Compositional Verification of Concurrent Systems by Combining Bisimulations

Frédéric Lang[1], Radu Mateescu[1], and Franco Mazzanti[2]

[1] Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP**, LIG, 38000 Grenoble, France
[2] ISTI-CNR, Pisa, Italy

Abstract. One approach to verify a property expressed as a modal μ-calculus formula on a system with several concurrent processes is to build the underlying state space compositionally (i.e., by minimizing and recomposing the state spaces of individual processes, keeping visible only the relevant actions occurring in the formula), and check the formula on the resulting state space. It was shown previously that, when checking the formulas of the L_μ^{dsbr} fragment of μ-calculus (consisting of weak modalities only), individual processes can be minimized modulo divergence-preserving branching (divbranching) bisimilation. In this paper, we refine this approach to handle formulas containing both strong and weak modalities, so as to enable a combined use of strong or divbranching bisimulation minimization on concurrent processes depending whether they contain or not the actions occurring in the strong modalities of the formula. We extend L_μ^{dsbr} with strong modalities and show that the combined minimization approach preserves the truth value of formulas of the extended fragment. We implemented this approach on top of the CADP verification toolbox and demonstrated how it improves the capabilities of compositional verification on realistic examples of concurrent systems.

1 Introduction

We consider the problem of verifying a temporal logic property on a concurrent system $P_1 \mid\mid ... \mid\mid P_n$ consisting of n processes composed in parallel. We work in the action-based setting, the property being specified as a formula φ of the modal μ-calculus (L_μ) [18] and the processes P_i being described in a language with process algebraic flavour. A well-known problem is the state-space explosion that happens when the system state space exceeds the available computer memory.

Compositional verification is a set of techniques and tools that have proven efficient to palliate state-space explosion in many situations [11]. These techniques may be either independent of the property, i.e., focus only on the construction of the system state space, such as compositional state space construction [22, 29, 32, 31, 14, 33, 19]. Alternatively, they may depend on the property, e.g., verification of the property on the full system is decomposed in the verification of properties on (expectedly smaller) sub-systems, such as in compositional reachability analysis [36, 4], assume-guarantee reasoning [28], or partial model checking [1].

** Institute of Engineering Univ. Grenoble Alpes

Nevertheless, the frontier between property-independent and property-dependent techniques is loose. In compositional state space construction, to be able to reduce the system size, a set of actions is selected and a suitable equivalence relation (e.g., strong bisimulation, branching bisimulation, or divergence-preserving branching bisimulation — divbranching for short) is chosen, restricting the set of properties preserved after hiding the selected actions and reducing the system w.r.t. the selected relation. Therefore, there is still a dependency between the state space construction and the set of properties that can be verified. Given a formula φ of L_μ to be verified on the system, Mateescu & Wijs [24] have pushed this idea and shown how to extract a maximal hiding set of actions and an equivalence relation (either strong or divbranching bisimulation) automatically from φ, thus inviting the compositional state space construction technique to the table of property-dependent reductions. To select the equivalence relation from the formula, they have identified an L_μ fragment named L_μ^{dsbr}, which is adequate with divbranching bisimulation [24]. This fragment consists of L_μ restricted to *weak* modalities, which match actions preceded by arbitrary sequences of hidden actions, as opposed to traditional strong modalities $\langle\alpha\rangle\,\varphi_0$ and $[\alpha]\,\varphi_0$, which match only a single action satisfying α. If φ belongs to L_μ^{dsbr}, then the system can be reduced for divbranching bisimulation; otherwise, it can be reduced for strong bisimulation, the weakest equivalence relation preserving full L_μ.

In this paper, we revisit and refine this approach to accommodate L_μ formulas containing both strong and weak modalities. To do so, we define a logic named $L_\mu^{strong}(A_s)$, which extends L_μ^{dsbr} with strong modalities matching only the actions belonging to a given set A_s of *strong* actions. The set A_s induces a partition of the processes $P_1 \mid\mid ... \mid\mid P_n$ into those containing at least one strong action, and those that do not. We show that a formula φ of $L_\mu^{strong}(A_s)$ is still preserved if the processes containing strong actions are reduced modulo strong bisimulation and the other ones modulo divbranching bisimulation. We also provide guidelines for extracting the set A_s from particular L_μ formulas encoding the operators of widely-used temporal logics, such as CTL [5], ACTL [26], PDL [9], and PDL-Δ [30]. This combined use of bisimulations to reduce different parts of the same system makes possible a fine-tuning of the compositional state space construction by going smoothly from strong bisimulation (when all modalities are strong) to divbranching bisimulation (when A_s is empty, as in the previous approach based on L_μ^{dsbr}). We implemented this approach on top of the CADP verification toolbox [12], and demonstrated how it improves the capabilities of compositional verification on two realistic case studies, namely the TFTP plane-ground communication protocol specified in [13] and the parallel CTL benchmark of the RERS'2018 challenge.

The paper is organized as follows. Section 2 recalls some definitions. Section 3 defines $L_\mu^{strong}(A_s)$ and proves the main result of its adequacy with the combined use of strong and divbranching bisimulations. Section 4 presents the experimental results obtained on the two case studies. Finally, Section 5 contains concluding remarks and directions of future work. Formal proofs and code of case studies are available at `http://doi.org/10.5281/zenodo.2634148`.

2 Background

2.1 LTS compositions and reductions

We consider systems whose behavioural semantics can be represented using an LTS (*Labelled Transition System*).

Definition 1 (LTS). *Let $\mathcal{A}$ denote an infinite set of actions, including the invisible action τ, which denotes internal behaviour. All other actions are called visible. An LTS is a tuple $(\Sigma, A, \longrightarrow, p_{init})$, where Σ is a set of states, $A \subseteq \mathcal{A}$ is a set of actions, $\longrightarrow \subseteq \Sigma \times A \times \Sigma$ is the (labelled) transition relation, and $p_{init} \in \Sigma$ is the initial state. We write $p \xrightarrow{a} p'$ if $(p, a, p') \in \longrightarrow$ and $p \xrightarrow{\tau^*} p'$ if there is a (possibly empty) sequence of τ-transitions from p to p', i.e., states $p_0, \ldots, p_n$ $(n \geq 0)$ such that $p = p_0$, $p' = p_n$, and $p_i \xrightarrow{\tau} p_{i+1}$ for $i = 0, \ldots, n-1$.*

LTS can be composed in parallel and their actions can be abstracted away using the parallel composition and hiding operators defined below. Prior to hiding, an action mapping operator is also introduced for the generality of the approach.

Definition 2 (Parallel composition of LTS). *Let $P = (\Sigma_P, A_P, \longrightarrow_P, p_{init})$, $Q = (\Sigma_Q, A_Q, \longrightarrow_Q, q_{init})$, and $A_{sync} \subseteq \mathcal{A} \setminus \{\tau\}$. The parallel composition of P and Q with synchronization on A_{sync}, "$P \,|[A_{sync}]|\, Q$", is defined as $(\Sigma_P \times \Sigma_Q, A_P \cup A_Q, \longrightarrow, (p_{init}, q_{init}))$, where $(p, q) \xrightarrow{a} (p', q')$ if and only if either (1) $p \xrightarrow{a} p'$, $q' = q$, and $a \notin A_{sync}$, or (2) $p' = p$, $q \xrightarrow{a} q'$, and $a \notin A_{sync}$, or (3) $p \xrightarrow{a} p'$, $q \xrightarrow{a} q'$, and $a \in A_{sync}$.*

Definition 3 (Action mapping). *Let $P = (\Sigma_P, A_P, \longrightarrow_P, p_{init})$ and a total function $\rho : A_P \rightarrow 2^{\mathcal{A}}$. We write $\rho(A_P)$ for the image of ρ, defined by $\bigcup_{a \in A_P} \rho(a)$. We write $\rho(P)$ for the action mapping ρ applied to P, defined as the LTS $(\Sigma_P, \rho(A_P), \longrightarrow'_P, p_{init})$ where $\longrightarrow'_P = \{(p, a', p') \mid p \xrightarrow{a}_P p' \wedge a' \in \rho(a)\}$. An action mapping ρ is admissible if $\tau \in A_P \implies \rho(\tau) = \{\tau\}$.*

Action mapping enables a single action a to be mapped onto the empty set of actions, onto a single action a', or onto more than one actions $a'_0, \ldots, a'_{n+1}$ $(n \geq 0)$. In the first case, every transition labelled by a is removed. In the second case, a is renamed into a'. In the third case, every transition labelled by a is replaced by $n + 2$ transitions with same source and target states, labelled by $a'_0, \ldots, a'_{n+1}$. Action hiding is a special case of admissible action mapping.

Definition 4 (Action hiding). *Let $P = (\Sigma_P, A_P, \longrightarrow_P, p_{init})$ and $A \subseteq \mathcal{A} \setminus \{\tau\}$. We write "**hide** A **in** P" for the LTS $\rho(P)$, where ρ is the admissible action mapping defined by $(\forall a \in A_P \cap A)\ \rho(a) = \{\tau\}$ and $(\forall a \in A_P \setminus A)\ \rho(a) = \{a\}$.*

Parallel composition and admissible action mapping subsume all abstraction and composition operators encodable as *networks of LTS* [20, 11, 7], such as the parallel composition, hiding, renaming, and cut (or restriction) operators of CCS [25], CSP [2], mCRL [15], LOTOS [16], E-LOTOS [17], and LNT [3], as well as synchronization vectors[3]. In the sequel, we write $P_1 \,||\, \ldots \,||\, P_n$ for any

[3] For instance, the composition of P and Q where action a of P synchronizes with either b or c of Q, can be written as $\rho(P) \,|[b, c]|\, Q$, where ρ maps a onto $\{b, c\}$.

expression composing $P_1, \ldots, P_n$ using these operators. Given any partition of $P_1, \ldots, P_n$ into arbitrary subsets $\mathcal{P}_1$ and $\mathcal{P}_2$, it is always possible to rewrite $P_1 \mid\mid \ldots \mid\mid P_n$ in the form $(||_{P_i \in \mathcal{P}_1} P_i) \mid\mid (||_{P_j \in \mathcal{P}_2} P_j)$, even for non-associative parallel composition operators (e.g., $|[\ldots]|$), using appropriate action mappings[4].

LTS can be compared and reduced with respect to well-known bisimulation relations. In this paper, we consider strong bisimulation [27] and divbranching bisimulation, which itself derives from branching bisimulation [34, 35].

Definition 5 (Bisimulations). *A strong bisimulation is a symmetric relation $R \subseteq \Sigma \times \Sigma$ such that if $(p_1, p_2) \in R$ then: for all $p_1 \xrightarrow{a} p_1'$, there exists p_2' such that $p_2 \xrightarrow{a} p_2'$ and $(p_1', p_2') \in R$. A branching bisimulation is a symmetric relation $R \subseteq \Sigma \times \Sigma$ such that if $(p_1, p_2) \in R$ then: for all $p_1 \xrightarrow{a} p_1'$, either $a = \tau$ and $(p_1', p_2) \in R$, or there exists a sequence $p_2 \xrightarrow{\tau^*} p_2' \xrightarrow{a} p_2''$ such that $(p_1, p_2') \in R$ and $(p_1', p_2'') \in R$. A divergence-preserving branching bisimulation (divbranching bisimulation for short) is a branching bisimulation R such that if $(p_1^0, p_2^0) \in R$ and there is an infinite sequence $p_1^0 \xrightarrow{\tau} p_1^1 \xrightarrow{\tau} p_1^2 \xrightarrow{\tau} \ldots$ with $(p_1^i, p_2^0) \in R$ for all $i \geq 0$, then there is an infinite sequence $p_2^0 \xrightarrow{\tau} p_2^1 \xrightarrow{\tau} p_2^2 \xrightarrow{\tau} \ldots$ such that $(p_1^i, p_2^j) \in R$ for all $i, j \geq 0$. Two states p_1 and p_2 are strongly (resp. branching, divbranching) bisimilar, written $p_1 \sim p_2$ (resp. $p_1 \sim_{br} p_2$, $p_1 \sim_{dsbr} p_2$), if there exists a strong (resp. branching, divbranching) bisimulation R such that $(p_1, p_2) \in R$. Two LTS P_1 and P_2 are strongly (resp. branching, divbranching) bisimilar, written $P_1 \sim P_2$ (resp. $P_1 \sim_{br} P_2$, $P_1 \sim_{dsbr} P_2$), if their initial states are strongly (resp. branching, divbranching) bisimilar.*

Strong, branching, and divbranching bisimulations are congruences for parallel composition and admissible action mapping. This allows reductions to be applied at any intermediate step during the state space construction, thus potentially reducing the overall cost of reduction. However, since processes may constrain each other by synchronization, composing LTS two by two following the algebraic structure of the composition expression and applying reduction after each composition can be orders of magnitude less efficient than other strategies in terms of the largest intermediate LTS. Finding an optimal strategy is difficult. One generally relies on heuristics to select a subset of LTS to compose at each step of the compositional reduction. In this paper, we will use the *smart reduction* heuristic [6, 11], which is implemented within the SVL [10] tool of CADP [12]. This heuristic tries to find an efficient composition order by analysing the synchronization and hiding structure of the composition expression.

