be an* LTS*, φ be a state formula, and $B \subseteq h_A(\varphi)$. Then:*

$$[\![\varphi]\!]_M \, \rho = [\![\varphi]\!]_{hide_B(M)} \, \rho$$

for any propositional context ρ.

Proof. We proceed by structural induction on φ. We give here the most interesting case $\varphi ::= \langle \alpha \rangle \, \varphi_1$, the other cases being handled in Appendix A. Since $B \subseteq h_A(\langle \alpha \rangle \, \varphi_1)$ by hypothesis and $h_A(\langle \alpha \rangle \, \varphi_1) = h_A(\alpha) \cap h_A(\varphi_1)$ by Definition 2, it follows that $B \subseteq h_A(\alpha)$ and $B \subseteq h_A(\varphi_1)$. Therefore, we can apply the induction hypothesis for φ_1, B and Lemma 1 for α, B, which yields:

$$\begin{array}{ll}
[\![\langle \alpha \rangle \, \varphi_1]\!]_{hide_B(M)} \, \rho & = \text{by definition of } [\![\ ]\!] \text{ and } hide_B(M) \\
\{s \in S \mid \exists s \stackrel{hide_B(b)}{\rightarrow} s'. hide_B(b) \models_A \alpha \wedge & \\
\qquad s' \in [\![\varphi_1]\!]_{hide_B(M)} \, \rho\} & = \text{by induction hyp. and Lemma 1} \\
\{s \in S \mid \exists s \stackrel{b}{\rightarrow} s'. b \models_A \alpha \wedge s' \in [\![\varphi_1]\!]_M \, \rho\} & = \text{by definition of } [\![\ ]\!] \\
[\![\langle \alpha \rangle \, \varphi_1]\!]_M \, \rho. &
\end{array}$$

□

In general, for a given property, there are several μ-calculus formulas φ specifying it, with different hiding sets $h_A(\varphi)$. To take advantage of Proposition 1, one must choose a formula φ with a hiding set as large as possible. Intuitively, in such *well-specified* formula φ, all action subformulas are relevant for the interpretation of φ on an LTS. For example, the following formula is not well-specified:

$$\varphi = \mu X. (\langle a_1 \rangle \, \mathsf{true} \vee ([a_2] \, \mathsf{false} \vee \langle a_2 \rangle \, \mathsf{true}) \wedge \langle a_3 \rangle \, X))$$

because its subformula $[a_2] \, \mathsf{false} \vee \langle a_2 \rangle \, \mathsf{true}$ is a tautology and could be deleted from φ without changing its meaning. The presence of this subformula yields the hiding set $h_A(\varphi) = A \setminus \{a_1, a_2, a_3\}$, whereas deleting it yields a larger hiding set $h_A(\varphi) = A \setminus \{a_1, a_3\}$. We do not attempt here to check well-specifiedness automatically, and will assume below that state formulas are well-specified.

For instance, consider the L_μ^{reg} formula below, expressing the inevitable reachability of a *recv* action after every *send* action:

$$\varphi = [\mathsf{true}^*.send]\, \mu X.(\neg deadlock \wedge [\neg recv]\, X)$$

When checking φ on an LTS, one can hide all actions in $h_A(\varphi) = h_A(send) \cap h_A(\neg recv) = (A \setminus [\![send]\!]_A) \cap [\![\neg recv]\!]_A = (A \setminus \{send\}) \cap (A \setminus \{recv\}) = A \setminus \{send, recv\}$, i.e., all actions other than *send* and *recv*, without changing the interpretation of the formula.

4 Mu-Calculus Fragment Compatible with $\approx_{br}^{ds}$

When minimizing an LTS modulo a weak bisimulation relation, such as $\approx_{br}^{ds}$ [15], the degree of reduction achieved is often directly proportional to the percentage of τ-transitions contained in the original LTS. Therefore, Proposition 1 provides, for a given L_μ formula, the highest potential for reduction, by enabling as many actions as possible to be hidden in the LTS. However, this proposition does not indicate which L_μ formulas are *compatible* with $\approx_{br}^{ds}$, i.e., are preserved by reduction modulo this relation. We propose below a fragment of L_μ satisfying this property.

4.1 Mu-calculus fragment L_μ^{dsbr}

The L_μ fragment we consider here, called L_μ^{dsbr}, is defined in Figure 3. Compared to standard L_μ, this fragment differs in two respects:

1. It introduces two new weak operators $\langle(\varphi_1?.\alpha_1)^*\rangle\, \psi$ and $\langle\varphi_1?.\alpha_1\rangle\, @$ expressed in PDL-Δ, where the action formulas α_1 must capture the invisible action. The weak possibility modality $\langle(\varphi_1?.\alpha_1)^*\rangle\, \psi$ characterizes the states having an outgoing sequence of (0 or more) α_1-transitions whose intermediate states satisfy φ_1 and whose terminal state satisfies ψ. The weak infinite looping operator $\langle\varphi_1?.\alpha_1\rangle\, @$ characterizes the states having an infinite outgoing sequence of α_1-transitions whose intermediate states satisfy φ_1. When the φ_1 subformula occurring in a weak operator is true, it can be omitted, because in this case the operator becomes $\langle\alpha_1^*\rangle\, \psi$ or $\langle\alpha_1\rangle\, @$.

2. The occurrence of strong modalities $\langle\alpha_2\rangle\, \varphi$ and $[\alpha_2]\, \varphi$ is restricted syntactically such that these modalities must contain action formulas α_2 denoting visible actions only, and that they can appear only after a weak possibility modality $\langle(\varphi_1?.\alpha_1)^*\rangle$ or weak necessity modality $[(\varphi_1?.\alpha_1)^*]$. The intuition is that visible transitions matched by a strong modality will remain in the LTS after maximal hiding and $\approx_{br}^{ds}$ minimization, and the transition sequences preceding them can become invisible or even disappear in the minimized LTS without affecting the interpretation of the formula, because these sequences are still captured by the weak modality immediately preceding the current strong modality.

$$\varphi ::= \langle (\varphi_1?.\alpha_1)^* \rangle\, \psi \mid \langle \varphi_1?.\alpha_1 \rangle\, @ \mid \mathsf{false} \mid \neg\varphi_1 \mid \varphi_1 \vee \varphi_2 \mid X \mid \mu X.\varphi_1$$
$$\psi ::= \varphi \mid \langle \alpha_2 \rangle\, \varphi \mid \neg\psi_1 \mid \psi_1 \vee \psi_2$$
where $\tau \in [\![\alpha_1]\!]_A$ and $\tau \notin [\![\alpha_2]\!]_A$

$$[\![\langle (\varphi_1?.\alpha_1)^* \rangle\, \psi]\!]_M\, \rho = \{s \in S \mid \exists m \geq 0. s = s_0 \wedge (\forall 0 \leq i < m. s_i \xrightarrow{b_{i+1}} s_{i+1} \in T \wedge\ b_{i+1} \in [\![\alpha_1]\!]_A \wedge\ s_i \in [\![\varphi_1]\!]_M\, \rho) \wedge\ s_m \in [\![\psi]\!]_M\, \rho\}$$
$$[\![\langle \varphi_1?.\alpha_1 \rangle\, @]\!]_M\, \rho = \{s \in S \mid s = s_0 \wedge \forall i \geq 0. (s_i \xrightarrow{b_{i+1}} s_{i+1} \in T \wedge b_{i+1} \in [\![\alpha]\!]_A \wedge\ s_i \in [\![\varphi_1]\!]_M\, \rho)\}$$

Figure 3: Syntax and semantics of the L_μ^{dsbr} fragment

The deadlock formula defined in Section 2 belongs to L_μ^{dsbr}, since it can be rewritten as follows by eliminating the concatenation operator:

$$\mathit{deadlock} = [\mathsf{true}^*.\neg\tau]\, \mathsf{false} \wedge [\tau]\, \dashv\ = [\mathsf{true}^*]\, [\neg\tau]\, \mathsf{false} \wedge [\tau]\, \dashv$$

The response formula given in Section 3 can also be reformulated in L_μ^{dsbr}:

$$[\mathsf{true}^*.\mathit{send}]\, \mu X. (\neg \mathit{deadlock} \wedge [\neg \mathit{recv}]\, X) =$$
$$[\mathsf{true}^*]\, [\mathit{send}]\, ([(\neg \mathit{recv})^*]\, \neg \mathit{deadlock} \wedge [\neg \mathit{recv}]\, \dashv)$$

The subformula stating the inevitable reachability of a *recv* action, initially expressed using a minimal fixed point operator, was replaced by the conjunction of a weak necessity modality forbidding the occurrence of deadlocks before a *recv* action has been reached, and a weak saturation operator forbidding the presence of cycles not passing through a *recv* action.

In [14, Corollary 4.4], it was shown that $\approx_{br}^{ds}$ is an equivalence with the so-called *stuttering property*:

Definition 4 (Stuttering) *Let $M = \langle S, A, T, s_0 \rangle$ be an* LTS *and let $s_1, s_2 \in S$ such that $s_1 \approx_{br}^{ds} s_2$. If $s_1 \xrightarrow{\tau} s_1^1 \xrightarrow{\tau} \cdots \xrightarrow{\tau} s_1^m \xrightarrow{\tau} s_1'$ $(m \geq 0)$ and $s_1' \approx_{br}^{ds} s_2$, then $\forall 1 \leq i \leq m. s_1^i \approx_{br}^{ds} s_2$.*

Using the stuttering property, we can prove the following lemma.

Lemma 2 *Let $M = \langle S, A, T, s_0 \rangle$ be an* LTS *and let $A' \subseteq A$ with $\tau \in A'$ and $s_1, s_2 \in S$ such that $s_1 \approx_{br}^{ds} s_2$. Then for all $m \geq 0$ with $s_1 = s_1^0$ and $\forall 0 \leq i < m. s_1^i \xrightarrow{b_i} s_1^{i+1} \in T$ $(b_i \in A')$, there exists $k \geq 0$ such that $s_2 = s_2^0$ and $\forall 0 \leq j < k. (s_2^j \xrightarrow{b_j'} s_2^{j+1} \in T$ $(b_j' \in A') \wedge \exists 0 \leq i < m. s_1^i \approx_{br}^{ds} s_2^j)$, and $s_1^m \approx_{br}^{ds} s_2^k$.*

Proof. We proceed by induction on m.

1. *Base case:* $m = 0$, hence $s_1 = s_1^0 = s_1^m$. Clearly, we can choose $k = 0$ and $s_2 = s_2^0 = s_2^k$.

2. *Inductive case:* $s_1^0 \xrightarrow{b_1} s_1^1 \cdots s_1^{m-1} \xrightarrow{b_m} s_1^m \xrightarrow{b_{m+1}} s_1^{m+1}$. By the induction hypothesis, there exists $k \geq 0$ such that $s_2 = s_2^0$ and $\forall 0 \leq j < k.(s_2^j \xrightarrow{b'_j} s_2^{j+1} \in T\ (b'_j \in A') \wedge \exists 0 \leq i < m.s_1^i \approx_{br}^{ds} s_2^j)$, and $s_1^m \approx_{br}^{ds} s_2^k$. We show that it also holds for $m+1$. We distinguish two cases for $s_1^m \xrightarrow{b_{m+1}} s_1^{m+1}$:

   (a) $b_{m+1} = \tau$. Since $s_1^m \approx_{br}^{ds} s_2^k$, by Definition 1, also $s_1^{m+1} \approx_{br}^{ds} s_2^k$.

   (b) $b_{m+1} \neq \tau$. Since $s_1^m \approx_{br}^{ds} s_2^k$, by Definition 1, $s_2^k \Longrightarrow \hat{s}_2 \xrightarrow{b_{m+1}} s'_2$, with $s_1^m \approx_{br}^{ds} \hat{s}_2$, and $s_1^{m+1} \approx_{br}^{ds} s'_2$. Say that $s_2^k \Longrightarrow \hat{s}_2$ consists of c τ-steps $s_2^k \xrightarrow{\tau_1} s_2^{k+1} \cdots s_2^{k+c-1} \xrightarrow{\tau_c} s_2^{k+c}$ with $s_2^{k+c} = \hat{s}_2$. By Definition 4, for all $k \leq i \leq k+c$, we have $s_1^m \approx_{br}^{ds} s_2^i$. Hence, there exists a matching sequence from s_2 of length $k+c+1$ with $s_2^{k+c+1} = s'_2$. Note that $\tau_1, \ldots, \tau_c \in A'$.