2.2 Temporal logics

Definition 6 (Modal μ-calculus [18]). *The modal μ-calculus (L_μ) is built from action formulas α and state formulas φ, whose syntax and semantics w.r.t.*

[4] For instance, $P_1 \,|[a]|\, (P_2 \,|[]|\, P_3)$ is equivalent to $\rho_0((\rho_1(P_1) \,|[a_1]|\, \rho_2(P_2)) \,|[a_2]|\, \rho_3(P_3))$ —observe the different associativity— where ρ_1 maps a onto $\{a_1, a_2\}$, ρ_2 renames a into a_1, ρ_3 renames a into a_2, and ρ_0 renames a_1 and a_2 into a.

an LTS $P = (\Sigma, A, \longrightarrow, p_{init})$ *are defined as follows:*

$$
\begin{array}{lll}
\alpha ::= a & & [\![a]\!]_A = \{a\} \\
\quad | \ \mathsf{false} & & [\![\mathsf{false}]\!]_A = \emptyset \\
\quad | \ \alpha_1 \vee \alpha_2 & & [\![\alpha_1 \vee \alpha_2]\!]_A = [\![\alpha_1]\!]_A \cup [\![\alpha_2]\!]_A \\
\quad | \ \neg\alpha_0 & & [\![\neg\alpha_0]\!]_A = A \setminus [\![\alpha_0]\!]_A \\
\varphi ::= \mathsf{false} & & [\![\mathsf{false}]\!]_P\delta = \emptyset \\
\quad | \ \varphi_1 \vee \varphi_2 & & [\![\varphi_1 \vee \varphi_2]\!]_P\delta = [\![\varphi_1]\!]_P\delta \cup [\![\varphi_2]\!]_P\delta \\
\quad | \ \neg\varphi_0 & & [\![\neg\varphi_0]\!]_P\delta = \Sigma \setminus [\![\varphi_0]\!]_P\delta \\
\quad | \ \langle\alpha\rangle\,\varphi_0 & & [\![\langle\alpha\rangle\,\varphi_0]\!]_P\delta = \{s \in \Sigma \mid \exists s \xrightarrow{a} s'.a \in [\![\alpha]\!]_A \wedge s' \in [\![\varphi_0]\!]_P\delta\} \\
\quad | \ X & & [\![X]\!]_P\delta = \delta(X) \\
\quad | \ \mu X.\varphi_0 & & [\![\mu X.\varphi_0]\!]_P\delta = \bigcup_{k\geq 0} \Phi_{0\,P,\delta}^{k}(\emptyset)
\end{array}
$$

where $X \in \mathcal{X}$ *are propositional variables denoting sets of states,* $\delta : \mathcal{X} \to 2^{\Sigma}$ *is a context mapping propositional variables to sets of states,* $[\,]$ *is the empty context,* $\delta[U/X]$ *is the context identical to* δ *except for variable* X*, which is mapped to state set* U*, the functional* $\Phi_{0\,P,\delta} : 2^{\Sigma} \to 2^{\Sigma}$ *associated to the formula* $\mu X.\varphi_0$ *is defined as* $\Phi_{0\,P,\delta}(U) = [\![\varphi_0]\!]_P\delta[U/X]$*, and* Φ^k *means* k*-fold application. We write* $P \models \varphi$ *(read* P *satisfies* φ*) for* $p_0 \in [\![\varphi]\!]_P[\,]$*.*

Action formulas α are built from actions and boolean operators. State formulas φ are built from boolean operators, the possibility modality $\langle\alpha\rangle\,\varphi_0$ denoting the states with an outgoing transition labeled by an action satisfying α and leading to a state satisfying φ_0, and the minimal fixed point operator $\mu X.\varphi_0$ denoting the least solution of the equation $X = \varphi_0$ interpreted over 2^{Σ}.

The usual derived operators are defined as follows: boolean connectors $\mathsf{true} = \neg\mathsf{false}$ and $\varphi_1 \wedge \varphi_2 = \neg(\neg\varphi_1 \vee \neg\varphi_2)$; necessity modality $[\alpha]\,\varphi_0 = \neg\langle\alpha\rangle\,\neg\varphi_0$; and maximal fixed point operator $\nu X.\varphi_0 = \neg\mu X.\neg\varphi_0[\neg X/X]$, where $\varphi_0[\neg X/X]$ is the syntactic substitution of X by $\neg X$ in φ_0. Syntactically, $\langle\rangle$ and $[]$ have the highest precedence, followed by $\wedge$, then $\vee$, and finally μ and ν. To have a well-defined semantics, state formulas are syntactically monotonic [18], i.e., in every subformula $\mu X.\varphi_0$, all occurrences of X in φ_0 fall in the scope of an even number of negations. Thus, negations can be eliminated by downward propagation.

Although L_μ subsumes most action-based logics, its operators are rather low-level and lead to complex formulas. In practice, temporal logics or extensions of L_μ with higher-level operators are used, avoiding (or at least reducing) the use of fixed point operators and modalities. We review informally some of these logics (whose operators can be translated to L_μ), which will be useful in the sequel.

Propositional Dynamic Logic with Looping The logic PDL-Δ [30] introduces the modalities $\langle\beta\rangle\,\varphi_0$ and $\langle\beta\rangle\,@$, where β is a regular formula defined as follows:

$$\beta ::= \varphi? \mid \alpha \mid \beta_1 \cdot \beta_2 \mid \beta_1 \,|\, \beta_2 \mid \beta_0^*$$

Regular formulas β denote sets of transition sequences in an LTS: the testing operator $\varphi?$ denotes all zero-step sequences consisting of states satisfying φ; α

denotes all one-step sequences consisting of a transition labeled by an action satisfying α; the concatenation $\beta_1 \cdot \beta_2$, choice $\beta_1 \,|\, \beta_2$, and transitive-reflexive closure β_0^* operators have their usual semantics transposed to transition sequences.

The regular diamond modality $\langle \beta \rangle \varphi_0$ denotes the states with an outgoing transition sequence satisfying β and leading to a state satisfying φ_0. The infinite looping operator $\langle \beta \rangle @$ denotes the states having an outgoing transition sequence consisting of an infinite concatenation of subsequences satisfying β.

Action Computation Tree Logic The logic ACTL\X (ACTL without next operator) [26] introduces four temporal operators, whose semantics can be found in terms of L_μ formulas in [8, 24], where α_1, α_2 are interpreted over visible actions:

$$\mathsf{E}(\varphi_1 \,{}_{\alpha_1}\mathsf{U}\, \varphi_2), \mathsf{E}(\varphi_1 \,{}_{\alpha_1}\mathsf{U}_{\alpha_2}\, \varphi_2), \mathsf{A}(\varphi_1 \,{}_{\alpha_1}\mathsf{U}\, \varphi_2), \mathsf{A}(\varphi_1 \,{}_{\alpha_1}\mathsf{U}_{\alpha_2}\, \varphi_2)$$

A transition sequence satisfies the path formula $\varphi_1 \,{}_{\alpha_1}\mathsf{U}_{\alpha_2}\, \varphi_2$ if it contains a visible transition whose action satisfies α_2 and whose target state satisfies φ_2, whereas at any moment before this transition, φ_1 holds and all visible actions satisfy α_1. A sequence satisfies $\varphi_1 \,{}_{\alpha_1}\mathsf{U}\, \varphi_2$ if it contains a state satisfying φ_2 and at any moment before, φ_1 holds and all visible actions satisfy α_1. A state satisfies $\mathsf{E}(\varphi_1 \,{}_{\alpha_1}\mathsf{U}_{\alpha_2}\, \varphi_2)$ (resp. $\mathsf{E}(\varphi_1 \,{}_{\alpha_1}\mathsf{U}\, \varphi_2)$) if it has an outgoing sequence satisfying $\varphi_1 \,{}_{\alpha_1}\mathsf{U}_{\alpha_2}\, \varphi_2$ (resp. $\varphi_1 \,{}_{\alpha_1}\mathsf{U}\, \varphi_2$). It satisfies $\mathsf{A}(\varphi_1 \,{}_{\alpha_1}\mathsf{U}_{\alpha_2}\, \varphi_2)$ (resp. $\mathsf{A}(\varphi_1 \,{}_{\alpha_1}\mathsf{U}\, \varphi_2)$) if all its outgoing sequences satisfy the corresponding path formula. The following abbreviations are often used:

$$\mathsf{EF}_\alpha(\varphi_0) = \mathsf{E}(\mathsf{true} \,{}_{\mathsf{true}}\mathsf{U}_\alpha\, \varphi_0) \qquad \mathsf{AG}_\alpha(\varphi_0) = \neg\mathsf{EF}_{\neg\alpha}(\mathsf{true}) \wedge \neg\mathsf{E}(\mathsf{true} \,{}_{\mathsf{true}}\mathsf{U}\, \neg\varphi_0)$$

A state satisfies $\mathsf{EF}_\alpha(\varphi_0)$ if it has an outgoing sequence leading to a transition whose action satisfies α and target state satisfies φ_0. A state satisfies $\mathsf{AG}_\alpha(\varphi_0)$ if none of its outgoing sequences leads to a transition labeled by an action not satisfying α or to a state not satisfying φ_0.

Computation Tree Logic The logic CTL [5] contains the following operators:

$$\mathsf{E}(\varphi_1\, \mathsf{U}\, \varphi_2), \mathsf{A}(\varphi_1\, \mathsf{U}\, \varphi_2), \mathsf{E}(\varphi_1\, \mathsf{W}\, \varphi_2), \mathsf{A}(\varphi_1\, \mathsf{W}\, \varphi_2), \mathsf{EF}(\varphi_0), \mathsf{AG}(\varphi_0), \mathsf{AF}(\varphi_0), \mathsf{EG}(\varphi_0)$$

A state satisfies $\mathsf{E}(\varphi_1\, \mathsf{U}\, \varphi_2)$ (resp. $\mathsf{A}(\varphi_1\, \mathsf{U}\, \varphi_2)$) if some of (resp. all) its outgoing sequences lead to states satisfying φ_2 after passing only through states satisfying φ_1. It satisfies $\mathsf{E}(\varphi_1\, \mathsf{W}\, \varphi_2)$ (resp. $\mathsf{A}(\varphi_1\, \mathsf{W}\, \varphi_2)$) if some of (resp. all) its outgoing sequences either contain only states satisfying φ_1, or lead to states satisfying φ_2 after passing only through states satisfying φ_1. A state satisfies $\mathsf{EF}(\varphi_0)$ (resp. $\mathsf{AF}(\varphi_0)$) if some of (resp. all) its outgoing sequences lead to states satisfying φ_0. A state satisfies $\mathsf{EG}(\varphi_0)$ (resp. $\mathsf{AG}(\varphi_0)$) if some of (resp. all) its outgoing sequences contain only states satisfying φ_0.