□

A propositional context $\rho : \mathcal{X} \to 2^S$ is said to be $\approx_{br}^{ds}$*-closed* if for all states $s_1, s_2 \in S$ such that $s_1 \approx_{br}^{ds} s_2$ and for any propositional variable $X \in \mathcal{X}$, $s_1 \in \rho(X) \Leftrightarrow s_2 \in \rho(X)$. Now we can state the main result about L_μ^{dsbr}, namely that this fragment is compatible with the $\approx_{br}^{ds}$ relation.

Proposition 2 (Compatibility with $\approx_{br}^{ds}$) *Let $M = \langle S, A, T, s_0 \rangle$ be an* LTS *and let $s_1, s_2 \in S$ such that $s_1 \approx_{br}^{ds} s_2$. Then:*

$$s_1 \in [\![\varphi]\!]_M\, \rho \Leftrightarrow s_2 \in [\![\varphi]\!]_M\, \rho$$

for any state formula φ of L_μ^{dsbr} and any $\approx_{br}^{ds}$-closed propositional context ρ.

Proof. We proceed by structural induction on φ. We give here the most interesting cases, the other cases being handled in Appendix A.

Case $\varphi ::= \langle (\varphi_1?.\alpha_1)^* \rangle\, \psi$. Let $s_1, s_2 \in S$ such that $s_1 \approx_{br}^{ds} s_2$ and assume that $s_1 \in [\![\langle (\varphi_1?.\alpha_1)^* \rangle\, \psi]\!]_M\, \rho$, i.e., $s_1 \in \{s \in S \mid \exists m \geq 0.s = s_0 \wedge (\forall 0 \leq i < m.s_i \xrightarrow{b_{i+1}} s_{i+1} \in T \wedge b_{i+1} \in [\![\alpha_1]\!]_A \wedge\ s_i \in [\![\varphi_1]\!]_M\, \rho) \wedge\ s_m \in [\![\psi]\!]_M\, \rho\}$. This means that:

$$\begin{aligned} &\exists m \geq 0.s_1 = s'_0\ \wedge\ (\forall 0 \leq i < m.s'_i \xrightarrow{b_{i+1}} s'_{i+1} \in T \\ &\quad \wedge\ b_{i+1} \in [\![\alpha_1]\!]_A \wedge\ s'_i \in [\![\varphi_1]\!]_M\, \rho) \wedge\ s'_m \in [\![\psi]\!]_M\, \rho\} \end{aligned} \tag{1}$$

We have to prove that $s_2 \in [\![\langle (\varphi_1?.\alpha_1)^* \rangle\, \psi]\!]_M\, \rho$, which means that:

$$\begin{aligned} &\exists k \geq 0.s_2 = s''_0\ \wedge\ (\forall 0 \leq j < k.s''_i \xrightarrow{b'_{j+1}} s''_{i+1} \in T \\ &\quad \wedge\ b'_{j+1} \in [\![\alpha_1]\!]_A \wedge\ s''_i \in [\![\varphi_1]\!]_M\, \rho) \wedge\ s''_k \in [\![\psi]\!]_M\, \rho\} \end{aligned} \tag{2}$$

First, since $s_1 \approx_{br}^{ds} s_2$, $\tau \in [\![\alpha_1]\!]_A$, and (1), by Lemma 2 with $A' = [\![\alpha_1]\!]_A$, there exists $k \geq 0$ with $s_2 = s''_0$ such that $\forall 0 \leq j < k.(s''_j \xrightarrow{b'_{j+1}} s''_{j+1} \in T(b'_{j+1} \in [\![\alpha_1]\!]_A) \wedge \exists 0 \leq i < m.s'_i \approx_{br}^{ds} s''_j)$

and $s'_m \approx^{ds}_{br} s''_k$. Furthermore, for all $0 \leq j < k$, since there exists $0 \leq i < m.s'_i \approx^{ds}_{br} s''_j$ and $s'_i \in [\![\varphi_1]\!]_M\,\rho$, by the induction hypothesis, it follows that $s''_j \in [\![\varphi_1]\!]_M\,\rho$. Finally, since $s'_m \approx^{ds}_{br} s''_k$ and $s'_m \in [\![\psi]\!]_M\,\rho$, we will show that $s''_k \in [\![\psi]\!]_M\,\rho$ by induction on the structure of ψ. First, we can assume that there is no $\hat{s}'' \in S$ such that $s''_k \xrightarrow{\tau} \hat{s}'' \in T$. If this is not true, since $\tau \in [\![\alpha_1]\!]_A$, we can choose $s''_{k+1} = \hat{s}''$ and increase k by one. This can be repeated until there is no $\hat{s}'' \in S$ such that $s''_k \xrightarrow{\tau} \hat{s}'' \in T$. For ψ, we distinguish four cases:

- $\psi ::= \varphi$. By the induction hypothesis, $s''_k \in [\![\psi]\!]_M\,\rho$.
- $\psi ::= \langle\alpha_2\rangle\,\varphi$. Since $s'_m \in [\![\langle\alpha_2\rangle\,\varphi]\!]_M\,\rho$, we have $s'_m \in \{s \in S \mid \exists s \xrightarrow{a} s' \in T.a \in [\![\alpha_2]\!]_A \wedge\ s' \in [\![\varphi]\!]_M\,\rho\}$, hence there exists $s'_m \xrightarrow{a} s' \in T$ with $a \in [\![\alpha_2]\!]_A$. Since $s'_m \approx^{ds}_{br} s''_k$, $\tau \notin [\![\alpha_2]\!]_A$, and $s''_k \xrightarrow{\tau} \hat{s}'' \notin T$, by Definition 1, there must exist $s''_k \xrightarrow{a} \hat{s}'' \in T$ with $s' \approx^{ds}_{br} \hat{s}''$. Since $s' \in [\![\varphi]\!]_M\,\rho$, by the induction hypothesis, $\hat{s}'' \in [\![\varphi]\!]_M\,\rho$, hence $s''_k \in \{s \in S \mid \exists s \xrightarrow{a} s' \in T.a \in [\![\alpha_2]\!]_A \wedge\ s' \in [\![\varphi]\!]_M\,\rho\}$, i.e., $s''_k \in [\![\psi]\!]_M\,\rho$.
- $\psi ::= \neg\psi_1$. Since $s'_m \in [\![\neg\psi_1]\!]_M\,\rho$, we have $s'_m \in S \setminus [\![\psi_1]\!]_M\,\rho$. By the induction hypothesis for ψ, also $s''_k \in S \setminus [\![\psi_1]\!]_M\,\rho$, hence $s''_k \in [\![\neg\psi_1]\!]_M\,\rho$.
- $\psi ::= \psi_1 \vee \psi_2$. Since $s'_m \in [\![\psi_1 \vee \psi_2]\!]_M\,\rho$, i.e., $s'_m \in [\![\psi_1]\!]_M\,\rho \cup [\![\psi_2]\!]_M\,\rho$, i.e., $s'_m \in [\![\psi_1]\!]_M\,\rho \vee s'_m \in [\![\psi_2]\!]_M\,\rho$, by the induction hypothesis for ψ, we have $s''_k \in [\![\psi_1]\!]_M\,\rho \vee s''_k \in [\![\psi_2]\!]_M\,\rho$, i.e., $s''_k \in [\![\psi_1]\!]_M\,\rho \cup [\![\psi_2]\!]_M\,\rho$, i.e., $s''_k \in [\![\psi_1 \vee \psi_2]\!]_M\,\rho$.

Hence, (2) holds. The converse implication (by considering $s_2 \in [\![\langle(\varphi_1?.\alpha_1)^*\rangle\,\psi]\!]_M\,\rho$) holds by a symmetric argument.

Case $\varphi ::= \langle\varphi_1?.\alpha_1\rangle\,@$. Let $s_1, s_2 \in S$ such that $s_1 \approx^{ds}_{br} s_2$ and assume that $s_1 \in [\![\langle\varphi_1?.\alpha_1\rangle\,@]\!]_M\,\rho$, i.e., $s_1 \in \{s \in S \mid s = s_0 \wedge \forall i \geq 0.(s_i \xrightarrow{b_i} s_{i+1} \wedge b_i \in [\![\alpha_1]\!]_A \wedge\ s_i \in [\![\varphi_1]\!]_M\,\rho)\}$. This means that:

$$s_1 = s'_0 \wedge \forall i \geq 0.(s'_i \xrightarrow{b_i} s'_{i+1} \wedge b_i \in [\![\alpha_1]\!]_A \wedge\ s'_i \in [\![\varphi_1]\!]_M\,\rho) \tag{3}$$

We have to prove that $s_2 \in [\![\langle\varphi_1?.\alpha_1\rangle\,@]\!]_M\,\rho$, which means that:

$$s_2 = s''_0 \wedge \forall j \geq 0.(s''_j \xrightarrow{b'_j} s''_{j+1} \wedge b'_j \in [\![\alpha_1]\!]_A \wedge\ s''_j \in [\![\varphi_1]\!]_M\,\rho) \tag{4}$$

Since $s_1 \approx^{ds}_{br} s_2$, $\tau \in [\![\alpha_1]\!]_A$, and (3), by Lemma 2 with $A' = [\![\alpha_1]\!]_A$, for any finite prefix of length $m \geq 0$ of the infinite path π from s_1, there exists a finite path of length $k \geq 0$ from s_2 such that $s_2 = s''_0 \wedge \forall 0 \leq j < k.(s''_j \xrightarrow{b'_j} s''_{j+1} \wedge b'_j \in [\![\alpha_1]\!]_A \wedge \exists 0 \leq i < m.s'_i \approx^{ds}_{br} s''_j)$ and $s'_m \approx^{ds}_{br} s''_k$, hence, by the induction hypothesis, for all $0 \leq j \leq k$, we have $s''_j \in [\![\varphi_1]\!]_M\,\rho$. We distinguish two cases:

1. π contains an infinite number of transitions with a label in $[\![\alpha_1]\!]_A \setminus \{\tau\}$. Repeatedly applying the above reasoning for intermediate states in π yields that (4) holds for s_2.

2. π contains a finite number of transitions with a label in $[\![\alpha_1]\!]_A \setminus \{\tau\}$. Then, there exists an $\hat{s}$ reachable from s_1 such that from $\hat{s}$, an infinite τ-path exists. By the earlier reasoning, there exists an $\hat{s}'$ reachable from s_2 such that $\hat{s} \approx_{br}^{ds} \hat{s}'$ and for all states s''_j in the path from s_2 to $\hat{s}'$, we have $s''_j \in [\![\varphi_1]\!]_M \rho$. Finally, since $\hat{s} \approx_{br}^{ds} \hat{s}'$, by the second clause of Definition 1, there also exists an infinite τ-path π' from $\hat{s}'$. Finally, by Definition 1 and repeated application of Definition 4, it follows that for all states s''_j in π', $\hat{s} \approx_{br}^{ds} s''_j$, hence by the induction hypothesis, $s''_j \in [\![\varphi_1]\!]_M \rho$. Therefore, (4) holds for s_2.

The converse implication (by considering $s_2 \in [\![\langle \varphi_1 ?.\alpha_1 \rangle \, @]\!]_M \rho$) holds by a symmetric argument. □

Proposition 2 makes it possible to reduce an LTS (after applying maximal hiding) modulo $\approx_{br}^{ds}$ before the verification of a closed L_μ^{dsbr} formula. It follows that L_μ^{dsbr} is also compatible with all equivalence relations weaker than $\approx_{br}^{ds}$, such as $\tau^*.a$ [10] and weak [25] bisimulations. For practical purposes, it is desirable to use a temporal logic sufficiently expressive to capture the essential classes of properties (safety, liveness, fairness). Thus, the question is whether L_μ^{dsbr} subsumes the existing temporal logics compatible with $\tau^*.a$ and weak bisimulations; in Subsections 4.2 and 4.3, we show that this is indeed the case.