2.3 Compositional property-dependent LTS reductions

Given a formula $\varphi \in L_\mu$ and a composition of processes $P_1 \,||\, \ldots \,||\, P_n$, [24] shows two results that allow an LTS equivalent to $P_1 \,||\, \ldots \,||\, P_n$ to be reduced compositionally, while preserving the truth value of φ. The first result is a procedure,

called *maximal hiding*, which extracts systematically from φ a set of actions $H(\varphi)$ that are not discriminated by any action formula occurring in φ. It is shown that $P_1 \parallel \ldots \parallel P_n \models \varphi$ if and only if **hide** $H(\varphi)$ **in** $(P_1 \parallel \ldots \parallel P_n) \models \varphi$. The second result is the identification of a fragment of L_μ, called L_μ^{dsbr}, which is strictly more expressive than μACTL\X[5] and adequate with divbranching bisimulation. This fragment is defined as follows.

Definition 7 (Modal μ-calculus fragment L_μ^{dsbr} [24]). *By convention, we use the symbols α_τ and α_a to denote action formulas such that $\tau \in [\![\alpha_\tau]\!]_A$ and $\tau \notin [\![\alpha_a]\!]_A$. The fragment L_μ^{dsbr} of L_μ is defined as the set of formulas that are semantically equivalent to some formula of the following language:*

$$\begin{array}{l} \varphi ::= \mathsf{false} \mid \varphi_1 \vee \varphi_2 \mid \neg\varphi_0 \mid X \mid \mu X.\varphi_0 \\ \quad \mid \ \langle (\varphi_1?.\alpha_\tau)^* \rangle\, \varphi_2 \mid \langle (\varphi_1?.\alpha_\tau)^*.\varphi_1?.\alpha_a \rangle\, \varphi_2 \mid \langle \varphi_1?.\alpha_\tau \rangle\, @ \end{array}$$

The ultra-weak modality $\langle (\varphi_1?.\alpha_\tau)^* \rangle\, \varphi_2$, weak modality $\langle (\varphi_1?.\alpha_\tau)^*.\varphi_1?.\alpha_a \rangle\, \varphi_2$, and weak infinite looping modality $\langle \varphi_1?.\alpha_\tau \rangle\, @$ are shorthand notations for the respective L_μ formulas $\mu X.\varphi_2 \vee (\varphi_1 \wedge \langle \alpha_\tau \rangle\, X)$, $\mu X.\varphi_1 \wedge (\langle \alpha_a \rangle\, \varphi_2 \vee \langle \alpha_\tau \rangle\, X)$, and $\nu X.\varphi_1 \wedge \langle \alpha_\tau \rangle\, X$. Derived operators are also defined as follows:

$$\begin{array}{c} [(\varphi_1?.\alpha_\tau)^*]\, \varphi_2 = \neg\langle (\varphi_1?.\alpha_\tau)^* \rangle\, \neg\varphi_2 \\ [\varphi_1?.\alpha_\tau]\, \dashv = \neg\langle \varphi_1?.\alpha_\tau \rangle\, @ \\ [(\varphi_1?.\alpha_\tau)^*.\varphi_1?.\alpha_a]\, \varphi_2 = \neg\langle (\varphi_1?.\alpha_\tau)^*.\varphi_1?.\alpha_a \rangle\, \neg\varphi_2 \end{array}$$

Depending on the L_μ fragment φ belongs to, it is thus possible to determine whether the system can or cannot be reduced for divbranching bisimulation.

3 Combining Bisimulations Compositionally

The above approach is a *mono-bisimulation* approach: either the formula is in L_μ^{dsbr} and then the system is entirely reduced for divbranching bisimulation, or it is not and then the system is entirely reduced for strong bisimulation. We show in this section that, even if the formula is not in L_μ^{dsbr}, it may still be possible to reduce some processes among the parallel processes $P_1, \ldots, P_n$ for divbranching instead of strong bisimulation. This approach relies on the fact that, in general, an arbitrary temporal logic formula φ may be rewritten in a form that contains both weak modalities, as those present in L_μ^{dsbr}, and non-weak modalities of L_μ (called *strong modalities* in this context).

To do so, we characterize a family of fragments of L_μ, each of which is written $L_\mu^{strong}(A_s)$, where A_s is the set of actions that can be matched by strong modalities. We then prove that if φ belongs to $L_\mu^{strong}(A_s)$ and some process P_i does not contain any action from the set A_s, then P_i can be reduced for divbranching bisimulation. Throughout this section, we assume that the concurrent system $P_1 \parallel \ldots \parallel P_n$ is fixed, and we write A for the set of actions occurring in the system.

[5] μACTL\X denotes ACTL\X plus fixed points. The authors of [24] claim that L_μ^{dsbr} is as expressive as μACTL\X, but they omit that the $\langle \varphi_1?.\alpha_\tau \rangle\, @$ weak infinite looping modality cannot be expressed in μACTL\X.

3.1 The $L_\mu^{strong}(A_s)$ fragments of L_μ

Definition 8 ($L_\mu^{strong}(A_s)$). *Let $A_s \subseteq A$ be a fixed set of actions, called strong actions, and let α_s denote any action formula such that $[\![\alpha_s]\!]_A \subseteq A_s$, called a strong action formula. The fragment $L_\mu^{strong}(A_s)$ of L_μ is defined as the set of formulas that are semantically equivalent to some formula of the following language:*

$$\begin{aligned}\varphi ::= {} & \mathsf{false} \mid \varphi_1 \vee \varphi_2 \mid \neg\varphi_0 \mid \langle\alpha_s\rangle\,\varphi_0 \mid X \mid \mu X.\varphi_0 \\ & \mid \langle(\varphi_1?.\alpha_\tau)^*\rangle\,\varphi_2 \mid \langle(\varphi_1?.\alpha_\tau)^*.\varphi_1?.\alpha_a\rangle\,\varphi_2 \mid \langle\varphi_1?.\alpha_\tau\rangle\,@\end{aligned}$$

We call $\langle\alpha_s\rangle\,\varphi_0$ a *strong modality*. $L_\mu^{strong}(A_s)$ is the fragment of L_μ consisting of formulas expressible in a form where strong modalities match only actions in A_s. Its formal relationship with L_μ^{dsbr} and L_μ is given in Theorem 1.

Theorem 1. *The following three propositions hold trivially: $L_\mu^{strong}(\emptyset) = L_\mu^{dsbr}$, $L_\mu^{strong}(\mathcal{A}) = L_\mu$, and if $A_s \subset A'_s$ then $L_\mu^{strong}(A_s) \subset L_\mu^{strong}(A'_s)$.*

Given $\varphi \in L_\mu$, there exists a (not necessarily unique, see Theorem 3 page 10) minimal set A_s such that $\varphi \in L_\mu^{strong}(A_s)$. Obviously, $L_\mu^{strong}(A_s)$ is not adequate with divbranching bisimulation when A_s is not empty, as illustrated by the following example.

Example 1. Consider the LTS P, P', Q, and Q' depicted in Figure 1. P' (resp. Q') denotes the minimal LTS equivalent to P (resp. Q) for divbranching bisimulation. The formula $\varphi = [(\mathsf{true}?.\mathsf{true})^*.\mathsf{true}?.a_1]\,[a_2]\,\mathsf{false}$ of $L_\mu^{strong}(\{a_2\})$ (which is equivalent to the PDL formula $[\mathsf{true}^*.a_1.a_2]\,\mathsf{false}$) expresses that the system does not contain two successive transitions labelled by a_1 and a_2 respectively. φ does not belong to L_μ^{dsbr}. Indeed, $P \;|[a_1]|\; Q$ satisfies φ because a_1 is necessarily followed by a τ transition, but $P' \;|[a_1]|\; Q'$ (which is isomorphic to Q') does not.

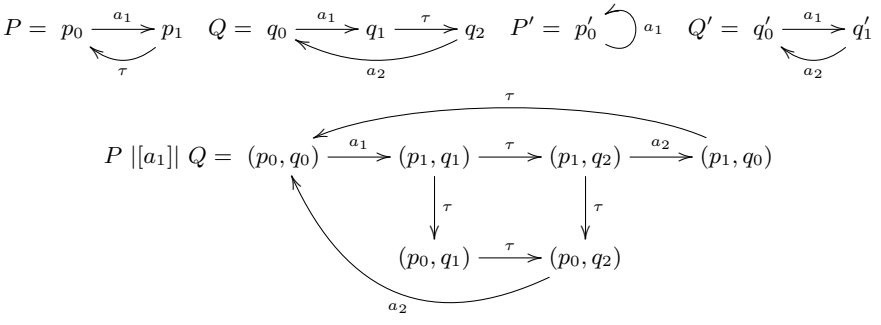

Fig. 1. LTS used in Examples 1 and 2

3.2 Applying divbranching bisimulation to selected components

The following theorem states the main result of this paper, namely that every component process containing no strong action can be replaced by any divbranching equivalent process, without affecting the truth value of the formula[6].

Theorem 2. *Let* $P = (\Sigma_P, A_P, \rightarrow_P, p_{init})$, $Q = (\Sigma_Q, A_Q, \longrightarrow_Q, q_{init})$, $Q' = (\Sigma_{Q'}, A_{Q'}, \longrightarrow_{Q'}, q'_{init})$, $A_{sync} \subseteq \mathcal{A}$, *and* $\varphi \in L_\mu^{strong}(A_s)$. *If* $A_Q \cap A_s = \emptyset$ *and* $Q \sim_{dsbr} Q'$, *then* $P\,|[A_{sync}]|\,Q \models \varphi$ *if and only if* $P\,|[A_{sync}]|\,Q' \models \varphi$.

Proof. The proof looks like the one in [24], showing that divbranching bisimulation preserves the properties of L_μ^{dsbr}, but reasoning concerns product states and additionally handles the case of strong modalities. □

Note that τ may belong to A_s. If so, every P_i containing τ cannot be reduced for divbranching bisimulation. On the contrary, processes that do not contain strong actions do not contain τ. Reducing them for divbranching bisimulation is thus allowed, but coincides with strong bisimulation reduction. In the end, all τ-transitions of the system are preserved, as expected for the truth value of formulas containing strong modalities matching τ to be preserved.

Example 2. In Example 1, P does not contain a_2, the only strong action of the system. Thus, φ can be checked on $P'\ |[a_1]|\ Q$ (which is isomorphic to Q and has only 3 states) instead of $P\ |[a_1]|\ Q$ (6 states), while preserving its truth value.

Theorem 2 is consistent with Andersen's *partial model checking* [1] and the mono-bisimulation approach [24]. Given $P\,||\,Q$ such that the strong actions of φ occur only in P, one can observe that the quotient $\varphi//P$ (defined in [1, 21]) belongs to L_μ^{dsbr}, because quotienting removes all strong modalities, leaving only weak modalities in the quotiented formula. It follows that Q, on which $\varphi//P$ has to be checked, can be reduced for divbranching bisimulation. This observation could serve as an alternative proof of Theorem 2.

3.3 Identifying strong actions in derived operators

In the general case, identifying a minimal set of strong actions is not easy, if even feasible. One cannot reasonably assume that formulas are written in the obscure $L_\mu^{strong}(A_s)$ syntax (see Ex. 1) and that the remaining strong modalities cannot be turned to weak ones. Instead, users shall continue to use "syntactic sugar" extensions of L_μ, with operators of e.g., CTL, ACTL, PDL, or PDL-Δ. In Lemma 1, we provide patterns that can be used to prove that a formula written using one of those operators belongs to a particular instance of $L_\mu^{strong}(A_s)$.