4.2 Subsuming μACTL\X

ACTL [26] is a branching-time logic similar to CTL [4], but interpreted on LTSs. It consists of action formulas (noted α) and state formulas (noted φ) expressing properties about actions and states of an LTS, respectively. The temporal operators of ACTL\X (the fragment of the logic without the next-time operators) are defined in Table 1 by means of their encodings in L_μ proposed in [9]. The operator $\mathsf{E}[\varphi_1{}_\alpha \mathsf{U} \varphi_2]$ (resp. $\mathsf{A}[\varphi_1{}_\alpha \mathsf{U} \varphi_2]$) denotes the states from which some (resp. all) outgoing sequences lead, after 0 or more α-transitions (or τ-transitions) whose source states satisfy φ_1, to a state satisfying φ_2. The operator $\mathsf{E}[\varphi_1{}_{\alpha_1} \mathsf{U}_{\alpha_2} \varphi_2]$ (resp. $\mathsf{A}[\varphi_1{}_{\alpha_1} \mathsf{U}_{\alpha_2} \varphi_2]$) denotes the states from which some (resp. all) outgoing sequences lead, after 0 or more α_1-transitions (or τ-transitions) whose source states satisfy φ_1, to an α_2-transition whose source state satisfies φ_1 and whose target state satisfies φ_2. The action subformulas α, α_1, and α_2 denote visible actions only.

Table 1: Syntax and semantics of the ACTL\X temporal operators

Operator	Translation
$\mathsf{E}[\varphi_1{}_\alpha \mathsf{U} \varphi_2]$	$\mu X.(\varphi_2 \vee (\varphi_1 \wedge \langle \alpha \vee \tau \rangle X))$
$\mathsf{E}[\varphi_1{}_{\alpha_1} \mathsf{U}_{\alpha_2} \varphi_2]$	$\mu X.(\varphi_1 \wedge (\langle \alpha_2 \rangle \varphi_2 \vee \langle \alpha_1 \vee \tau \rangle X))$
$\mathsf{A}[\varphi_1{}_\alpha \mathsf{U} \varphi_2]$	$\mu X.(\varphi_2 \vee (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha \vee \tau)] \, \mathsf{false} \wedge [\alpha \vee \tau] \, X))$
$\mathsf{A}[\varphi_1{}_{\alpha_1} \mathsf{U}_{\alpha_2} \varphi_2]$	$\mu X.(\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha_1 \vee \alpha_2 \vee \tau)] \, \mathsf{false} \wedge [\alpha_2 \wedge \neg \alpha_1] \, \varphi_2 \wedge [\alpha_1 \wedge \alpha_2] \, (\varphi_2 \vee X) \wedge [\neg \alpha_2] \, X)$

ACTL\X was shown to be adequate with $\approx_{br}^{ds}$ [26]. Moreover, this logic was extended in [9] with fixed point operators, yielding a fragment of L_μ called μACTL\X, which is still adequate with $\approx_{br}^{ds}$. The temporal operators of ACTL\X can be translated in L_μ^{dsbr}, as stated by the following proposition.

Proposition 3 (Translation from ACTL\X to L_μ^{dsbr}) *The following identities relating formulas of L_μ and formulas of L_μ^{dsbr} hold:*

$$\mu X.(\varphi_2 \vee (\varphi_1 \wedge \langle \alpha \vee \tau \rangle X)) = \langle (\varphi_1?.\alpha \vee \tau)^* \rangle \varphi_2$$

$$\mu X.(\varphi_1 \wedge (\langle \alpha_2 \rangle \varphi_2 \vee \langle \alpha_1 \vee \tau \rangle X)) = \langle (\varphi_1?.\alpha \vee \tau)^* \rangle (\varphi_1 \wedge \langle \alpha_2 \rangle \varphi_2)$$

$$\begin{aligned}&\mu X.(\varphi_2 \vee (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha \vee \tau)]\, \mathsf{false} \wedge [\alpha \vee \tau]\, X)) = \\ &[(\neg\varphi_2?.\alpha \vee \tau)^*]\, (\varphi_2 \vee (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha \vee \tau)]\, \mathsf{false})) \wedge [\neg\varphi_2?.\alpha \vee \tau]\, \dashv\end{aligned}$$

$$\begin{aligned}&\mu X.(\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha_1 \vee \alpha_2 \vee \tau)]\, \mathsf{false} \wedge [\alpha_2 \wedge \neg\alpha_1]\, \varphi_2 \wedge \\ &\quad [\alpha_1 \wedge \alpha_2]\, (\varphi_2 \vee X) \wedge [\neg\alpha_2]\, X) = \\ &\nu X.\, [(\neg\alpha_2)^*]\, (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha_1 \vee \alpha_2 \vee \tau)]\, \mathsf{false} \wedge [\alpha_2 \wedge \neg\alpha_1]\, \varphi_2 \wedge \\ &\qquad [\alpha_1 \wedge \alpha_2]\, (\varphi_2 \vee X) \wedge X) \wedge \\ &\nu X.([\neg\alpha_2]\, \dashv \wedge [(\neg\alpha_2)^*]\, [\alpha_1 \wedge \alpha_2]\, (\varphi_2 \vee X)) \wedge \mu X.\, [(\neg\alpha_2)^*]\, [\alpha_1 \wedge \alpha_2]\, (\varphi_2 \vee X).\end{aligned}$$

Proposition 3 (proven in Appendix A) also ensures that μACTL\X is subsumed by L_μ^{dsbr}, since the fixed point operators are present in both logics. The L_μ^{dsbr} formulas corresponding to the $\mathsf{A}[\varphi_1{}_\alpha \mathsf{U} \varphi_2]$ and $\mathsf{A}[\varphi_1{}_{\alpha_1} \mathsf{U}_{\alpha_2} \varphi_2]$ operators are complex, and they serve solely for the purpose of establishing the translation to L_μ^{dsbr}. In practice, we will use the simpler L_μ encodings of the ACTL\X operators given in Table 1.

The response formula given in Section 3 can also be expressed in ACTL\X:

$$\mathsf{AG}_{\mathsf{true},send} \mathsf{A}[\mathsf{true}_{\mathsf{true}} \mathsf{U}_{recv} \mathsf{true}]$$

where $\mathsf{AG}_{\alpha_1,\alpha_2}\varphi = \neg\mathsf{EF}_{\alpha_1,\alpha_2}\neg\varphi = \neg\mathsf{E}[\mathsf{true}_{\alpha_1} \mathsf{U}_{\alpha_2} \neg\varphi]$ is the ACTL counterpart of the AG operator of CTL.

Finally, we claim that L_μ^{dsbr} is more powerful than μACTL\X. Indeed, the formula $\langle (\neg a)^* \rangle (\langle b \rangle \mathsf{true} \wedge \langle c \rangle \mathsf{true})$ does not seem to be expressible in μACTL\X because the occurrences of strong modalities expressing the existence of neighbor b- and c-transitions cannot be coupled individually with the preceding weak modality in order to use only the four temporal operators given in Table 1.

4.3 Subsuming selective and weak μ-calculus

The selective μ-calculus [3] introduces modalities indexed by sets of actions (represented here as action formulas) specifying the reachability of certain actions after sequences of (0 or more)

actions not belonging to the indexing set. The selective possibility modality can be encoded in L_μ^{dsbr} as follows:

$$\langle \alpha_1 \rangle_\alpha \, \varphi = \langle (\neg(\alpha_1 \vee \alpha))^* \rangle \, \langle \alpha_1 \rangle \, \varphi$$

where α, α_1 denote visible actions only. Selective μ-calculus is adequate w.r.t. the $\tau^*.a$ bisimulation: for each selective formula φ, one can hide all LTS actions other than those occurring in the modalities of φ and their index sets, and then minimize the LTS modulo $\tau^*.a$ without changing the interpretation of φ.

Selective μ-calculus was shown to be equivalent to L_μ, because the strong possibility modality of L_μ can be expressed in terms of the selective one: $\langle \alpha \rangle \, \varphi = \langle \alpha \rangle_{\mathsf{true}} \, \varphi$. However, this way of translating would yield no possibility of hiding actions, because the index sets would contain all actions of the LTS. For instance, the response formula given in Section 3 can be reformulated in selective μ-calculus as follows:

$$[send]_{\mathsf{false}} \, \mu X.(\langle \mathsf{true} \rangle_{\mathsf{true}} \, \mathsf{true} \wedge [\neg recv]_{\mathsf{true}} \, X)$$

The minimal fixed point subformula expressing the inevitable reachability of a *recv* action cannot be mapped to selective μ-calculus modalities, which forces the use of strong modalities (represented by selective modalities indexed by true). Therefore, the set of actions that can be hidden according to [3] without disturbing the interpretation of this formula is $A \setminus (\{send, recv\} \cup A) = \emptyset$, i.e., no hiding of actions prior to verification would be possible in that setting.

The weak (or observational) μ-calculus [28] is a fragment of L_μ adequate w.r.t. weak bisimulation. It introduces weak modalities specifying the reachability of certain actions preceded and followed by 0 or more τ-transitions. These weak modalities can be encoded in L_μ^{dsbr} as follows:

$$\langle\langle \alpha \rangle\rangle \, \varphi = \langle \tau^* \rangle \, \langle \alpha \rangle \, \langle \tau^* \rangle \, \varphi \qquad \langle\langle \rangle\rangle \, \varphi = \langle \tau^* \rangle \, \varphi$$

where α denotes visible actions only. The weak μ-calculus is able to express only weak safety and liveness properties; in particular, it does not capture the inevitable reachability of *recv* actions present in the example above.

5 Implementation and Experiments

We have implemented the maximal hiding and associated on-the-fly reduction machinery within the CADP verification toolbox [13]. We experimented on the effect of these optimizations on the EVALUATOR [21, 22] model checker, which evaluates formulas of the alternation-free fragment of L_μ^{reg} on LTSs on-the-fly. The tool works by first translating the L_μ^{reg} formulas into plain L_μ by eliminating the PDL regular operators, and then reformulating the verification problem as the resolution of a Boolean equation system (BES) [1], which is solved locally using the algorithms of the CÆSAR_SOLVE library [20] of CADP. EVALUATOR makes possible the definition of reusable libraries of derived operators (e.g., those of ACTL) and property patterns (e.g., the pattern system of [8]).

For the sake of efficiency, we focus on L_μ^{dsbr} formulas having a linear-time model checking complexity, namely the alternation-free fragment [6] extended with the infinite looping and saturation operators of PDL-Δ [29], which can be evaluated in linear time using the algorithms proposed in [22]. In the formulas below, we use the operators of PDL and ACTL\X, and the L_μ^{dsbr} formula $inev(a) = [(\neg a)^*]\,\neg deadlock \wedge [\neg a] \dashv$ as a shorthand for expressing the inevitable execution of an action a. For each verification experiment, we applied maximal hiding as stated in Proposition 1, and then carried out LTS reductions either prior to, or simultaneously with, the verification of the formula.

Strong bisimulation reduction. We considered first global verification, which consists in generating the LTS, applying maximal hiding, minimizing the LTS modulo strong bisimulation, and then verifying the properties on the minimized LTS. LTSs are represented as files in the compact BCG (*Binary Coded Graphs*) format of CADP. Hiding and minimization were carried out using the BCG_LABELS and BCG_MIN tools [7], the whole process being automated using SVL [12] scripts.

We considered the alternating bit protocol, implemented in LOTOS (demo 02 of CADP), and checked the following property, stating that the protocol behaves as a one-place buffer (initially empty) regarding the emission and reception of messages:

$$
\begin{array}{l}
[\mathsf{true}^*]\,(\\
\quad [get]\,(\mathsf{A}[\mathsf{true}_{\neg put}\mathsf{U}\langle\tau\rangle\,@] \wedge [(\neg put)^*.get]\,\mathsf{false}) \\
\quad \wedge \\
\quad [put]\,(\mathsf{A}[\mathsf{true}_{\neg get}\mathsf{U}\langle\tau\rangle\,@] \wedge [(\neg get)^*.put]\,\mathsf{false}))
\end{array}
$$

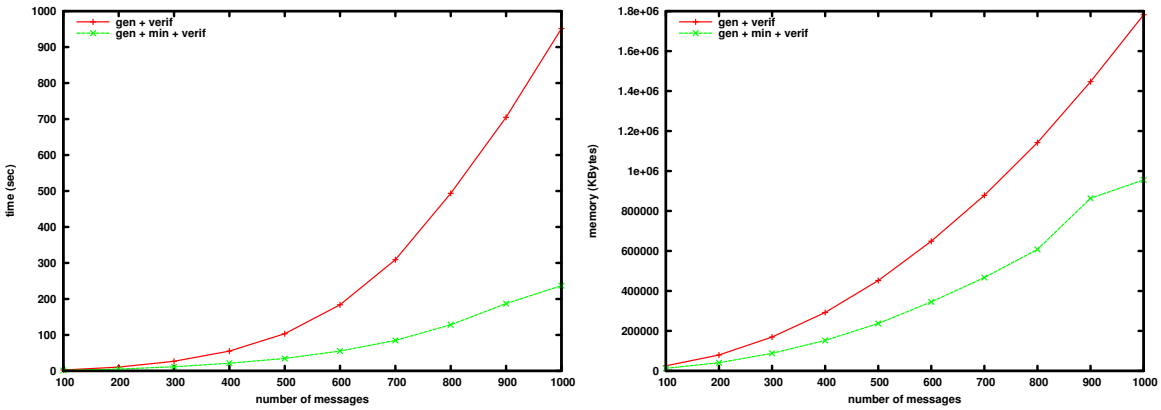

Figure 4: Effect of strong bisimulation minimization (Alternating Bit Protocol)

This formula allows hiding of every action other than *put* and *get*. The subformulas $\langle\tau\rangle\,@$ capture the divergences due to unreliable communication channels.

The overall time and peak memory needed for verification are shown in Figure 4 for increasingly larger configurations of the protocol. When strong bisimulation minimization is carried out before verification, we observe gains both in speedup and memory (factors 4 and 2 for the LTS corresponding to 1000 messages, having 12,196,201 states and 46,639,612 transitions), which become larger with the size of the LTS.

We also considered a token ring leader election protocol, implemented in LOTOS (experiment 6 in demo 17 of CADP), and checked the following property, stating that each station i on the ring accesses a shared resource (actions $open_i$ and $close_i$) in mutual exclusion with the other stations and each access is reachable (modulo the divergences due to unreliable communication channels):

$$[\mathsf{true}^*]\,([open_i.(\neg close_i)^*.open_j]\mathsf{false} \wedge \mathsf{A}[\mathsf{true}_{\mathsf{true}}\mathsf{U}\langle(\langle\mathsf{true}^*.open_i\rangle\,\mathsf{true})?.\tau\rangle\,@])$$

This formula belongs to L_μ^{dsbr} (after eliminating the concatenation operators and expanding the $\mathsf{A}[\mathsf{U}]$ operator) and allows hiding of every action other than *open* and *close*. The "$\langle ... \rangle\,@$" subformula of $\mathsf{A}[\mathsf{U}]$ expresses the existence of infinite τ-sequences whose intermediate states enable the potential reachability of an $open_i$ action.

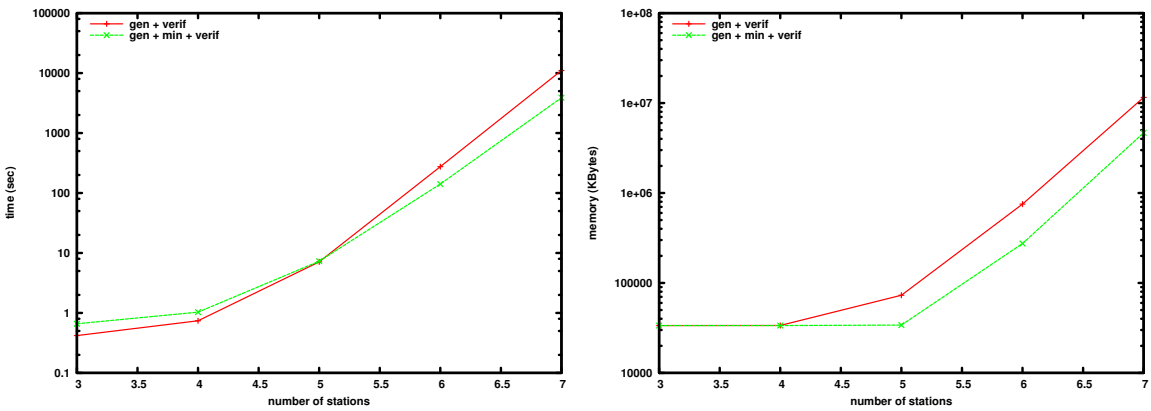

Figure 5: Effect of strong bisimulation minimization (Token Ring Protocol)

The overall time and peak memory needed for verification are shown in Figure 5 for increasingly larger configurations of the protocol. When strong bisimulation minimization is carried out before verification, we observe gains both in speedup and memory (factors 2.8 and 2.5 for the LTS corresponding to 7 stations, having 53,848,492 states and 214,528,176 transitions), which become larger with the size of the LTS.

Divergence-sensitive branching bisimulation reduction. To study the effect of $\approx_{br}^{ds}$ minimization, we considered Philips' Bounded Retransmission Protocol, implemented in LOTOS (demo 16 of CADP), and checked the following response property, expressing that

every emission of a data chunk from a packet is eventually followed by the reception of a confirmation:

$$[\mathsf{true}^*.in_data]\,\mathsf{A}[\mathsf{true}_{\neg in_data}\mathsf{U}_{in_conf}\mathsf{true}]$$

This formula belongs to L_μ^{dsbr} (after eliminating the concatenation operator and expanding the $\mathsf{A}[\mathsf{U}]$ operator) and allows hiding of every action other than *in_ data* and *in_ conf*.

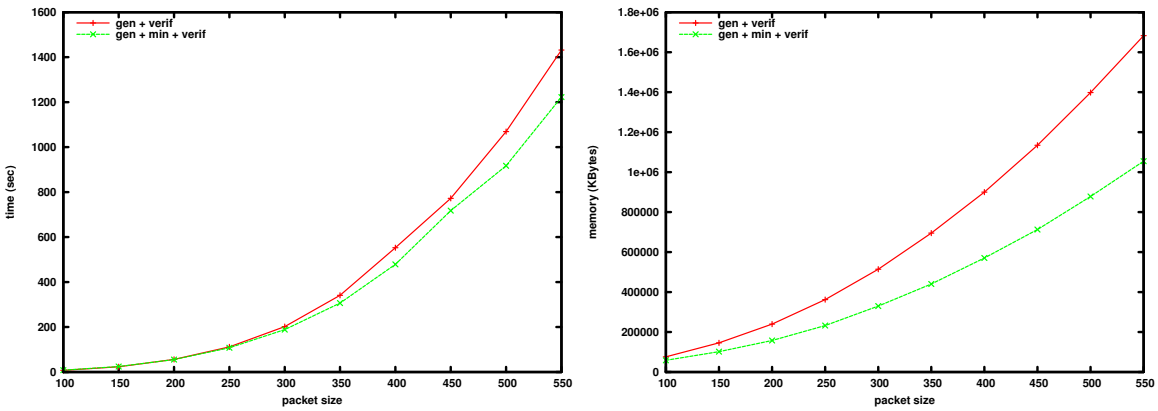

Figure 6: Effect of $\approx_{br}^{ds}$ minimization (Bounded Retransmission Protocol)

The overall time and peak memory needed for verification are shown in Figure 6 for increasingly larger configurations of the protocol. For this example, the presence of $\approx_{br}^{ds}$ bisimulation minimization yields mainly memory reductions (factor 1.6 for the LTS corresponding to data packets of length 550 and two retransmissions, having $12,450,383$ states and $14,880,828$ transitions).

On-the-fly τ-confluence reduction. Lastly, we examined the effect of τ-confluence reduction [16] carried out on-the-fly during the verification of formulas. This reduction, which preserves branching bisimulation, consists in identifying confluent τ-transitions (i.e., whose execution does not alter the observable behavior of the system), and giving them priority over their neighbors during the LTS traversal. The detection of confluent τ-transitions is done on-the-fly by reformulating the problem as a BES resolution [27, 23], which is performed locally using the algorithms of CÆSAR_SOLVE. In order to make the reduction compatible with $\approx_{br}^{ds}$, we enhanced the τ-confluence detection with the bookkeeping of divergence, by exploiting the τ-cycle compression algorithm proposed in [19].

We considered the distributed version of Erathosthene's sieve, implemented using LOTOS processes and EXP networks of automata (demo 36 of CADP). We checked the following formula, expressing that each prime number p fed as input to the sieve will be eventually delivered as output and each non-prime number q will be filtered:

$$[\mathsf{true}^*]\,([gen_p]inev(output_p) \wedge [gen_q.\mathsf{true}^*.\neg output_q]\mathsf{false})$$

This formula belongs to L_μ^{dsbr} (after eliminating the concatenation operators) and allows hiding of every action other than *gen* and *output*.

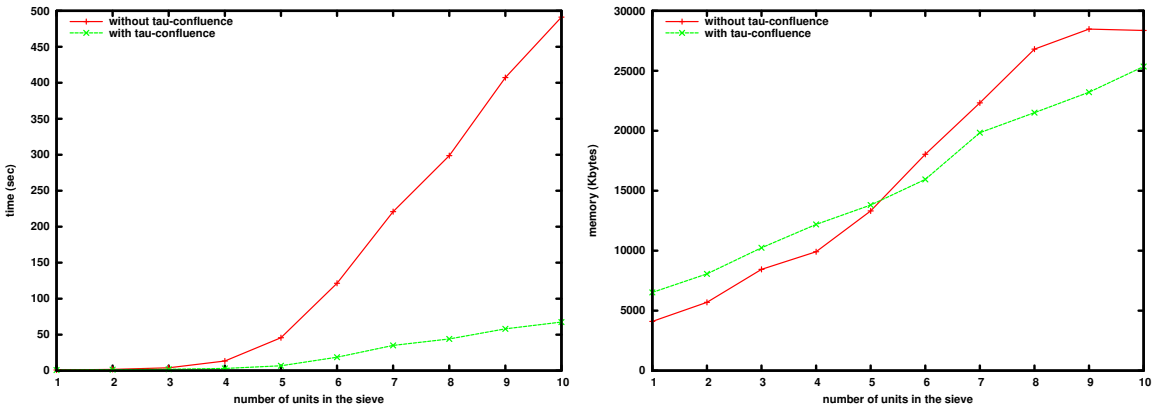

Figure 7: Effect of on-the-fly τ-confluence reduction (Erathostene's sieve)

The overall time and peak memory needed for verification are shown in Figure 7 for increasingly larger configurations of the sieve. We observe a substantial increase in speed in the presence of τ-confluence reduction (about one order of magnitude for a sieve with 10 units). The reduction in memory usage becomes apparent once the size of the system becomes sufficiently large, such that the memory overhead induced by the presence of the on-the-fly reduction machinery is compensated by the memory required for verifying the formula.

6 Conclusion and Future Work

We have presented two automatic techniques to improve the effectiveness of LTS reductions, both before and during system verification. The first technique involves maximal hiding of LTSs based on given L_μ formulas, such that the LTSs can be minimized modulo strong bisimulation. This technique is not intrusive, meaning that the user is not forced to write formulas in a specific way. In the second technique, formulas written in a specific fragment of L_μ, called L_μ^{dsbr}, are used to maximally hide LTSs such that they can be minimized modulo $\approx_{br}^{ds}$. Experimental results show the effectiveness of these techniques.

In future work, we plan to study which property patterns of the system [8] can be translated in L_μ^{dsbr}, so as to provide useful information about the possible reductions modulo $\approx_{br}^{ds}$. We also plan to continue experimenting with maximal hiding and on-the-fly reduction by using *weak* forms of divergence-sensitive τ-confluence implemented in a distributed setting [24], i.e., by employing clusters of machines for both LTS reduction and verification.