Lemma 1. *Let* $\varphi_0, \varphi_1, \varphi_2 \in L_\mu^{strong}(A_s)$, $\tau \in [\![\alpha_\tau]\!]_A$, $\tau \notin [\![\alpha_a]\!]_A$, $[\![\alpha_s]\!]_A \subseteq A_s$, *and* $\alpha_0, \alpha_1, \alpha_2$ *be any action formulas. The following formulas belong to* $L_\mu^{strong}(A_s)$ *(the list may be not exhaustive):*

[6] Theorem 2 generalizes easily to more general compositions $P\,||\,Q$ (with admissible action mappings) if Q does not contain any action that maps onto a strong action.

1. ***Modal μ-calculus:***
   $\langle \alpha_s \rangle \varphi_0 \quad [\alpha_s] \varphi_0 \quad \neg \varphi_0 \quad \varphi_1 \vee \varphi_2 \quad \varphi_1 \wedge \varphi_2 \quad \varphi_1 \Rightarrow \varphi_2$
2. ***Propositional Dynamic Logic:***
   $\langle \alpha_\tau^* \rangle \varphi_0 \quad [\alpha_\tau^*] \varphi_0 \quad \langle \alpha_\tau^* \cdot \alpha_a \rangle \varphi_0 \quad [\alpha_\tau^* \cdot \alpha_a] \varphi_0 \quad \langle \alpha_\tau \rangle @ \quad [\alpha_\tau] \dashv$
3. ***Action Computation Tree Logic:***
   $\mathsf{A}(\varphi_1 \, {}_{\alpha_1}\mathsf{U} \, \varphi_2) \quad \mathsf{A}(\varphi_1 \, {}_{\alpha_1}\mathsf{U}_{\alpha_2} \, \varphi_2) \quad \mathsf{AG}_{\alpha_0}(\varphi_0)$
   $\mathsf{E}(\varphi_1 \, {}_{\alpha_1}\mathsf{U} \, \varphi_2) \quad \mathsf{E}(\varphi_1 \, {}_{\alpha_1}\mathsf{U}_{\alpha_2} \, \varphi_2) \quad \mathsf{EF}_{\alpha_0}(\varphi_0)$
4. ***Computation Tree Logic:***
   $\mathsf{A}(\varphi_1 \, \mathsf{U} \, \varphi_2) \quad \mathsf{A}(\varphi_1 \, \mathsf{W} \, \varphi_2) \quad \mathsf{AG}(\varphi_0) \quad \mathsf{AF}(\varphi_0)$
   $\mathsf{E}(\varphi_1 \, \mathsf{U} \, \varphi_2) \quad \mathsf{E}(\varphi_1 \, \mathsf{W} \, \varphi_2) \quad \mathsf{EF}(\varphi_0) \quad \mathsf{EG}(\varphi_0)$
   $\mathsf{A}([\alpha_a] \varphi_1 \, \mathsf{U} \, \varphi_2) \quad \mathsf{A}([\alpha_a] \varphi_1 \, \mathsf{W} \, \varphi_2) \quad \mathsf{AG}([\alpha_a] \varphi_0) \quad \mathsf{EF}(\langle \alpha_a \rangle \varphi_0)$
   $\mathsf{AG}(\varphi_1 \vee [\alpha_a] \varphi_2) \quad \mathsf{EF}(\varphi_1 \wedge \langle \alpha_a \rangle \varphi_2)$

Example 3. Let a_1, a_2, and a_3 be visible actions and recall that A denotes the set of all actions of the system. Lemma 1 allows the following to be shown (this is left as an exercise):

$$\langle \mathsf{true}^*.a_1.(\neg a_2)^*.a_3 \rangle \, \mathsf{true} \in L_\mu^{strong}(\emptyset) \qquad [\mathsf{true}] \, \mathsf{false} \in L_\mu^{strong}(A)$$
$$\mathsf{A}(\langle a_1 \rangle \, \mathsf{true} \, {}_{\neg a_2}\mathsf{U}_{a_3} \, \mathsf{true}) \in L_\mu^{strong}(\{a_1\}) \qquad \mathsf{AG}([a_1] \, \mathsf{false}) \in L_\mu^{strong}(\emptyset)$$
$$\mathsf{E}(\mathsf{true} \, {}_{\mathsf{true}}\mathsf{U}_\tau \, \mathsf{true}) \in L_\mu^{strong}(A) \qquad \langle a_1^*.a_2 \rangle \, \mathsf{true} \in L_\mu^{strong}(\{a_1, a_2\})$$
$$\mathsf{E}(\mathsf{true} \, {}_{\mathsf{true}}\mathsf{U} \, \langle \tau \rangle \, \mathsf{true}) \in L_\mu^{strong}(\{\tau\}) \qquad [a_1.a_2] \, \mathsf{false} \in L_\mu^{strong}(\{a_1, a_2\})$$
$$\mathsf{EF}(\langle a_1 \rangle \, \mathsf{true} \wedge \langle a_2 \rangle \, \mathsf{true}) \in L_\mu^{strong}(\{a_1\}) \qquad \mathsf{EF}(\langle a_1 \rangle \, \langle a_2 \rangle \, \mathsf{true}) \in L_\mu^{strong}(\{a_2\})$$
$$\mathsf{EF}(\langle a_1 \rangle \, \mathsf{true} \wedge \langle a_2 \rangle \, \mathsf{true}) \in L_\mu^{strong}(\{a_2\}) \qquad \mathsf{EF}(\langle \neg a_1 \rangle \, \mathsf{true}) \in L_\mu^{strong}(A \setminus \{a_1\})$$

Theorem 3. *There is not a unique minimal set A_s such that $\varphi \in L_\mu^{strong}(A_s)$.*

Proof. $\mathsf{EF}(\langle a_1 \rangle \, \mathsf{true} \wedge \langle a_2 \rangle \, \mathsf{true}) \in L_\mu^{strong}(\{a_1\}) \cap L_\mu^{strong}(\{a_2\})$, because it is semantically equivalent to both formulas $\mathsf{EF}(\langle \langle a_1 \rangle \, \mathsf{true}?.\mathsf{true})^*.\langle a_1 \rangle \, \mathsf{true}?.a_2 \rangle \, \mathsf{true})$ and $\mathsf{EF}(\langle \langle \langle a_2 \rangle \, \mathsf{true}?.\mathsf{true})^*.\langle a_2 \rangle \, \mathsf{true}?.a_1 \rangle \, \mathsf{true})$. Yet, it is not in $L_\mu^{strong}(\emptyset)$ as it has not the same truth value on the divbranching equivalent LTS P and P' below:

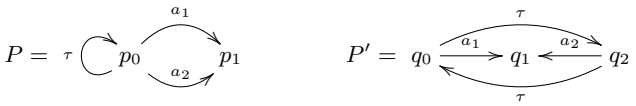

Thus, $\{a_1\}$ and $\{a_2\}$ are non-unique minimal sets of strong actions. □

4 Applications

We consider two examples to illustrate our new verification approach combining strong and divbranching bisimulation and show how it can reduce both time and memory usage when associated to the smart reduction heuristic. In both examples, the aim is to perform a set of verification tasks, each consisting in checking a formula φ on a system of parallel processes $P_1 \parallel \ldots \parallel P_n$. Since our approach can only improve the verification of formulas containing both strong and weak modalities, we consider only the pairs of formulas and systems such

that the formula is part of $L_\mu^{strong}(A_s)$ for some minimal A_s that is not empty and that is strictly included in the set of visible actions of the system[7]. For each verification task, we compare the largest LTS size, the verification time, and the memory peak obtained using the following two approaches:

Mono-bisimulation approach: φ is verified on **hide** $H(\varphi)$ **in** $(P_1 \mid\mid \ldots \mid\mid P_n)$ (where $H(\varphi)$ is the maximal hiding set mentioned in Sect. 2.3) reduced compositionally for strong bisimulation (since φ is not in L_μ^{dsbr}) using the smart reduction heuristic.

Refined approach combining bisimulations: The set $\{P_1, \ldots, P_n\}$ is partitioned in two groups $\mathcal{P}_s$ and $\mathcal{P}_w$ such that $P_i \in \mathcal{P}_s$ if it contains actions in A_s and $P_i \in \mathcal{P}_w$ otherwise, so that $P_1 \mid\mid \ldots \mid\mid P_n$ can be rewritten in the equivalent form $(||_{P_i \in \mathcal{P}_s} P_i) \mid\mid (||_{P_j \in \mathcal{P}_w} P_j)$. The set A_I of actions on which at least one process of $\mathcal{P}_s$ and one process of $\mathcal{P}_w$ synchronize (*inter-group* synchronization) is then identified. Using the smart reduction heuristic, **hide** $H(\varphi) \setminus A_I$ **in** $||_{P_i \in \mathcal{P}_s} P_i$ (corresponding to the processes containing strong actions) is reduced compositionally for strong bisimulation, leading to a first LTS P_s, and **hide** $H(\varphi) \setminus A_I$ **in** $||_{P_j \in \mathcal{P}_w} P_j$ (corresponding to the processes containing no strong action) is reduced compositionally for divbranching bisimulation, leading to a second LTS P_w. Finally, φ is verified on **hide** $H(\varphi) \cap A_I$ **in** $(P_s \mid [A_I] \mid P_w)$ reduced for strong bisimulation.

All experiments were done on a 3GHz/12GB RAM/8-core Intel Xeon computer running Linux, using the specification languages and 32-bit versions of tools provided in the CADP toolbox [12] version 2019-a "Pisa".

4.1 Trivial File Transfer Protocol

The TFTP (*Trivial File Transfer Protocol*) case-study[8] addresses the verification of an avionic communication protocol between a plane and the ground [13]. It comprises two instances (A and B) of a process named TFTP, connected through a FIFO buffer. Since the state space is very large in the general case, the authors defined five scenarios named A, B, C, D, and E, depending on whether each instance may write and/or read a file. The system corresponding to each scenario is a parallel composition of eight processes. The requirements consist of 29 properties parameterized by the identity of a TFTP instance, defined in MCL [23] (an implementation of the alternation-free modal μ-calculus including PDL-Δ modalities and macro definitions enabling the construction of libraries of operators), 24 of which belong to L_μ^{dsbr}. The remaining five, namely properties 08, 09, 14, 16, and 17, contain both weak and strong modalities. The shape of these properties is described in Table 1, where we do not provide the details of the action formulas, but instead denote them by letters $a_1, a_2, \ldots$, where $\tau \notin [\![a_i]\!]_A$

[7] Otherwise, our approach coincides with the mono-bisimulation approach of [24]. In all the examples addressed in this section, there is always a unique minimal set A_s, whose identification is made easy using Lemma 1.

[8] Specification available at `ftp://ftp.inrialpes.fr/pub/vasy/demos/demo_05`

Nr.	Property
08	$[\mathsf{true}^* \cdot a_1 \cdot a_2]\ \mathsf{false}$
09	$[\mathsf{true}^* \cdot a_1 \cdot a_2 \cdot ((a_3 \cdot (\neg a_4)^* \cdot a_5) \mid (a_6 \cdot (\neg a_7)^* \cdot a_8))]\ \mathsf{false}$
14	$[\mathsf{true}^* \cdot a_1 \cdot a_2 \cdot (\neg a_3)^* \cdot a_4 \cdot a_5]\ \mathsf{false}$
16	$[(\neg a_1)^* \cdot a_2 \cdot (\neg a_3)^* \cdot a_4]\ \langle((\neg a_5)^* \cdot a_6 \cdot a_7) \cdot ((\neg a_5)^* \cdot a_6 \cdot a_7)\rangle\ \mathsf{true}$
17	Same shape as property Nr. 16

Table 1. TFTP properties (strong action formulas are highlighted)

for all i. Strong action formulas are highlighted and one shows easily that the other are weak using Lemma 1-2.

We consider 31 among a potential of 50 verification tasks (five properties, five scenarios, and two instances) as some properties are not relevant to every TFTP instance and scenario (e.g., in a scenario where one TFTP instance only receives messages, checking a property concerning a message emission is irrelevant). All 31 verification tasks return true and the strong actions occur in only three (although not the same three) out of the eight parallel processes.