References

[1] Henrik R. Andersen. Model checking and boolean graphs. *Theoretical Computer Science*, 126(1):3–30, April 1994.

[2] J.C.M. Baeten. A Brief History of Process Algebra. *Theoretical Computer Science*, 335(2-3):131–146, 2005.

[3] Roberto Barbuti, Nicoletta de Francesco, Antonella Santone, and Gigliola Vaglini. Selective mu-calculus and formula-based equivalence of transition systems. *Journal of Computer and System Science*, 59(3):537–556, 1999.

[4] E. M. Clarke, E. A. Emerson, and A. P. Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM Transactions on Programming Languages and Systems*, 8(2):244–263, April 1986.

[5] E.M. Clarke, O. Grumberg, and D.A. Peled. *Model Checking*. The MIT Press, 1999.

[6] Rance Cleaveland and Bernhard Steffen. A linear-time model-checking algorithm for the alternation-free modal mu-calculus. *Formal Methods in System Design*, 2(2):121–147, April 1993.

[7] Nicolas Coste, Hubert Garavel, Holger Hermanns, Frédéric Lang, Radu Mateescu, and Wendelin Serwe. Ten years of performance evaluation for concurrent systems using CADP. In Tiziana Margaria and Bernhard Steffen, editors, *Proceedings of the 4th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation, ISoLA'2010 (Heraclion, Crete), Part II*, volume 6416 of *Lecture Notes in Computer Science*, pages 128–142. Springer Verlag, October 2010.

[8] Matthew B. Dwyer, George S. Avrunin, and James C. Corbett. Patterns in property specifications for finite-state verification. In *Proceedings of the 21st International Conference on Software Engineering ICSE'99 (Los Angeles, CA, USA)*, May 1999.

[9] Alessandro Fantechi, Stefania Gnesi, and Gioia Ristori. From ACTL to Mu-Calculus. In IEI-CNR, editor, *Proceedings of the ERCIM Workshop on Theory and Practice in Verification (Pisa, Italy)*, 1992.

[10] Jean-Claude Fernandez and Laurent Mounier. "on the fly" verification of behavioural equivalences and preorders. In K. G. Larsen and A. Skou, editors, *Proceedings of the 3rd Workshop on Computer-Aided Verification CAV'91 (Aalborg, Denmark)*, volume 575 of *Lecture Notes in Computer Science*, Berlin, 1991. Springer Verlag.

[11] Michael J. Fischer and Richard E. Ladner. Propositional dynamic logic of regular programs. *Journal of Computer and System Sciences*, 18(2):194–211, September 1979.

[12] Hubert Garavel and Frédéric Lang. Svl: a scripting language for compositional verification. In Myungchul Kim, Byoungmoon Chin, Sungwon Kang, and Danhyung Lee,

editors, *Proceedings of the 21st IFIP WG 6.1 International Conference on Formal Techniques for Networked and Distributed Systems FORTE'2001 (Cheju Island, Korea)*, pages 377–392. IFIP, Kluwer Academic Publishers, August 2001. Full version available as INRIA Research Report RR-4223.

[13] Hubert Garavel, Frédéric Lang, Radu Mateescu, and Wendelin Serwe. Cadp 2010: A toolbox for the construction and analysis of distributed processes. In Parosh A. Abdulla and K. Rustan M. Leino, editors, *Proceedings of the 17th International Conference on Tools and Algorithms for the Construction and Analysis of Systems TACAS'2011 (Saarbrücken, Germany)*, volume 6605 of *Lecture Notes in Computer Science*, pages 372–387. Springer Verlag, March 2011.

[14] R.J. van Glabbeek, B. Luttik, and N. Trčka. Branching Bisimilarity with Explicit Divergence. *Fundamenta Informaticae*, 93(4):371–392, 2009.

[15] R.J. van Glabbeek and W.P. Weijland. Branching Time and Abstraction in Bisimulation Semantics. *Journal of the ACM*, 43(3):555–600, 1996.

[16] Jan Friso Groote and M. P. A. Sellink. Confluence for process verification. *Theoretical Computer Science*, 170(1–2):47–81, 1996.

[17] S. C. Kleene. *Introduction to Metamathematics*. North-Holland, 1952.

[18] D. Kozen. Results on the propositional μ-calculus. *Theoretical Computer Science*, 27:333–354, 1983.

[19] Radu Mateescu. On-the-fly state space reductions for weak equivalences. In Tiziana Margaria and Mieke Massink, editors, *Proceedings of the 10th International Workshop on Formal Methods for Industrial Critical Systems FMICS'05 (Lisbon, Portugal)*, pages 80–89. ERCIM, ACM Computer Society Press, September 2005.

[20] Radu Mateescu. Caesar_solve: A generic library for on-the-fly resolution of alternation-free boolean equation systems. *Springer International Journal on Software Tools for Technology Transfer (STTT)*, 8(1):37–56, February 2006. Full version available as INRIA Research Report RR-5948, July 2006.

[21] Radu Mateescu and Mihaela Sighireanu. Efficient on-the-fly model-checking for regular alternation-free mu-calculus. *Science of Computer Programming*, 46(3):255–281, March 2003.

[22] Radu Mateescu and Damien Thivolle. A model checking language for concurrent value-passing systems. In Jorge Cuellar, Tom Maibaum, and Kaisa Sere, editors, *Proceedings of the 15th International Symposium on Formal Methods FM'08 (Turku, Finland)*, number 5014 in Lecture Notes in Computer Science, pages 148–164. Springer Verlag, May 2008.

[23] Radu Mateescu and Anton Wijs. Efficient on-the-fly computation of weak tau-confluence. Research Report RR-7000, INRIA, 2009.

[24] Radu Mateescu and Anton Wijs. Sequential and distributed on-the-fly computation of weak tau-confluence. *Science of Computer Programming*, 2011.

[25] Robin Milner. *Communication and Concurrency.* Prentice-Hall, 1989.

[26] Rocco De Nicola and Frits W. Vaandrager. *Action versus State Based Logics for Transition Systems*, volume 469 of *Lecture Notes in Computer Science*, pages 407–419. Springer Verlag, April 1990.

[27] Gordon Pace, Frédéric Lang, and Radu Mateescu. Calculating τ-confluence compositionally. In Jr Warren A. Hunt and Fabio Somenzi, editors, *Proceedings of the 15th International Conference on Computer Aided Verification CAV'2003 (Boulder, Colorado, USA)*, volume 2725 of *Lecture Notes in Computer Science*, pages 446–459. Springer Verlag, July 2003. Full version available as INRIA Research Report RR-4918.

[28] Colin Stirling. *Modal and Temporal Properties of Processes.* Springer Verlag, 2001.

[29] R. Streett. Propositional dynamic logic of looping and converse. *Information and Control*, 1982.

A Proofs

We provide in this annex the proofs of all lemmas and propositions stated in the main text.

Proof (Proposition 1). We proceed by structural induction on φ.

Case $\varphi ::= \mathsf{false}$.

$$\begin{aligned}
&[\![\mathsf{false}]\!]_{hide_B(M)}\,\rho &&= \text{by definition of } [\![\ ]\!] \\
&\emptyset &&= \text{by definition of } [\![\ ]\!] \\
&[\![\mathsf{false}]\!]_M\,\rho.
\end{aligned}$$

Case $\varphi ::= \neg\varphi_1$. Since $B \subseteq h_A(\neg\varphi_1)$ by hypothesis and $h_A(\neg\varphi_1) = h_A(\varphi_1)$ by Definition 2, it follows that $B \subseteq h_A(\varphi_1)$. Therefore, we can apply the induction hypothesis for φ_1 and B, which yields:

$$\begin{aligned}
&[\![\neg\varphi_1]\!]_{hide_B(M)}\,\rho &&= \text{by definition of } [\![\ ]\!] \\
&S \setminus [\![\varphi_1]\!]_{hide_B(M)}\,\rho &&= \text{by induction hypothesis} \\
&S \setminus [\![\varphi_1]\!]_M\,\rho &&= \text{by definition of } [\![\ ]\!] \\
&[\![\neg\varphi_1]\!]_M\,\rho.
\end{aligned}$$

Case $\varphi ::= \varphi_1 \vee \varphi_2$. Since $B \subseteq h_A(\varphi_1 \vee \varphi_2)$ by hypothesis and $h_A(\varphi_1 \vee \varphi_2) = h_A(\varphi_1) \cap h_A(\varphi_2)$ by Definition 2, it follows that $B \subseteq h_A(\varphi_1)$ and $B \subseteq h_A(\varphi_2)$. Therefore, we can apply the induction hypothesis for φ_1, φ_2, and B, which yields:

$$\begin{array}{ll}
[\![\varphi_1 \vee \varphi_2]\!]_{hide_B(M)}\,\rho & = \text{by definition of } [\![\ ]\!] \\
[\![\varphi_1]\!]_{hide_B(M)}\,\rho \cup [\![\varphi_2]\!]_{hide_B(M)}\,\rho & = \text{by induction hypothesis} \\
[\![\varphi_1]\!]_M\,\rho \cup [\![\varphi_2]\!]_M\,\rho & = \text{by definition of } [\![\ ]\!] \\
[\![\varphi_1 \vee \varphi_2]\!]_M\,\rho. &
\end{array}$$

Case $\varphi ::= X$.

$$\begin{array}{ll}
[\![X]\!]_{hide_B(M)}\,\rho & = \text{by definition of } [\![\ ]\!] \\
\rho(X) & = \text{by definition of } [\![\ ]\!] \\
[\![X]\!]_M\,\rho. &
\end{array}$$

Case $\varphi ::= \mu X.\varphi_1$. Since $B \subseteq h_A(\mu X.\varphi_1)$ by hypothesis and $h_A(\mu X.\varphi_1) = h_A(\varphi_1)$ by Definition 2, it follows that $B \subseteq h_A(\varphi_1)$. Therefore, we can apply the induction hypothesis for φ_1 and B, which yields:

$$\begin{array}{ll}
[\![\mu X.\varphi_1]\!]_{hide_B(M)}\,\rho & = \text{by definition of } [\![\ ]\!] \\
\bigcap\{U \subseteq S \mid [\![\varphi_1]\!]_{hide_B(M)}\,(\rho \oslash [U/X]) \subseteq U\} & = \text{by induction hypothesis} \\
\bigcap\{U \subseteq S \mid [\![\varphi_1]\!]_M\,(\rho \oslash [U/X]) \subseteq U\} & = \text{by definition of } [\![\ ]\!] \\
[\![\mu X.\varphi_1]\!]_M\,\rho. &
\end{array}$$

□

Proof (Proposition 2). We proceed by structural induction on φ.

Case $\varphi ::= \mathsf{false}$. Since $[\![\mathsf{false}]\!]_M\,\rho = \emptyset$ by definition of $[\![\ ]\!]$, none of s_1, s_2 belong to the interpretation of false.

Case $\varphi ::= \neg\varphi_1$. Let $s_1, s_2 \in S$ such that $s_1 \approx_{br}^{ds} s_2$ and assume that $s_1 \in [\![\neg\varphi_1]\!]_M\,\rho$, i.e., $s_1 \notin [\![\varphi_1]\!]_M\,\rho$ by definition of $[\![\ ]\!]$. By the induction hypothesis, this is equivalent to $s_2 \notin [\![\varphi_1]\!]_M\,\rho$, i.e., $s_2 \in [\![\neg\varphi_1]\!]_M\,\rho$.

Case $\varphi ::= \varphi_1 \vee \varphi_2$. Let $s_1, s_2 \in S$ such that $s_1 \approx_{br}^{ds} s_2$ and assume that $s_1 \in [\![\varphi_1 \vee \varphi_2]\!]_M\,\rho$. By definition of $[\![\ ]\!]$, this means $s_1 \in [\![\varphi_1]\!]_M\,\rho \cup [\![\varphi_2]\!]_M\,\rho$, i.e., $s_1 \in [\![\varphi_1]\!]_M\,\rho \vee s_1 \in [\![\varphi_2]\!]_M\,\rho$. By the induction hypothesis, this is equivalent to $s_2 \in [\![\varphi_1]\!]_M\,\rho \vee s_2 \in [\![\varphi_2]\!]_M\,\rho$, i.e., $s_2 \in [\![\varphi_1]\!]_M\,\rho \cup [\![\varphi_2]\!]_M\,\rho$. By definition of $[\![\ ]\!]$, this means $s_2 \in [\![\varphi_1 \vee \varphi_2]\!]_M\,\rho$.

Case $\varphi ::= X$. Let $s_1, s_2 \in S$ such that $s_1 \approx_{br}^{ds} s_2$ and assume that $s_1 \in [\![X]\!]_M\,\rho$, i.e., $s_1 \in \rho(X)$ by definition of $[\![\ ]\!]$. Since $s_1 \approx_{br}^{ds} s_2$ and ρ is $\approx_{br}^{ds}$-closed by hypothesis, this is equivalent to $s_2 \in \rho(X)$, i.e., $s_2 \in [\![X]\!]_M\,\rho$. The converse implication (by considering $s_2 \in [\![X]\!]_M\,\rho$) holds by a symmetric argument.

Case $\varphi ::= \mu X.\varphi_1$. Since we consider finite LTSs, we can use the alternative characterization

of minimal fixed point formulas [17]:

$$[\![\mu X.\varphi_1]\!]_M\,\rho = \bigcup_{k\geq 0} \Phi^k_{M,\rho}(\emptyset), \qquad \Phi^0_{M,\rho}(\emptyset) = \emptyset$$

where $\Phi_{M,\rho} : 2^S \to 2^S$, $\Phi_{M,\rho}(U) = [\![\varphi_1]\!]_M\,(\rho \oslash [U/X])$.

We show first the following statement by induction on k:

$$\forall s_1, s_2 \in S. \forall k \geq 0. s_1 \approx^{ds}_{br} s_2 \Rightarrow (s_1 \in \Phi^k_{M,\rho}(\emptyset) \Leftrightarrow s_2 \in \Phi^k_{M,\rho}(\emptyset)) \tag{5}$$

1. *Base case:* $k = 0$. We have $s_1 \in \Phi^0_{M,\rho}(\emptyset)$, i.e., $s_1 \in \emptyset$, which is equivalent to false. This is equivalent in turn to $s_2 \in \emptyset$, i.e., $s_2 \in \Phi^0_{M,\rho}(\emptyset)$.

2. *Inductive case:* Let $s_1 \in \Phi^{k+1}_{M,\rho}(\emptyset)$, i.e., $s_1 \in \Phi_{M,\rho}(\Phi^k_{M,\rho}(\emptyset))$, which is equivalent to $s_1 \in [\![\varphi_1]\!]_M\,(\rho \oslash [\Phi^k_{M,\rho}(\emptyset)/X])$ by definition of $\Phi_{M,\rho}$. We show that the context $\rho \oslash [\Phi^k_{M,\rho}(\emptyset)/X]$ is $\approx^{ds}_{br}$-closed. Since ρ is $\approx^{ds}_{br}$-closed by hypothesis, it is sufficient to show the closedness of $\rho \oslash [\Phi^k_{M,\rho}(\emptyset)/X]$ for variable X. Let $s'_1, s'_2 \in S$ such that $s'_1 \approx^{ds}_{br} s'_2$ and $s'_1 \in (\rho \oslash [\Phi^k_{M,\rho}(\emptyset)/X])(X)$, i.e., $s'_1 \in \Phi^k_{M,\rho}(\emptyset)$. By the induction hypothesis of (5), this is equivalent to $s'_2 \in \Phi^k_{M,\rho}(\emptyset)$, i.e., $s'_2 \in (\rho \oslash [\Phi^k_{M,\rho}(\emptyset)/X])(X)$.

   Since $\rho \oslash [\Phi^k_{M,\rho}(\emptyset)/X]$ is $\approx^{ds}_{br}$-closed, we can apply the induction hypothesis of the proposition to s_1, φ, and $\rho \oslash [\Phi^k_{M,\rho}(\emptyset)/X]$, and conclude that $s_2 \in [\![\varphi_1]\!]_M\,(\rho \oslash [\Phi^k_{M,\rho}(\emptyset)/X])$, i.e., $s_2 \in \Phi^{k+1}_{M,\rho}(\emptyset)$. The converse implication (by considering $s_2 \in \Phi^{k+1}_{M,\rho}(\emptyset)$) holds by a symmetric argument.

Let $s_1, s_2 \in S$ such that $s_1 \approx^{ds}_{br} s_2$ and assume $s_1 \in [\![\mu X.\varphi_1]\!]_M\,\rho$, i.e., $s_1 \in \bigcup_{k\geq 0} \Phi^k_{M,\rho}(\emptyset)$. This means there exists $k \geq 0$ such that $s_1 \in \Phi^k_{M,\rho}(\emptyset)$ and by applying (5), this is equivalent to $s_2 \in \Phi^k_{M,\rho}(\emptyset)$. This implies $s_2 \in \bigcup_{k\geq 0} \Phi^k_{M,\rho}(\emptyset)$, i.e., $s_2 \in [\![\mu X.\varphi_1]\!]_M\,\rho$. The converse implication (by considering $s_2 \in [\![\mu X.\varphi_1]\!]_M\,\rho$) holds by a symmetric argument.

□

In order to prove Proposition 3, we show first two lemmas.

Lemma 3 *Let $X \in \mathcal{X}$ be a propositional variable and let φ be a state formula of L_μ, which may contain free occurrences of X. Then:*

$$\nu X.(\varphi \wedge [\beta]\,X) = \nu X.\,[\beta^*]\,(\varphi \wedge X)$$

for any regular formula β of PDL.

Proof. The right-hand side of the identity can be rewritten by applying the PDL identity $[\beta^*]\,\psi = \psi \wedge [\beta]\,[\beta^*]\,\psi$ [11]:

$$\nu X.\,[\beta^*]\,(\varphi \wedge X) = \nu X.(\varphi \wedge X \wedge [\beta]\,[\beta^*]\,(\varphi \wedge X))$$

The first occurrence of X can be replaced with true by applying absorption, which yields the identity below:

$$\nu X.\,[\beta^*]\,(\varphi \wedge X) = \nu X.(\varphi \wedge [\beta]\,[\beta^*]\,(\varphi \wedge X)) \tag{6}$$

Consider an LTS $M = \langle S, A, T, s_0\rangle$ and a propositional context ρ. The functionals $\Phi_{M,\rho} : 2^S \to 2^S$ and $\Psi_{M,\rho} : 2^S \to 2^S$ are defined as follows:

$$\Phi_{M,\rho}(U) = [\![\varphi \wedge [\beta]\,X]\!]_M\,(\rho \oslash [U/X])$$
$$\Psi_{M,\rho}(U) = [\![\varphi \wedge [\beta]\,[\beta^*]\,(\varphi \wedge X)]\!]_M\,(\rho \oslash [U/X])$$

Let $\theta, \theta' \subseteq S$ be the minimal fixed points of $\Phi_{M,\rho}, \Psi_{M,\rho}$, respectively. We must show that $\theta = \theta'$. We show first that $\theta \subseteq \theta'$ by using Tarski's theorem, which requires to check that $\theta \subseteq \Psi_{M,\rho}(\theta)$. By definition, θ satisfies the fixed point equation $\theta = [\![\varphi \wedge [\beta]\,X]\!]_M\,(\rho \oslash [\theta/X])$, which implies that $\theta \subseteq [\![\varphi]\!]_M\,(\rho \oslash [\theta/X])$ and $\theta \subseteq [\![[\beta]\,X]\!]_M\,(\rho \oslash [\theta/X])$.

$$\begin{array}{ll}
\Psi_{M,\rho}(\theta) & = \text{by definition of } \Psi_{M,\rho} \text{ and (6)} \\
[\![[\beta^*]\,(\varphi \wedge X)]\!]_M\,(\rho \oslash [\theta/X]) & = \text{by introducing } Y \\
[\![[\beta^*]\,Y]\!]_M\,(\rho \oslash [\theta/X, [\![\varphi \wedge X]\!]_M\,(\rho \oslash [\theta/X])/Y]) & = \text{by definition of } [\![\ ]\!] \\
[\![[\beta^*]\,Y]\!]_M\,(\rho \oslash [\theta/X, [\![\varphi]\!]_M\,(\rho \oslash [\theta/X]) \cap [\![X]\!]_M\,(\rho \oslash [\theta/X])/Y]) & = \text{by def. of } [\![\ ]\!] \\
[\![[\beta^*]\,Y]\!]_M\,(\rho \oslash [\theta/X, [\![\varphi]\!]_M\,(\rho \oslash [\theta/X]) \cap \theta/Y]) & = \text{by } \theta \subseteq [\![\varphi]\!]_M\,(\rho \oslash [\theta/X]) \\
[\![[\beta^*]\,Y]\!]_M\,(\rho \oslash [\theta/X, \theta/Y]) & = \text{by replacing } Y \text{ with } X \\
[\![[\beta^*]\,X]\!]_M\,(\rho \oslash [\theta/X]). &
\end{array}$$

It remains to show that $\theta \subseteq [\![[\beta^*]\,X]\!]_M\,(\rho \oslash [\theta/X])$. Let $\Gamma_{M,\rho} : 2^S \to 2^S$ be the functional associated to the formula $[\beta^*]\,X$, defined as follows: $\Gamma_{M,\rho}(U) = [\![X \wedge [\beta]\,Y]\!]_M\,(\rho \oslash [U/Y])$. The semantics of this formula when X is replaced by θ is characterized iteratively as follows [17]:

$$[\![[\beta^*]\,X]\!]_M\,(\rho \oslash [\theta/X]) = [\![\nu Y.(X \wedge [\beta]\,Y)]\!]_M\,(\rho \oslash [\theta/X]) = \bigcap_{k \geq 0} \Gamma^k_{M,\rho \oslash [\theta/X]}(S)$$

To show the desired inclusion, we prove that $\theta \subseteq \Gamma^k_{M,\rho \oslash [\theta/X]}(S)$ by induction on k.

1. *Base case:* $\theta \subseteq S = \Gamma^0_{M,\rho \oslash [\theta/X]}(S)$.
2. *Inductive case:*

$$\begin{array}{ll}
\Gamma^{k+1}_{M,\rho \oslash [\theta/X]}(S) & = \text{by definition of } \Gamma_{M,\rho} \\
\Gamma_{M,\rho \oslash [\theta/X]}(\Gamma^k_{M,\rho \oslash [\theta/X]}(S)) & = \text{by definition of } \Gamma_{M,\rho} \\
[\![X \wedge [\beta]\,Y]\!]_M\,(\rho \oslash [\theta/X, \Gamma^k_{M,\rho \oslash [\theta/X]}(S)/Y] & \supseteq \text{by induction hypothesis} \\
[\![X \wedge [\beta]\,Y]\!]_M\,(\rho \oslash [\theta/X, \theta/Y]) & = \text{by definition of } [\![\ ]\!] \\
\theta \cap [\![[\beta]\,Y]\!]_M\,(\rho \oslash [\theta/Y]) & = \text{by } \theta \subseteq [\![[\beta]\,Y]\!]_M\,(\rho \oslash [\theta/Y]) \\
\theta. &
\end{array}$$

To show that $\theta' \subseteq \theta$, we check that θ' satisfies the fixed point equation of $\Phi_{M,\rho}$:

$$
\begin{array}{ll}
\theta' & = \text{by fixed point def.} \\
\Psi_{M,\rho}(\theta') & = \text{by def. of } \Psi_{M,\rho} \\
[\![\varphi \wedge [\beta]\,[\beta^*]\,(\varphi \wedge X)]\!]_M\,(\rho \oslash [\theta'/X]) & = \text{by introducing } Y \\
[\![\varphi \wedge [\beta]\,Y]\!]_M\,(\rho \oslash [\theta'/X, [\![[\beta^*]\,(\varphi \wedge X)]\!]_M\,(\rho \oslash [\theta'/X])/Y]) & = \text{by (6)} \\
[\![\varphi \wedge [\beta]\,Y]\!]_M\,(\rho \oslash [\theta'/X, \theta'/Y]) & = \text{by removing } Y \\
[\![\varphi \wedge [\beta]\,X]\!]_M\,(\rho \oslash [\theta'/X]) & = \text{by def. of } \Phi_{M,\rho} \\
\Phi_{M,\rho}(\theta'). &
\end{array}
$$

□

Lemma 4 *Let β_1, β_2 be regular formulas of* PDL *such that they denote one-step sequences (i.e., β_1 and β_2 are satisfied by transition sequences containing only one transition) and they are disjoint (i.e., no transition can satisfy both β_1 and β_2). Then:*

$$\langle \beta_1 | \beta_2 \rangle\, @ = \langle (\beta_1^*.\beta_2)^* \rangle\, \langle \beta_1 \rangle\, @ \vee \langle \beta_1^*.\beta_2 \rangle\, @.$$

Proof. Implication "$\Leftarrow$". Both disjuncts in the right-hand side of the equality are included in the left-hand side term because the ω-regular languages $(\beta_1^*.\beta_2)^*.\beta_1^\omega$ and $(\beta_1^*.\beta_2)^\omega$ are both included in $(\beta_1|\beta_2)^\omega$, which consists of all infinite sequences made from transitions satisfying β_1 or β_2.