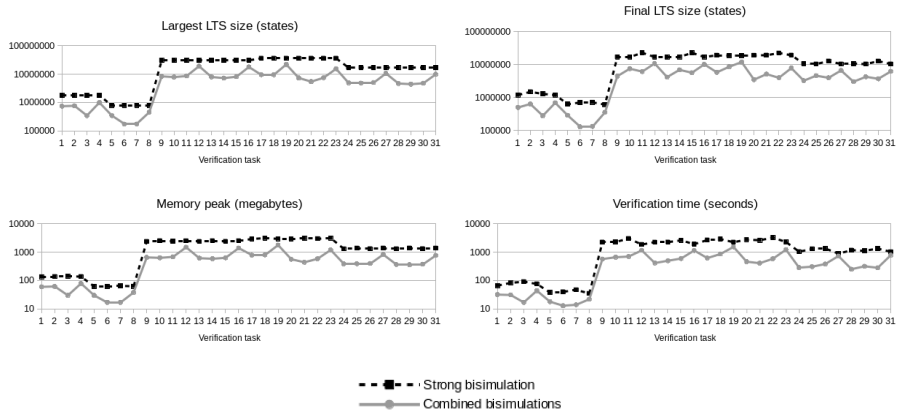

Fig. 2. Experimental results of the TFTP case-study

Figure 2 shows that the refined approach always reduces LTS size (for both intermediate and final LTS), memory and time following similar curves, up to a factor 7 (the vertical axis is on a logarithmic scale). Time does not include LTS generation of the component processes from their LNT specification, which takes only a few seconds and is common to both approaches. In these experiments, time is dominated by the last step of generation and minimization, whereas memory usage is dominated by minimization.

Nr.	Property	Result
101#21	$\mathsf{AG}([A21]\ [A23]\ [A4]\ [\mathsf{true}]\ \mathsf{false})$	false
101#22	$\mathsf{AG}([A3]\ \mathsf{AF}(\langle A2\rangle\ \mathsf{true}))$	false
101#23	$\mathsf{AG}(\langle A20\rangle\ \mathsf{true} \Rightarrow \langle A20\rangle\ \mathsf{A}([A23]\ \mathsf{false}\ \mathsf{W}\ \langle A8\rangle\ \mathsf{true}))$	true
102#21	$\mathsf{EF}(\mathsf{AG}([A5]\ \mathsf{false}))$	true
102#22	$\mathsf{EG}([A35]\ \mathsf{E}([A23]\ \mathsf{false}\ \mathsf{U}\ \langle A35\rangle\ \mathsf{true}))$	false
102#23	$\mathsf{AG}([A22]\ \mathsf{A}([A8]\ \mathsf{false}\ \mathsf{U}\ \langle A22\rangle\ \mathsf{true}))$	false
103#21	$\mathsf{AG}([A11]\ \mathsf{A}([A2]\ \mathsf{false}\ \mathsf{W}\ \langle A6\rangle\ \mathsf{true}) \Rightarrow [A11]\ \mathsf{A}([A5]\ \mathsf{false}\ \mathsf{W}\ \langle A6\rangle\ \mathsf{true}))$	true
103#22	$\mathsf{EG}([A14]\ \mathsf{false} \wedge (\langle A18\rangle\ \mathsf{true} \Rightarrow [A18]\ \mathsf{EG}([A21]\ \mathsf{false} \wedge \mathsf{EF}(\langle A19\rangle\ \mathsf{true}))))$ $= \mathsf{EG}([A14]\ \mathsf{false} \wedge [A18]\ \mathsf{EG}([A21]\ \mathsf{false} \wedge \mathsf{EF}(\langle A19\rangle\ \mathsf{true})))$	true
103#23	$\mathsf{AG}(\langle A34\rangle\ \mathsf{true} \Rightarrow [A34]\ \mathsf{A}([A68]\ \mathsf{false}\ \mathsf{W}\ \langle A59\rangle\ \mathsf{true}))$ $= \mathsf{AG}([A34]\ \mathsf{A}([A68]\ \mathsf{false}\ \mathsf{W}\ \langle A59\rangle\ \mathsf{true}))$	false

Table 2. RERS 2018 properties (strong action formulas are highlighted)

4.2 Parallel benchmark of the RERS 2018 challenge

The RERS (Rigorous Examination of Reactive Systems)[9] challenge is an international competition on a benchmark of verification tasks. Since 2018 (8th edition), the challenge features a set of parallel problems where systems are synchronizing LTS and properties are expressed using CTL and modalities. This section illustrates the benefits of our approach on these problems.

The benchmark comprises three specifications of concurrent systems, numbered 101, 102, and 103, each accompanied by three properties to be checked, numbered $p\#21$, $p\#22$, and $p\#23$, where p is the system number. Thus, nine verification tasks have to be solved. The properties are presented in Table 2, where the strong action formulas are highlighted. One easily shows that all other action formulas are weak using Lemmas 1-1 and 1-4. However, for 103#22 and 103#23, the identity $(\langle\alpha\rangle\,\mathsf{true} \Rightarrow [\alpha]\,\varphi) = ([\alpha]\,\mathsf{false} \vee [\alpha]\,\varphi) = [\alpha]\,\varphi$ (because $[\alpha]\,\mathsf{false} \implies [\alpha]\,\varphi$ for all φ) was applied to obtain the simplified formulas occurring after the = sign in the table. For 102#23, this simplification allowed us to prove that $A34$ is not a strong action, unlike what appears at first sight.

Table 3 gives, for each of the nine verification tasks, the number #act of actions in the system, the number #hide of actions in the maximal hiding set, the number #sact of strong actions, the number #proc of parallel processes, the number #sproc of processes in the strong group, the number #sync of inter-group actions, and the best reduction relation among strong bisimulation, divbranching bisimulation, or a combination of both. We observe that:

- The set of weak actions of 101#21 is empty due to the presence of the "true" strong action formula, whereas the set of strong actions of 102#21 is empty, i.e., the property belongs to L_μ^{dsbr}. In both cases, our approach coincides with the mono-bisimulation approach. The verification of 101#21 (reduced

[9] `http://rers-challenge.org`

Task	#act	#hide	#sact	#proc	#sproc	#sync	relation
101#21	24	21	24	9	9	-	strong
101#22	24	22	1	9	4	11	combination
101#23	24	21	2	9	3	9	combination
102#21	28	27	0	20	0	-	divbranching
102#22	28	26	2	20	10	14	combination
102#23	28	26	1	20	4	12	combination
103#21	70	66	2	34	8	12	combination
103#22	70	66	3	34	6	18	combination
103#23	70	67	1	34	7	10	combination

Table 3. Some numbers about the RERS 2018 parallel benchmark

for strong bisimulation) takes 75 seconds, with a memory peak of 11 MB and a largest LTS of 83,964 states and 374,809 transitions. The verification of 102#21 (reduced for divbranching bisimulation) takes 261 seconds, with a memory peak of 22 MB and a largest LTS of 243 states and 975 transitions.

- 101#22, 101#23, 102#22, 102#23, 103#21, 103#22, and 103#23 contain both weak and strong actions. They are used to evaluate our approach.

Table 4 compares the performance of verifying the latter seven verification tasks using the approaches described above. LTS sizes are given in kilostates, memory in megabytes, and time in seconds. Tasks using more than 3 GB of memory were aborted. We see that our approach reduces both time and memory usage and allows all problems of the challenge to be solved, whereas using strong bisimulation alone fails in five out of those seven tasks.

	Strong bisimulation				Combined bisimulations			
	Kstates		verif.		Kstates		verif.	
Task	largest	final	MB	sec.	largest	final	MB	sec.
101#22	84	77	10	77	1.4	1.4	10	72
101#23	84	77	11	80	0.5	0.5	8	73
102#22	-	-	-	-	611	585	57	295
102#23	-	-	-	-	17	9.8	22	260
103#21	-	-	-	-	734	313	101	604
103#22	-	-	-	-	14,143	14,141	1575	2533
103#23	-	-	-	-	122	122	35	566

Table 4. Experimental results of the RERS 2018 parallel benchmark

The negligible reductions in time and memory usage observed for tasks 101#22 and 101#23 are due to the fact that time and memory usage are dominated by the algorithm in charge of selecting a sub-set of processes to be composed and reduced (implemented in smart reduction). The complexity of this

algorithm does not depend on the state space size, but on the number of actions and parallel processes, which is almost the same using both approaches. When considering larger examples, memory usage gets dominated by minimisation. In particular, for tasks 102#22, 102#23, 103#21, and 103#23 (and likely also 103#22), memory usage is reduced by several orders of magnitude.

Note that some of these tasks can be verified more efficiently using non-compositional approaches, such as on-the-fly model checking, in cases where proofs or counter-examples can be detected much before having explored the full state space. The main drawback of maximal hiding is that the generated counter-examples are given only in terms of the actions visible in the formula, which abstracts out a lot of intermediate transitions. However, this is the price to pay for being able to verify most of the tasks, such as 103#21, for which on-the-fly verification aborts due to memory exhaustion.

5 Conclusion and Future Work

In this paper, we proposed a compositional verification approach that extends the state of the art [24] and consists of three steps: First, so-called strong actions are identified, corresponding to those actions of the system that the formula cannot match using weak modalities in the sense of the L_μ fragment L_μ^{dsbr} adequate with divbranching bisimulation. These actions are used to partition the parallel processes into those containing strong actions and the others. Second, maximal hiding and compositional reduction are used to minimize the composition of processes not containing strong actions for divbranching bisimulation, and the other processes for strong bisimulation. Finally, the property is verified on the reduced system.

The originality of this approach is to combine strong and divbranching bisimulation, as opposed to the mono-bisimulation approach of [24]. We proved it correct by characterizing a family of fragments of the logic L_μ, called $L_\mu^{strong}(A_s)$, parameterized by the set A_s of strong actions. We also showed under which conditions action-based branching-time temporal logic formulas containing well-known operators from the logics CTL, ACTL, PDL, and PDL-Δ are part of $L_\mu^{strong}(A_s)$ when A_s is fixed. In the future, it might be worth investigating whether more operators can be considered, e.g., from the linear-time logic LTL.

This approach may significantly improve the verification performance for systems containing both processes with and without strong actions, as illustrated by two case-studies. In particular, it allowed the whole parallel CTL benchmark of the RERS 2018 challenge to be solved on a standard computer.

Identifying (close to minimal) sets of strong actions for arbitrary formulas manually is a cumbersome task, prone to errors. We shall investigate ways to compute such sets automatically. As illustrated by verification task 103#23 of RERS 2018, the problem is not purely syntactic: considering non-trivial semantic equivalences may prove useful to eliminate actions that appear strong at first sight. Yet, we trust that the presented approach has potential to be implemented in automated software tools, such as those available in the CADP toolbox.