Implication "$\Rightarrow$". By expanding the infinite looping operators in terms of maximal fixed points, this implication becomes:

$$\nu X.\, \langle \beta_1 | \beta_2 \rangle\, X \Rightarrow \langle (\beta_1^*.\beta_2)^* \rangle\, \nu Y.\, \langle \beta_1 \rangle\, Y \vee \nu Z.\, \langle \beta_1^*.\beta_2 \rangle\, Z$$

Let $M = \langle S, A, T, s_0 \rangle$ be an LTS. The functionals $\Phi_M, \Psi_M, \Gamma_M : 2^S \to 2^S$ associated to the three maximal fixed point operators are defined as follows:

$$
\begin{array}{rcl}
\Phi_M(U) & = & [\![\langle \beta_1 | \beta_2 \rangle\, X]\!]_M\, [U/X] \\
\Psi_M(U) & = & [\![\langle \beta_1 \rangle\, Y]\!]_M\, [U/Y] \\
\Gamma_M(U) & = & [\![\langle \beta_1^*.\beta_2 \rangle\, Z]\!]_M\, [U/Z]
\end{array}
$$

Using the interpretation of fixed point formulas, the implication to show is equivalent to the inclusion below:

$$\nu \Phi_M \subseteq [\![\langle (\beta_1^*.\beta_2)^* \rangle\, Y]\!]_M\, [\nu \Psi_M / Y] \cup \nu \Gamma_M$$

Since we consider finite LTSs, we can rewrite this inclusion as follows using the alternative characterization of maximal fixed point formulas [17]:

$$\bigcap_{k \geq 0} \Phi_M^k(S) \subseteq [\![\langle (\beta_1^*.\beta_2)^* \rangle\, Y]\!]_M\, [\bigcap_{k \geq 0} \Psi_M^k(S)/Y] \cup \bigcap_{k \geq 0} \Gamma_M^k(S)$$

We show this inclusion by reasoning in two complementary cases, depending on the fact that the relation below holds or not:

$$\forall k \geq 0. \exists n \geq k. \Phi_M^n(S) \subseteq \Gamma_M^k(S) \tag{7}$$

1. *Case when (7) holds.* The limit of the series $\Phi_M^k(S)$ is calculated as follows:

$$\begin{array}{ll} \bigcap_{k\geq 0} \Phi_M^k(S) & = \\ \bigcap_{k\geq 0} \bigcap_{n\geq k} \Phi_M^n(S) & \subseteq \text{by (7)} \\ \bigcap_{k\geq 0} \Gamma_M^k(S) & \end{array}$$

and therefore the desired inclusion holds.

2. *Case when (7) fails.* By using the definition of Γ_M, the negation of (7) can be written as follows:

$$\exists k \geq 0. \forall n \geq k. \Phi_M^n(S) \not\subseteq [\![\left\langle (\beta_1^*.\beta_2)^k \right\rangle Y]\!]_M [S/Y]$$

where the notation β^k stands for the concatenation $\beta \ldots \beta$ k times. From the definition of Φ_M, it follows that $\Phi_M^n(S) = [\![\langle (\beta_1 | \beta_2)^n \rangle X]\!]_M [S/X]$, i.e., $\Phi_M^n(S)$ denotes the states having an outgoing transition sequence of length n whose transitions satisfy β_1 or β_2. Let $k \geq 0$ satisfying the negation of (7) above. This means that for all $n \geq k$, the outgoing transition sequence cannot contain more than k transitions satisfying β_2, and therefore there exists $0 < j \leq k$ such that the prefix of the sequence contains at most $k - j$ transitions satisfying β_2:

$$\Phi_M^n(S) = [\![\langle (\beta_1^*.\beta_2)^{k-j}.\beta_1^{n-k+j} \rangle Y]\!]_M [S/Y]$$

By expanding the last concatenation operator in the modality above and introducing the auxiliary variable X, this implies the following inclusion:

$$\Phi_M^n(S) \subseteq [\![\langle (\beta_1^*.\beta_2)^{k-j} \rangle X]\!]_M [[\![\langle \beta_1^{n-k+j} \rangle Y]\!]_M [S/Y]/X]$$

which can in turn be rewritten using the definition of Ψ_M and the fact that the series $\Psi_M^k(S)$ is decreasing:

$$\Phi_M^n(S) \subseteq [\![\langle (\beta_1^*.\beta_2)^* \rangle X]\!]_M [\Psi_M^{n-k}(S)/X] \tag{8}$$

Now we can calculate the limit of the series $\Phi_M^n(S)$ as follows:

$$\begin{array}{ll} \bigcap_{n\geq 0} \Phi_M^n(S) & = \text{since } \Phi_M^n(S) \text{ is decreasing} \\ \bigcap_{n\geq k} \Phi_M^n(S) & \subseteq \text{by (8)} \\ \bigcap_{n\geq k} [\![\langle (\beta_1^*.\beta_2)^* \rangle X]\!]_M [\Psi_M^{n-k}(S)/X] & = \text{by replacing } n-k \text{ by } l \\ \bigcap_{l\geq 0} [\![\langle (\beta_1^*.\beta_2)^* \rangle X]\!]_M [\Psi_M^{l}(S)/X] & = \text{since } \Psi_M^l(S) \text{ is decreasing to } \nu\Psi_M \\ [\![\langle (\beta_1^*.\beta_2)^* \rangle X]\!]_M [\nu\Psi_M/X] & \end{array}$$

and therefore the desired inclusion holds.

□

Proof (Proposition 3). Starting from the L_μ^{dsbr} formulations of the ACTL\X temporal operators stated in the proposition, we expand the weak modalities to obtain plain L_μ formulas, and then we show that these formulas are equivalent to the L_μ formulas given in Table 1.

Operator $\mathsf{E}[\varphi_1 {}_\alpha\mathsf{U} \varphi_2]$.

$$
\begin{array}{ll}
\mathsf{E}[\varphi_{1_\alpha}\mathsf{U}\varphi_2] & = \text{by hypothesis} \\
\langle(\varphi_1?.\alpha \vee \tau)^*\rangle\, \varphi_2 & = \text{by expansion of the } * \text{ operator} \\
\mu X.(\varphi_2 \vee \langle\varphi_1?.\alpha \vee \tau\rangle\, X) & = \text{by expansion of the . operator} \\
\mu X.(\varphi_2 \vee \langle\varphi_1?\rangle\, \langle\alpha \vee \tau\rangle\, X) & = \text{by expansion of the ? operator} \\
\mu X.(\varphi_2 \vee (\varphi_1 \wedge \langle\alpha \vee \tau\rangle\, X)). &
\end{array}
$$

Operator $\mathsf{E}[\varphi_{1_{\alpha_1}}\mathsf{U}_{\alpha_2}\varphi_2]$.

$$
\begin{array}{ll}
\mathsf{E}[\varphi_{1_{\alpha_1}}\mathsf{U}_{\alpha_2}\varphi_2] & = \text{by hypothesis} \\
\langle(\varphi_1?.\alpha_1 \vee \tau)^*\rangle\, (\varphi_1 \wedge \langle\alpha_2\rangle\, \varphi_2) & = \text{by expansion of the } * \text{ operator} \\
\mu X.((\varphi_1 \wedge \langle\alpha_2\rangle\, \varphi_2) \vee \langle\varphi_1?.\alpha_1 \vee \tau\rangle\, X) & = \text{by expansion of the . operator} \\
\mu X.((\varphi_1 \wedge \langle\alpha_2\rangle\, \varphi_2) \vee \langle\varphi_1?\rangle\, \langle\alpha_1 \vee \tau\rangle\, X) & = \text{by expansion of the ? operator} \\
\mu X.((\varphi_1 \wedge \langle\alpha_2\rangle\, \varphi_2) \vee (\varphi_1 \wedge \langle\alpha_1 \vee \tau\rangle\, X)) & = \text{by propositional calculus} \\
\mu X.(\varphi_1 \wedge (\langle\alpha_2\rangle\, \varphi_2 \vee \langle\alpha_1 \vee \tau\rangle\, X)). &
\end{array}
$$

Operator $\mathsf{A}[\varphi_{1_\alpha}\mathsf{U}\varphi_2]$.

$$
\begin{array}{ll}
\mathsf{A}[\varphi_{1_\alpha}\mathsf{U}\varphi_2] & = \text{ by hypothesis} \\
[(\neg\varphi_2?.\alpha \vee \tau)^*]\, (\varphi_2 \vee (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha \vee \tau)]\, \mathsf{false})) \wedge & \\
\quad [\neg\varphi_2?.\alpha \vee \tau] \dashv & = \text{ by expansion of the * operator} \\
\nu X.((\varphi_2 \vee (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha \vee \tau)]\, \mathsf{false})) \wedge & \\
\quad [\neg\varphi_2?.\alpha \vee \tau]\, X) \wedge [\neg\varphi_2?.\alpha \vee \tau] \dashv & = \text{ by expansion of the . operator} \\
\nu X.((\varphi_2 \vee (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha \vee \tau)]\, \mathsf{false})) \wedge & \\
\quad [\neg\varphi_2?]\, [\alpha \vee \tau]\, X) \wedge [\neg\varphi_2?.\alpha \vee \tau] \dashv & = \text{ by expansion of the ? operator} \\
\nu X.((\varphi_2 \vee (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha \vee \tau)]\, \mathsf{false})) \wedge & \\
\quad (\varphi_2 \vee [\alpha \vee \tau]\, X)) \wedge [\neg\varphi_2?.\alpha \vee \tau] \dashv & = \text{ by propositional calculus} \\
\nu X.(\varphi_2 \vee (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha \vee \tau)]\, \mathsf{false} \wedge [\alpha \vee \tau]\, X)) \wedge & \\
\quad [\neg\varphi_2?.\alpha \vee \tau] \dashv & = \text{ by expansion of the [ ] } \dashv \text{ operator} \\
\nu X.(\varphi_2 \vee (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha \vee \tau)]\, \mathsf{false} \wedge [\alpha \vee \tau]\, X)) \wedge & \\
\mu X.\, [\neg\varphi_2?.\alpha \vee \tau]\, X & = \text{ by expansion of the . operator} \\
\nu X.(\varphi_2 \vee (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha \vee \tau)]\, \mathsf{false} \wedge [\alpha \vee \tau]\, X)) \wedge & \\
\quad \mu X.\, [\neg\varphi_2?]\, [\alpha \vee \tau]\, X & = \text{ by expansion of the ? operator} \\
\nu X.(\varphi_2 \vee (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha \vee \tau)]\, \mathsf{false} \wedge [\alpha \vee \tau]\, X)) \wedge & \\
\quad \mu X.(\varphi_2 \vee [\alpha \vee \tau]\, X). &
\end{array}
$$