References

1. Henrik Reif Andersen. Partial Model Checking. In *Proceedings of the 10th Annual IEEE Symposium on Logic in Computer Science LICS (San Diego, California, USA)*, pages 398–407. IEEE Computer Society Press, June 1995.
2. Stephen D. Brookes, C. A. R. Hoare, and A. W. Roscoe. A Theory of Communicating Sequential Processes. *Journal of the ACM*, 31(3):560–599, July 1984.
3. David Champelovier, Xavier Clerc, Hubert Garavel, Yves Guerte, Christine McKinty, Vincent Powazny, Frédéric Lang, Wendelin Serwe, and Gideon Smeding. Reference Manual of the LNT to LOTOS Translator (Version 6.7). INRIA, Grenoble, France, July 2017.
4. S. C. Cheung and J. Kramer. Enhancing Compositional Reachability Analysis with Context Constraints. In *Proceedings of the 1st ACM SIGSOFT International Symposium on the Foundations of Software Engineering (Los Angeles, CA, USA)*, pages 115–125. ACM Press, December 1993.
5. 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.
6. Pepijn Crouzen and Frédéric Lang. Smart Reduction. In Dimitra Giannakopoulou and Fernando Orejas, editors, *Proceedings of Fundamental Approaches to Software Engineering (FASE'11), Saarbrücken, Germany*, volume 6603 of *Lecture Notes in Computer Science*, pages 111–126. Springer, March 2011.
7. Sander de Putter, Anton Wijs, and Frédéric Lang. Compositional Model Checking is Lively — Extended Version, 2018. Submitted to Science of Computer Programming.
8. Alessandro Fantechi, Stefania Gnesi, and Gioia Ristori. From ACTL to μ-calculus (extended abstract). In *Proceedings of the Workshop on Theory and Practice in Verification*. ERCIM, 1992.
9. 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.
10. 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'01), Cheju Island, Korea*, pages 377–392. Kluwer Academic Publishers, August 2001. Full version available as INRIA Research Report RR-4223.
11. Hubert Garavel, Frédéric Lang, and Radu Mateescu. Compositional Verification of Asynchronous Concurrent Systems Using CADP. *Acta Informatica*, 52(4):337–392, April 2015.
12. Hubert Garavel, Frédéric Lang, Radu Mateescu, and Wendelin Serwe. CADP 2011: A Toolbox for the Construction and Analysis of Distributed Processes. *Springer International Journal on Software Tools for Technology Transfer (STTT)*, 15(2):89–107, April 2013.
13. Hubert Garavel and Damien Thivolle. Verification of GALS Systems by Combining Synchronous Languages and Process Calculi. In Corina Pasareanu, editor, *Proceedings of the 16th International SPIN Workshop on Model Checking of Software (SPIN'09), Grenoble, France*, volume 5578 of *Lecture Notes in Computer Science*, pages 241–260. Springer, June 2009.

14. Susanne Graf and Bernhard Steffen. Compositional Minimization of Finite State Systems. In Edmund M. Clarke and Robert P. Kurshan, editors, *Proceedings of the 2nd Workshop on Computer-Aided Verification (CAV'90), Rutgers, New Jersey, USA*, volume 531 of *Lecture Notes in Computer Science*, pages 186–196. Springer, June 1990.
15. Jan Friso Groote and Alban Ponse. The Syntax and Semantics of μCRL. CS-R 9076, Centrum voor Wiskunde en Informatica, Amsterdam, 1990.
16. ISO/IEC. LOTOS – A Formal Description Technique Based on the Temporal Ordering of Observational Behaviour. International Standard 8807, International Organization for Standardization – Information Processing Systems – Open Systems Interconnection, Geneva, September 1989.
17. ISO/IEC. Enhancements to LOTOS (E-LOTOS). International Standard 15437:2001, International Organization for Standardization – Information Technology, Geneva, September 2001.
18. D. Kozen. Results on the Propositional μ-calculus. *Theoretical Computer Science*, 27:333–354, 1983.
19. Jean-Pierre Krimm and Laurent Mounier. Compositional State Space Generation from LOTOS Programs. In Ed Brinksma, editor, *Proceedings of the 3rd International Workshop on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'97), University of Twente, Enschede, The Netherlands*, volume 1217 of *Lecture Notes in Computer Science*. Springer, April 1997. Extended version with proofs available as Research Report VERIMAG RR97-01.
20. Frédéric Lang. EXP.OPEN 2.0: A Flexible Tool Integrating Partial Order, Compositional, and On-the-fly Verification Methods. In Judi Romijn, Graeme Smith, and Jaco van de Pol, editors, *Proceedings of the 5th International Conference on Integrated Formal Methods (IFM'05), Eindhoven, The Netherlands*, volume 3771 of *Lecture Notes in Computer Science*, pages 70–88. Springer, November 2005. Full version available as INRIA Research Report RR-5673.
21. Frédéric Lang and Radu Mateescu. Partial Model Checking using Networks of Labelled Transition Systems and Boolean Equation Systems. *Logical Methods in Computer Science*, 9(4):1–32, October 2013.
22. J. Malhotra, S. A. Smolka, A. Giacalone, and R. Shapiro. A Tool for Hierarchical Design and Simulation of Concurrent Systems. In *Proceedings of the BCS-FACS Workshop on Specification and Verification of Concurrent Systems, Stirling, Scotland, UK*, pages 140–152. British Computer Society, July 1988.
23. 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*, volume 5014 of *Lecture Notes in Computer Science*, pages 148–164. Springer, May 2008.
24. Radu Mateescu and Anton Wijs. Property-Dependent Reductions Adequate with Divergence-Sensitive Branching Bisimilarity. *Science of Computer Programming*, 96(3):354–376, 2014.
25. Robin Milner. *Communication and Concurrency*. Prentice-Hall, 1989.
26. R. De Nicola and F. W. Vaandrager. *Action versus State Based Logics for Transition Systems*. In *Semantics of Concurrency*, volume 469 of *Lecture Notes in Computer Science*, pages 407–419. Springer, 1990.
27. David Park. Concurrency and Automata on Infinite Sequences. In Peter Deussen, editor, *Theoretical Computer Science*, volume 104 of *Lecture Notes in Computer Science*, pages 167–183. Springer, March 1981.

28. Amir Pnueli. In Transition from Global to Modular Temporal Reasoning about Programs. *Logic and Models of Concurrent Systems*, 13:123–144, 1984.
29. Krishan K. Sabnani, Aleta M. Lapone, and M. Ümit Uyar. An Algorithmic Procedure for Checking Safety Properties of Protocols. *IEEE Transactions on Communications*, 37(9):940–948, September 1989.
30. R. Streett. Propositional Dynamic Logic of Looping and Converse. *Information and Control*, (54):121–141, 1982.
31. Kuo-Chung Tai and Pramod V. Koppol. An Incremental Approach to Reachability Analysis of Distributed Programs. In *Proceedings of the 7th International Workshop on Software Specification and Design, Los Angeles, CA, USA*, pages 141–150, Piscataway, NJ, December 1993. IEEE Press.
32. Kuo-Chung Tai and Pramod V. Koppol. Hierarchy-Based Incremental Reachability Analysis of Communication Protocols. In *Proceedings of the IEEE International Conference on Network Protocols, San Francisco, CA, USA*, pages 318–325, Piscataway, NJ, October 1993. IEEE Press.
33. Antti Valmari. Compositional State Space Generation. In Grzegorz Rozenberg, editor, *Advances in Petri Nets 1993 – Papers from the 12th International Conference on Applications and Theory of Petri Nets (ICATPN'91), Gjern, Denmark*, volume 674 of *Lecture Notes in Computer Science*, pages 427–457. Springer, 1993.
34. R. J. van Glabbeek and W. Peter Weijland. Branching-Time and Abstraction in Bisimulation Semantics (extended abstract). CS R8911, Centrum voor Wiskunde en Informatica, Amsterdam, 1989. Also in proc. IFIP 11th World Computer Congress, San Francisco, 1989.
35. Rob J. van Glabbeek and W. Peter Weijland. Branching Time and Abstraction in Bisimulation Semantics. *Journal of the ACM*, 43(3):555–600, 1996.
36. Wei Jen Yeh and Michal Young. Compositional Reachability Analysis Using Process Algebra. In *Proceedings of the ACM SIGSOFT Symposium on Testing, Analysis, and Verification (SIGSOFT'91), Victoria, British Columbia, Canada*, pages 49–59. ACM Press, October 1991.

Appendix

This appendix is organized as follows. Section A contains the proof of Theorem 2 (page 9). Section B contains the proof of Lemma 1 (page 9). Section C contains the detailed performance data collected from the TFTP experiment presented in Section 4.1.

Note to referees

An archive was created at `http://doi.org/10.5281/zenodo.2634149`. It contains the experimental material (TFTP and RERS case-studies) and a PDF of this appendix. If the paper is accepted, the URL of the archive will be given in the paper, so that readers can access this companion material and this appendix will be corrected according to the reviewer's comments (if any).

A Proof of Theorem 2 (page 9)

To show Theorem 2, we first need the following Lemma, which simply lifts a standard property of branching and divbranching bisimulations up to the level of product states:

Lemma 2. *In* $P \,|[A_{sync}]|\, Q$*, if* $(p_0, q_0) \xrightarrow{a} (p_1, q_1)$ *and* $q'_0 \sim_{dsbr} q_0$ *then there exists a state* q'_1 *such that* $q'_1 \sim_{dsbr} q_1$ *and either:*

- $a = \tau$, $p_0 = p_1$, $q_0 \sim_{dsbr} q_1$, *and* $q'_1 = q'_0$, *or*
- *there exists* $q''_0, ..., q''_n$ $(n \geq 0)$ *such that* $q''_0 = q'_0$, *for all* $i \in 0..n-1$, $q''_{i+1} \sim_{dsbr} q'_0$, $(p_0, q''_i) \xrightarrow{\tau} (p_0, q''_{i+1})$, *and* $(p_0, q''_n) \xrightarrow{a} (p_1, q'_1)$.

Graphically, this means that one of the diagrams below holds, where solid lines denote universal quantification whereas dotted lines denote existential quantification:

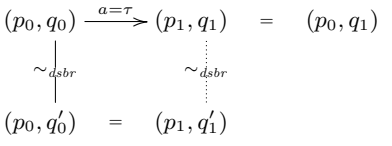

– or –

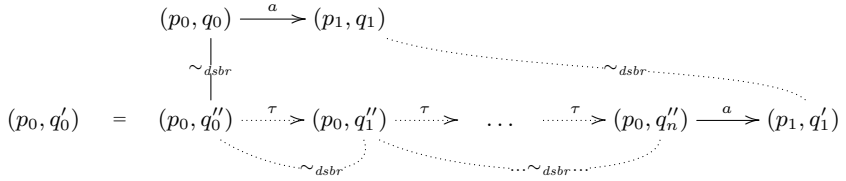

Proof. From the definition of $P \,|[A_{sync}]|\, Q$ and the fact that $(p_0, q_0) \xrightarrow{a} (p_1, q_1)$, there are three possible cases:

1. $a \in A_{sync}$, $p_0 \xrightarrow{a} p_1$, and $q_0 \xrightarrow{a} q_1$, or
2. $a \notin A_{sync}$, $p_0 \xrightarrow{a} p_1$, and $q_1 = q_0$, or
3. $a \notin A_{sync}$, $p_1 = p_0$, and $q_0 \xrightarrow{a} q_1$

Instead of considering those three cases, we merge cases 1 and 3 (where $q_0 \xrightarrow{a} q_1$, whatever p_0 does), which have a similar proof, into a single case. Thus we consider the following two cases, the first one corresponding to case 2 of the above enumeration, and the second one corresponding to cases 1 and 3:

- $q_1 = q_0$, $a \notin A_{sync}$, and $p_0 \xrightarrow{a} p_1$. Take $q'_1 = q'_0$. We have $q'_1 \sim_{dsbr} q_1$ because $q'_1 = q'_0$, $q'_0 \sim_{dsbr} q_0$, and $q_0 = q_1$. The second item of the lemma (bottom diagram) is verified with $n = 0$, $q''_0 = q''_n = q'_0$, and $(p_0, q''_0) = (p_0, q''_n) = (p_0, q'_0) \xrightarrow{a} (p_1, q'_0) = (p_1, q'_1)$.
- $q_0 \xrightarrow{a} q_1$ (and either $a \notin A_{sync}$ and $p_1 = p_0$, or $a \in A_{sync}$ and $p_0 \xrightarrow{a} p_1$). Then since $q'_0 \sim_{dsbr} q_0$, we have by definition of $\sim_{dsbr}$:
  - Either $a = \tau$ and $q_0 \sim_{dsbr} q_1$. Since $\tau \notin A_{sync}$, we have $p_0 = p_1$. Then, we can take $q'_1 = q'_0$. Indeed, $q'_1 \sim_{dsbr} q_1$ because $q'_1 = q'_0$, $q'_0 \sim_{dsbr} q_0$, and $q_0 \sim_{dsbr} q_1$. The first item of the lemma (top diagram) is verified.
  - Or there exists $q''_0, ..., q''_n$ $(n \geq 0)$ such that $q''_0 = q'_0$, for all $i \in 0..n-1$, $q''_{i+1} \sim_{dsbr} q'_0$, $q''_i \xrightarrow{\tau} q''_{i+1}$, and $q''_n \xrightarrow{a} q'_1$. In this case, we also have $(p_0, q''_i) \xrightarrow{\tau} (p_0, q''_{i+1})$ and $(p_0, q''_n) \xrightarrow{a} (p_1, q'_1)$, i.e., the second item of the lemma (bottom diagram) is verified.

□

We draw the reader's attention to the fact that the top diagram of Lemma 2 concerns only the case $(p_0, q_0) \xrightarrow{\tau} (p_0, q_1)$ deriving from a transition $q_0 \xrightarrow{\tau} q_1$ where $q_0 \sim_{dsbr} q_1$. The bottom diagram concerns all other cases, including $(p_0, q_0) \xrightarrow{\tau} (p_1, q_0)$ deriving from $p_0 \xrightarrow{\tau} p_1$ where $p_0 \sim_{dsbr} p_1$. We now show the main theorem:

Theorem 2 *Let $P = (\Sigma_P, A_P, \rightarrow_P, p_{init})$, $Q = (\Sigma_Q, A_Q, \longrightarrow_Q, q_{init})$, $Q' = (\Sigma_{Q'}, A_{Q'}, \longrightarrow_{Q'}, q'_{init})$, $A_{sync} \subseteq \mathcal{A}$, and $\varphi \in L_\mu^{strong}(A_s)$. If $A_Q \cap A_s = \emptyset$ and $Q \sim_{dsbr} Q'$, then $P \,|[A_{sync}]|\, Q \models \varphi$ if and only if $P \,|[A_{sync}]|\, Q' \models \varphi$.*

Proof. We prove that, given $p \in \Sigma_P$, $q \in \Sigma_Q$, and $q' \in \Sigma_{Q'}$ such that $q \sim_{dsbr} q'$, (p, q) satisfies φ if and only if (p, q') satisfies φ. We show this claim by structural induction on the formula φ. Since we work on finite LTS, every formula containing a fixed point operator μ can be expanded into an equivalent formula by unfolding the fixed point a bounded number of times. This way, we do not have to consider fixed points and propositional variables in this proof, and contexts δ mapping propositional variables to sets of states are always empty in the semantics of formulas. For conciseness, we write $[\![\varphi]\!]$ instead of $[\![\varphi]\!]_{P \,|[A_{sync}]|\, Q}[\,]$ and $[\![\varphi]\!]_{P \,|[A_{sync}]|\, Q'}[\,]$, respectively, as the LTS "$P \,|[A_{sync}]|\, Q$" or "$P \,|[A_{sync}]|\, Q'$"

to which the semantics apply is clear from the context. Note that since the claim is symmetric, it is sufficient to prove one implication only. We thus assume that $(p, q) \in [\![\varphi]\!]$ and prove that $(p, q') \in [\![\varphi]\!]$.

- Case $\varphi = \mathsf{false}$. It is obvious that both $(p, q) \notin [\![\varphi]\!]$ and $(p, q') \notin [\![\varphi]\!]$.
- Case $\varphi = \varphi_1 \vee \varphi_2$. By definition, either $(p, q) \in [\![\varphi_1]\!]$ or $(p, q) \in [\![\varphi_2]\!]$. If $(p, q) \in [\![\varphi_1]\!]$ (resp. $[\![\varphi_2]\!]$) then $(p, q') \in [\![\varphi_1]\!]$ (resp. $[\![\varphi_2]\!]$) by the induction hypothesis. Therefore, $(p, q') \in [\![\varphi_1 \vee \varphi_2]\!]$.
- Case $\varphi = \neg\varphi_0$. By definition, $(p, q) \notin [\![\varphi_0]\!]$. By the induction hypothesis, $(p, q') \notin [\![\varphi_0]\!]$ and thus $(p, q') \in [\![\neg\varphi_0]\!]$.
- Case $\varphi = \langle\alpha_s\rangle\, \varphi_0$. By definition, there exists $a \in [\![\alpha_s]\!]_A$ such that $(p, q) \xrightarrow{a} (p_1, q_1)$ and $(p_1, q_1) \in [\![\varphi_0]\!]$.
  Since $[\![\alpha_s]\!]_A \subseteq A_s$, $A_Q \cap A_s = \emptyset$, and $a \in [\![\alpha_s]\!]_A$, we have $a \notin A_Q$ (and hence, $a \notin A_{sync}$). Therefore, $q = q_1$, $p \xrightarrow{a} p_1$, and $(p, q') \xrightarrow{a} (p_1, q')$. Since $q' \sim_{dsbr} q$ and $(p_1, q) \in [\![\varphi_0]\!]$, we have by the induction hypothesis $(p_1, q') \in [\![\varphi_0]\!]$. As a consequence, $(p, q') \in [\![\langle\alpha_s\rangle\, \varphi_0]\!]$.
- Case $\varphi = \langle(\varphi_1?.\alpha_\tau)^*\rangle\, \varphi_2$. By definition, there exist $m \geq 0, p_i \in \Sigma_P, q_i \in \Sigma_Q$, and $a_i \in [\![\alpha_\tau]\!]_A$ $(i \in 0..m)$, such that $(p, q) = (p_0, q_0)$, $(p_m, q_m) \in [\![\varphi_2]\!]$, and for all $0 \leq i < m$, $(p_i, q_i) \xrightarrow{a_i} (p_{i+1}, q_{i+1})$ and $(p_i, q_i) \in [\![\varphi_1]\!]$. We show that the same holds from state (p, q'), i.e., there exists a sequence of transitions matching α_τ and passing through states satisfying φ_1, until reaching a state satisfying φ_2. By the induction hypothesis, $(p, q') = (p_0, q'_0) \in [\![\varphi_1]\!]$. Consider one transition $(p_i, q_i) \xrightarrow{a_i} (p_{i+1}, q_{i+1})$. We know that $(p_i, q_i) \in [\![\varphi_1]\!]$. Assume that there exists $q'_i \sim_{dsbr} q_i$. Substituting a_i for a, (p_i, q_i) for (p_0, q_0), and (p_{i+1}, q_{i+1}) for (p_1, q_1) in the two diagrams of Lemma 2, it appears clearly that there exists a state q'_{i+1} such that $q'_{i+1} \sim_{dsbr} q_{i+1}$. Moreover, by the induction hypothesis, every state divbranching equivalent to (p_i, q_i) in the bottom sequence of the diagrams is in $[\![\varphi_1]\!]$. Finally, note that $\tau \in [\![\alpha_\tau]\!]$ by definition, so that each of the τ-transitions in the bottom sequence of the second diagram thus matches α_τ. Moreover, $(p_m, q_m) \in [\![\varphi_2]\!]$ and $q'_m \sim_{dsbr} q_m$, so by the induction hypothesis, we have also $(p_m, q'_m) \in [\![\varphi_2]\!]$.
- Case $\varphi = \langle(\varphi_1?.\alpha_\tau)^*.\varphi_1?.\alpha_a\rangle\, \varphi_2$. Since $\langle(\varphi_1?.\alpha_\tau)^*.\varphi_1?.\alpha_a\rangle\, \varphi_2$ is equivalent to $\langle(\varphi_1?.\alpha_\tau)^*\rangle\, \langle(\varphi_1?.\tau)^*.\varphi_1?.\alpha_a\rangle\, \varphi_2$ and since we have already considered the case $\varphi = \langle(\varphi_1?.\alpha_\tau)^*\rangle\, \varphi_2$ above, we only have to consider the case $\varphi = \langle(\varphi_1?.\tau)^*.\varphi_1?.\alpha_a\rangle\, \varphi_2$, i.e., $\alpha_\tau = \tau$. By definition, there exist $m \geq 0, p_i \in \Sigma_P, q_i \in \Sigma_Q$ $(i \in 0..m+1)$ and $a \in [\![\alpha_a]\!]_A$, such that $(p, q) = (p_0, q_0)$, $(p_m, q_m) \xrightarrow{a} (p_{m+1}, q_{m+1})$, $(p_{m+1}, q_{m+1}) \in [\![\varphi_2]\!]$, and for all $0 \leq i < m$, $(p_i, q_i) \xrightarrow{\tau} (p_{i+1}, q_{i+1})$ and $(p_i, q_i) \in [\![\varphi_1]\!]$. The reasoning is similar to the previous case. It is important to note that the transition $(p_m, q_m) \xrightarrow{a} (p_{m+1}, q_{m+1})$ necessarily has a counterpart in the built sequence, since $a \in [\![\alpha_a]\!]_A$ and $\tau \notin [\![\alpha_a]\!]$. Therefore this transition cannot be an inert transition and the first case of Lemma 2 does not apply here.
- Case $\varphi = \langle\varphi_1?.\alpha_\tau\rangle\, @$. By definition, there exists an infinite sequence of states $(p_0, q_0) \xrightarrow{a_0} (p_1, q_1) \xrightarrow{a_1} \ldots (p_i, q_i) \xrightarrow{a_i} \ldots$, such that $p_0 = p, q_0 = q$, and for all $i > 0$, $(p_i, q_i) \in [\![\varphi_1]\!]$ and $a_i \in [\![\alpha_\tau]\!]$. The reasoning is similar

to the previous case. It is important to note that a sequence of infinite τ-transitions (when $a_i = \tau$ for all $i > 0$) cannot collapse into an empty sequence, as guaranteed by divbranching bisimulation.

□

B Proof of Lemma 1 (page 9)

Lemma 1-1 (Modal μ-calculus) *Assume $\varphi_i \in L_\mu^{strong}(A_s)$ $(i = 0, 1, 2)$ and $[\![\alpha_s]\!]_A \subseteq A_s$. Then the following hold:*

1. $\langle \alpha_s \rangle \varphi_0 \in L_\mu^{strong}(A_s)$
2. $[\alpha_s] \varphi_0 \in L_\mu^{strong}(A_s)$
3. $\neg\varphi_0 \in L_\mu^{strong}(A_s)$
4. $\varphi_1 \vee \varphi_2 \in L_\mu^{strong}(A_s)$
5. $\varphi_1 \wedge \varphi_2 \in L_\mu^{strong}(A_s)$
6. $\varphi_1 \Rightarrow \varphi_2 \in L_\mu^{strong}(A_s)$

Proof. Immediate from the definition of $L_\mu^{strong}(A_s)$ (see Def. 8) and the well-known identities $\varphi_1 \wedge \varphi_2 = \neg(\neg\varphi_1 \vee \neg\varphi_2)$ and $\varphi_1 \Rightarrow \varphi_2 = \neg\varphi_1 \vee \varphi_2$. □

Lemma 1-2 (Propositional Dynamic Logic) *Assume that $\varphi_0 \in L_\mu^{strong}(A_s)$, $\tau \in [\![\alpha_\tau]\!]_A$, and $\tau \notin [\![\alpha_a]\!]$. Then the following hold:*

1. $\langle \alpha_\tau^* \cdot \alpha_a \rangle \varphi_0 \in L_\mu^{strong}(A_s)$
2. $[\alpha_\tau^* \cdot \alpha_a] \varphi_0 \in L_\mu^{strong}(A_s)$
3. $\langle \alpha_\tau^* \rangle \varphi_0 \in L_\mu^{strong}(A_s)$
4. $[\alpha_\tau^*] \varphi_0 \in L_\mu^{strong}(A_s)$
5. $\langle \alpha_\tau \rangle @ \in L_\mu^{strong}(A_s)$
6. $[\alpha_\tau] \dashv \in L_\mu^{strong}(A_s)$

Proof. It was shown in [24] that these weak PDL-Δ modalities can be encoded in L_μ^{dsbr} when $\varphi_0 \in L_\mu^{dsbr}$. The proposed encoding also holds when $\varphi_0 \in L_\mu^{strong}(A_s)$. As this encoding does not add more strong modalities than already present in φ_0, then these properties belong to $L_\mu^{strong}(A_s)$. □