To show the equivalence between the last formula above and the translation of $\mathsf{A}[\varphi_{1_\alpha}\mathsf{U}\varphi_2]$ in L_μ given in Table 1, it remains to show the following equality:

$$
\begin{array}{l}
\mu X.(\varphi_2 \vee (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha \vee \tau)]\, \mathsf{false} \wedge [\alpha \vee \tau]\, X)) = \\
\nu X.(\varphi_2 \vee (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha \vee \tau)]\, \mathsf{false} \wedge [\alpha \vee \tau]\, X)) \wedge \mu X.(\varphi_2 \vee [\alpha \vee \tau]\, X)
\end{array}
$$

The "$\Rightarrow$" implication follows immediately by monotonicity. For the converse implication, we consider an LTS $M = \langle S, A, T, s_0\rangle$ and we show the following inequality between the interpretations on M of the formulas in the left- and right-hand sides (note that φ_1, φ_2 are closed and therefore there is no need for a propositional context ρ):

$[\![\nu X.(\varphi_2 \vee (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha \vee \tau)]\,\mathsf{false} \wedge [\alpha \vee \tau]\,X))]\!]_M \cap$
$[\![\mu X.(\varphi_2 \vee [\alpha \vee \tau]\,X)]\!]_M \subseteq$
$[\![\mu X.(\varphi_2 \vee (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha \vee \tau)]\,\mathsf{false} \wedge [\alpha \vee \tau]\,X))]\!]_M$

Let $\Phi_M, \Psi_M : 2^S \to 2^S$ the functionals defined below:

$\Phi_M(U) = [\![\varphi_2 \vee (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha \vee \tau)]\,\mathsf{false} \wedge [\alpha \vee \tau]\,X)]\!]_M\,[U/X]$
$\Psi_M(U) = [\![\varphi_2 \vee [\alpha \vee \tau]\,X]\!]_M\,[U/X]$

Using the iterative characterization of fixed point operators [17], the inequality above can be reformulated in terms of these functionals as follows:

$\bigcap_{n\geq 0} \Phi_M^n(S) \cap \bigcup_{n\geq 0} \Psi_M^n(\emptyset) \subseteq \bigcup_{n\geq 0} \Phi_M^n(\emptyset)$

or, equivalently:

$\bigcup_{k\geq 0} \left(\left(\bigcap_{n\geq 0} \Phi_M^n(S) \right) \cap \Psi_M^k(\emptyset) \right) \subseteq \bigcup_{k\geq 0} \Phi_M^k(\emptyset)$

To prove the last inequality, we first show the relation below by induction on k:

$$\forall k \geq 0. \left(\bigcap_{n\geq 0} \Phi_M^n(S) \right) \cap \Psi_M^k(\emptyset) \subseteq \Phi_M^k(\emptyset) \tag{9}$$

1. *Base case:* $(\bigcap_{n\geq 0} \Phi_M^n(S)) \cap \Psi_M^0(\emptyset) = (\bigcap_{n\geq 0} \Phi_M^n(S)) \cap \emptyset = \emptyset \subseteq \Phi_M^0(\emptyset)$.
2. *Inductive case:*

$$\begin{array}{lll}
\Phi_M^{k+1}(\emptyset) & = & \\
\Phi_M(\Phi_M^k(\emptyset)) & \supseteq & \text{by induction hypothesis} \\
\Phi_M((\bigcap_{n\geq 0} \Phi_M^n(S)) \cap \Psi_M^k(\emptyset)) & = & \text{by definition of } \Phi_M \\
[\![\varphi_2 \vee (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha \vee \tau)]\,\mathsf{false} \wedge [\alpha \vee \tau]\,X)]\!]_M & & \\
\quad [(\bigcap_{n\geq 0} \Phi_M^n(S)) \cap \Psi_M^k(\emptyset)/X] & = & \text{by introducing } Y \\
[\![\varphi_2 \vee (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha \vee \tau)]\,\mathsf{false} \wedge [\alpha \vee \tau]\,(X \wedge Y))]\!]_M & & \\
\quad [\bigcap_{n\geq 0} \Phi_M^n(S)/X, \Psi_M^k(\emptyset)/Y] & = & \text{by modal calculus} \\
[\![(\varphi_2 \vee (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha \vee \tau)]\,\mathsf{false} \wedge [\alpha \vee \tau]\,X)) \wedge (\varphi_2 \vee [\alpha \vee \tau]\,Y)]\!]_M & & \\
\quad [\bigcap_{n\geq 0} \Phi_M^n(S)/X, \Psi_M^k(\emptyset)/Y] & = & \text{by definition of } [\![\ ]\!] \\
[\![\varphi_2 \vee (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha \vee \tau)]\,\mathsf{false} \wedge [\alpha \vee \tau]\,X)]\!]_M & & \\
\quad [\bigcap_{n\geq 0} \Phi_M^n(S)/X] \cap & & \\
[\![\varphi_2 \vee [\alpha \vee \tau]\,Y]\!]_M\,[\Psi_M^k(\emptyset)/Y] & = & \text{by definition of } \Phi_M, \Psi_M \\
\Phi_M(\bigcap_{n\geq 0} \Phi_M^n(S)) \cap \Psi_M(\Psi_M^k(\emptyset)) & = & \text{by definition of } \nu\Phi_M \\
(\bigcap_{n\geq 0} \Phi_M^n(S)) \cap \Psi_M^{k+1}(\emptyset). & &
\end{array}$$

By applying union for all $k \geq 0$ on the left- and right-hand sides of (9), we obtain the desired inequality.

Operator $\mathsf{A}[\varphi_{1_{\alpha_1}} \mathsf{U}_{\alpha_2} \varphi_2]$.

$$
\begin{array}{ll}
\mathsf{A}[\varphi_{1_{\alpha_1}} \mathsf{U}_{\alpha_2} \varphi_2] & = \text{ by hypothesis} \\
\nu X.\, [(\neg\alpha_2)^*]\, (\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha_1 \vee \alpha_2 \vee \tau)]\, \mathsf{false} \wedge [\alpha_2 \wedge \neg\alpha_1]\, \varphi_2 \wedge & \\
\qquad [\alpha_1 \wedge \alpha_2]\, (\varphi_2 \vee X) \wedge X) \wedge & \\
\nu X.([\neg\alpha_2]\, \dashv \wedge [(\neg\alpha_2)^*]\, [\alpha_1 \wedge \alpha_2]\, (\varphi_2 \vee X)) \wedge & \\
\mu X.\, [(\neg\alpha_2)^*]\, [\alpha_1 \wedge \alpha_2]\, (\varphi_2 \vee X) & = \text{ by Lemma 3} \\
\nu X.(\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha_1 \vee \alpha_2 \vee \tau)]\, \mathsf{false} \wedge [\alpha_2 \wedge \neg\alpha_1]\, \varphi_2 \wedge & \\
\qquad [\alpha_1 \wedge \alpha_2]\, (\varphi_2 \vee X) \wedge [\neg\alpha_2]\, X) \wedge & \\
\nu X.([\neg\alpha_2]\, \dashv \wedge [(\neg\alpha_2)^*]\, [\alpha_1 \wedge \alpha_2]\, (\varphi_2 \vee X)) \wedge & \\
\mu X.\, [(\neg\alpha_2)^*]\, [\alpha_1 \wedge \alpha_2]\, (\varphi_2 \vee X) & = \text{ by PDL semantics} \\
\nu X.(\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha_1 \vee \alpha_2 \vee \tau)]\, \mathsf{false} \wedge [\alpha_2 \wedge \neg\alpha_1]\, \varphi_2 \wedge & \\
\qquad [\alpha_1 \wedge \alpha_2]\, (\varphi_2 \vee X) \wedge [\neg\alpha_2]\, X) \wedge & \\
[((\neg\alpha_2)^*.(\alpha_1 \wedge \alpha_2.\neg\varphi_2?))^*]\, [\neg\alpha_2]\, \dashv\ \wedge & \\
[(\neg\alpha_2)^*.(\alpha_1 \wedge \alpha_2.\neg\varphi_2?)]\, \dashv & = \text{ by negation of Lemma 4} \\
\nu X.(\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha_1 \vee \alpha_2 \vee \tau)]\, \mathsf{false} \wedge [\alpha_2 \wedge \neg\alpha_1]\, \varphi_2 \wedge & \\
\qquad [\alpha_1 \wedge \alpha_2]\, (\varphi_2 \vee X) \wedge [\neg\alpha_2]\, X) \wedge & \\
[\neg\alpha_2 | (\alpha_1 \wedge \alpha_2.\neg\varphi_2?)]\, \dashv & = \text{ by PDL semantics} \\
\nu X.(\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha_1 \vee \alpha_2 \vee \tau)]\, \mathsf{false} \wedge [\alpha_2 \wedge \neg\alpha_1]\, \varphi_2 \wedge & \\
\qquad [\alpha_1 \wedge \alpha_2]\, (\varphi_2 \vee X) \wedge [\neg\alpha_2]\, X) \wedge & \\
\mu X.\, [\neg\alpha_2 | (\alpha_1 \wedge \alpha_2.\neg\varphi_2?)]\, X & = \text{ by PDL semantics} \\
\nu X.(\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha_1 \vee \alpha_2 \vee \tau)]\, \mathsf{false} \wedge [\alpha_2 \wedge \neg\alpha_1]\, \varphi_2 \wedge & \\
\qquad [\alpha_1 \wedge \alpha_2]\, (\varphi_2 \vee X) \wedge [\neg\alpha_2]\, X) \wedge & \\
\mu X.([\neg\alpha_2]\, X \wedge [\alpha_1 \wedge \alpha_2.\neg\varphi_2?]\, X) & = \text{ by PDL semantics} \\
\nu X.(\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha_1 \vee \alpha_2 \vee \tau)]\, \mathsf{false} \wedge [\alpha_2 \wedge \neg\alpha_1]\, \varphi_2 \wedge & \\
\qquad [\alpha_1 \wedge \alpha_2]\, (\varphi_2 \vee X) \wedge [\neg\alpha_2]\, X) \wedge & \\
\mu X.([\alpha_1 \wedge \alpha_2]\, (\varphi_2 \vee X) \wedge [\neg\alpha_2]\, X). &
\end{array}
$$

To show the equivalence between the last formula above and the translation of $\mathsf{A}[\varphi_{1_{\alpha_1}} \mathsf{U}_{\alpha_2} \varphi_2]$ in L_μ given in Table 1, it remains to show the following equality:

$$
\begin{array}{l}
\mu X.(\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha_1 \vee \alpha_2 \vee \tau)]\, \mathsf{false} \wedge [\alpha_2 \wedge \neg\alpha_1]\, \varphi_2 \wedge \\
\qquad [\alpha_1 \wedge \alpha_2]\, (\varphi_2 \vee X) \wedge [\neg\alpha_2]\, X) = \\
\nu X.(\varphi_1 \wedge \neg \mathit{deadlock} \wedge [\neg(\alpha_1 \vee \alpha_2 \vee \tau)]\, \mathsf{false} \wedge [\alpha_2 \wedge \neg\alpha_1]\, \varphi_2 \wedge \\
\qquad [\alpha_1 \wedge \alpha_2]\, (\varphi_2 \vee X) \wedge [\neg\alpha_2]\, X) \wedge \mu X.([\alpha_1 \wedge \alpha_2]\, (\varphi_2 \vee X) \wedge [\neg\alpha_2]\, X)
\end{array}
$$

The proof of this last equality is very similar to the proof of the corresponding equality for the $\mathsf{A}[\varphi_{1_\alpha} \mathsf{U} \varphi_2]$ operator above, and is omitted here. □

Centre de recherche INRIA Grenoble – Rhône-Alpes
Inovallée, 655, avenue de l'Europe, Montbonnot - 38334 Saint Ismier Cedex (France)

Centre de recherche INRIA Futurs : Parc Club Orsay Université - ZAC des Vignes, 4, rue Jacques Monod - Bât. G - 91893 Orsay Cedex (France)
Centre de recherche INRIA Nancy – Grand Est : 615, rue du Jardin Botanique - 54600 Villers-lès-Nancy (France)
Centre de recherche INRIA Rennes – Bretagne Atlantique : Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)
Centre de recherche INRIA Sophia Antipolis – Méditerranée : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
http://www.inria.fr
ISSN 0249-6399