Lemma 1-3 (Action Computation Tree Logic) *Assume that $\varphi_i \in L_\mu^{strong}(A_s)$ $(i = 1, 2)$. Then the following hold:*

1. $\mathsf{E}(\varphi_1 \, {}_{\alpha_1}\mathsf{U} \, \varphi_2) \in L_\mu^{strong}(A_s)$
2. $\mathsf{E}(\varphi_1 \, {}_{\alpha_1}\mathsf{U}_{\alpha_2} \, \varphi_2) \in L_\mu^{strong}(A_s)$
3. $\mathsf{A}(\varphi_1 \, {}_{\alpha_1}\mathsf{U} \, \varphi_2) \in L_\mu^{strong}(A_s)$
4. $\mathsf{A}(\varphi_1 \, {}_{\alpha_1}\mathsf{U}_{\alpha_2} \, \varphi_2) \in L_\mu^{strong}(A_s)$
5. $\mathsf{AG}_{\alpha_0}(\varphi_0) \in L_\mu^{strong}(A_s)$

6. $\mathsf{EF}_{\alpha_0}(\varphi_0) \in L_\mu^{strong}(A_s)$

Proof. It was shown in [24] that the operators of ACTL\X can be encoded using the weak modalities of L_μ^{dsbr}. Therefore, the only strong modalities that remain in these formulas are those occurring in φ_1 and φ_2, which by definition match only labels in A_s. Therefore, these formulas belong to $L_\mu^{strong}(A_s)$. □

In order to prove Lemma 1-4, we need the following lemma:

Lemma 3. $\mathsf{E}(\varphi_1 \,{}_{\alpha_1}\mathsf{U}_{\alpha_2}\, \varphi_2) = \mathsf{E}(\varphi_1 \,{}_{\alpha_1}\mathsf{U}\, \varphi_1 \wedge \langle\alpha_2\rangle\, \varphi_2)$ *if* $\tau \notin [\![\alpha_2]\!]_A$.

Proof. This lemma is easily proven by showing that the modal μ-calculus definitions of the left-hand-side and right-hand-side formulas are equivalent. First note that since $\tau \notin [\![\alpha_2]\!]_A$, we have $[\![\alpha_2]\!]_A = [\![\alpha_2]\!]_A \setminus \{\tau\} = [\![\alpha_2 \wedge \neg\tau]\!]_A$ and thus $\langle\alpha_2\rangle\, \varphi = \langle\alpha_2 \wedge \neg\tau\rangle\, \varphi$ for any formula φ. We then have the following:

$$
\begin{array}{lll}
\mathsf{E}(\varphi_1 \,{}_{\alpha_1}\mathsf{U}\, \varphi_1 \wedge \langle\alpha_2\rangle\, \varphi_2) & & \\
\quad = \mu X.(\varphi_1 \wedge \langle\alpha_2\rangle\, \varphi_2) \vee (\varphi_1 \wedge \langle\alpha_1 \vee \tau\rangle\, X) & \text{by definition of } \mathsf{E}(_\,{}__\mathsf{U}\,_) \\
\quad = \mu X.\varphi_1 \wedge (\langle\alpha_2\rangle\, \varphi_2 \vee \langle\alpha_1 \vee \tau\rangle\, X) & \text{by factorization of } \varphi_1 \\
\quad = \mu X.\varphi_1 \wedge (\langle\alpha_2 \wedge \neg\tau\rangle\, \varphi_2 \vee \langle\alpha_1 \vee \tau\rangle\, X) & \text{from } \langle\alpha_2\rangle\, \varphi_2 = \langle\alpha_2 \wedge \neg\tau\rangle\, \varphi_2 \\
\quad = \mathsf{E}(\varphi_1 \,{}_{\alpha_1}\mathsf{U}_{\alpha_2}\, \varphi_2) & \text{by definition of } \mathsf{E}(_\,{}__\mathsf{U}{}__\,_)
\end{array}
$$

□

Lemma 1-4 (Computation Tree Logic) *Assume that* $\varphi_i \in L_\mu^{strong}(A_s)$ *(*$i = 0, 1, 2$*) and* $\tau \notin [\![\alpha_a]\!]_A$. *Then the following hold:*

1. $\mathsf{E}(\varphi_1\, \mathsf{U}\, \varphi_2) \in L_\mu^{strong}(A_s)$
2. $\mathsf{A}(\varphi_1\, \mathsf{U}\, \varphi_2) \in L_\mu^{strong}(A_s)$
3. $\mathsf{AG}(\varphi_0) \in L_\mu^{strong}(A_s)$
4. $\mathsf{EF}(\varphi_0) \in L_\mu^{strong}(A_s)$
5. $\mathsf{AF}(\varphi_0) \in L_\mu^{strong}(A_s)$
6. $\mathsf{EG}(\varphi_0) \in L_\mu^{strong}(A_s)$
7. $\mathsf{E}(\varphi_1\, \mathsf{W}\, \varphi_2) \in L_\mu^{strong}(A_s)$
8. $\mathsf{A}(\varphi_1\, \mathsf{W}\, \varphi_2) \in L_\mu^{strong}(A_s)$
9. $\mathsf{A}([\alpha_a]\, \varphi_1\, \mathsf{U}\, \varphi_2) \in L_\mu^{strong}(A_s)$
10. $\mathsf{A}([\alpha_a]\, \varphi_1\, \mathsf{W}\, \varphi_2) \in L_\mu^{strong}(A_s)$
11. $\mathsf{AG}([\alpha_a]\, \varphi_0) \in L_\mu^{strong}(A_s)$
12. $\mathsf{EF}(\langle\alpha_a\rangle\, \varphi_0) \in L_\mu^{strong}(A_s)$
13. $\mathsf{AG}(\varphi_1 \vee [\alpha_a]\, \varphi_2) \in L_\mu^{strong}(A_s)$
14. $\mathsf{EF}(\varphi_1 \wedge \langle\alpha_a\rangle\, \varphi_2) \in L_\mu^{strong}(A_s)$

Proof. CTL operators are defined on LTS in terms of ACTL\X operators as follows:

$$\begin{aligned}
\mathsf{E}(\varphi_1 \,\mathsf{U}\, \varphi_2) &= \mathsf{E}(\varphi_1 \,{}_{\mathsf{true}}\mathsf{U}\, \varphi_2) && \text{(Ctl1)}\\
\mathsf{A}(\varphi_1 \,\mathsf{U}\, \varphi_2) &= \mathsf{A}(\varphi_1 \,{}_{\mathsf{true}}\mathsf{U}\, \varphi_2) && \text{(Ctl2)}\\
\mathsf{EF}(\varphi_0) &= \mathsf{E}(\mathsf{true} \,{}_{\mathsf{true}}\mathsf{U}\, \varphi_0) && \text{(Ctl3)}\\
\mathsf{AG}(\varphi_0) &= \neg\mathsf{E}(\mathsf{true} \,{}_{\mathsf{true}}\mathsf{U}\, \neg\varphi_0) && \text{(Ctl4)}\\
\mathsf{AF}(\varphi_0) &= \mathsf{A}(\mathsf{true} \,{}_{\mathsf{true}}\mathsf{U}\, \varphi) && \text{(Ctl5)}\\
\mathsf{EG}(\varphi_0) &= \neg\mathsf{A}(\mathsf{true} \,{}_{\mathsf{true}}\mathsf{U}\, \neg\varphi) && \text{(Ctl6)}\\
\mathsf{E}(\varphi_1 \,\mathsf{W}\, \varphi_2) &= \mathsf{E}(\varphi_1 \,\mathsf{U}\, \varphi_2) \vee \mathsf{EG}(\varphi_1) && \text{(Ctl7)}\\
\mathsf{A}(\varphi_1 \,\mathsf{W}\, \varphi_2) &= \neg\mathsf{E}(\neg\varphi_2 \,\mathsf{U}\, \neg\varphi_1 \wedge \neg\varphi_2) && \text{(Ctl8)}
\end{aligned}$$

In addition, there is the following identity, which is easily proven by showing that the modal μ-calculus definitions of the left-hand-side and right-hand-side formulas are equivalent:

$$\mathsf{A}(\varphi_1 \,\mathsf{U}\, \varphi_2) = \neg\mathsf{E}(\neg\varphi_2 \,\mathsf{W}\, \neg\varphi_1 \wedge \neg\varphi_2) \quad \text{(Ctl9)}$$

We also use the following standard identity of the modal μ-calculus:

$$\neg[\alpha]\,\varphi_0 = \langle\alpha\rangle\,\neg\varphi_0 \qquad \text{(Mcl1)}$$

Items 1 to 8 are a direct consequence of Lemma 1-3 and the fact that these operators can be translated to ACTL$\setminus X$, as shown by equations (Ctl1) to (Ctl8). The proof of Items 9 to 14 follows:

9. $\mathsf{A}([\alpha_a]\,\varphi_1 \,\mathsf{U}\, \varphi_2)$
   $= \neg\mathsf{E}(\neg\varphi_2 \,\mathsf{W}\, \neg[\alpha_a]\,\varphi_1 \wedge \neg\varphi_2)$ By (Ctl9)
   $= \neg(\mathsf{E}(\neg\varphi_2 \,\mathsf{U}\, \neg[\alpha_a]\,\varphi_1 \wedge \neg\varphi_2) \vee \mathsf{EG}(\neg\varphi_2))$ By (Ctl7)
   $= \neg(\mathsf{E}(\neg\varphi_2 \,\mathsf{U}\, \langle\alpha_a\rangle\,\neg\varphi_1 \wedge \neg\varphi_2) \vee \mathsf{EG}(\neg\varphi_2))$ By (Mcl1)
   $= \neg(\mathsf{E}(\neg\varphi_2 \,{}_{\mathsf{true}}\mathsf{U}\, \langle\alpha_a\rangle\,\neg\varphi_1 \wedge \neg\varphi_2) \vee \mathsf{EG}(\neg\varphi_2))$ By (Ctl1)
   $= \neg(\mathsf{E}(\neg\varphi_2 \,{}_{\mathsf{true}}\mathsf{U}_{\alpha_a}\, \neg\varphi_1) \vee \mathsf{EG}(\neg\varphi_2))$ By Lemma 3 and $\tau \notin [\![\alpha_a]\!]_A$
   which is in $L_\mu^{strong}(A_s)$ following Item 6 and Lemmas 1-1 and 1-3.
10. $\mathsf{A}([\alpha_a]\,\varphi_1 \,\mathsf{W}\, \varphi_2)$
    $= \neg\mathsf{E}(\neg\varphi_2 \,\mathsf{U}\, \neg[\alpha_a]\,\varphi_1 \wedge \neg\varphi_2)$ By (Ctl8)
    $= \neg\mathsf{E}(\neg\varphi_2 \,\mathsf{U}\, \langle\alpha_a\rangle\,\neg\varphi_1 \wedge \neg\varphi_2)$ By (Mcl1)
    $= \neg\mathsf{E}(\neg\varphi_2 \,{}_{\mathsf{true}}\mathsf{U}\, \langle\alpha_a\rangle\,\neg\varphi_1 \wedge \neg\varphi_2)$ By (Ctl1)
    $= \neg\mathsf{E}(\neg\varphi_2 \,{}_{\mathsf{true}}\mathsf{U}_{\alpha_a}\, \neg\varphi_1)$ By Lemma 3 and $\tau \notin [\![\alpha_a]\!]_A$
    which is in $L_\mu^{strong}(A_s)$ following Lemmas 1-1 and 1-3.
11. $\mathsf{AG}([\alpha_a]\,\varphi_0)$
    $= \neg\mathsf{E}(\mathsf{true} \,{}_{\mathsf{true}}\mathsf{U}\, \neg[\alpha_a]\,\varphi_0)$ By (Ctl4)
    $= \neg\mathsf{E}(\mathsf{true} \,{}_{\mathsf{true}}\mathsf{U}\, \langle\alpha_a\rangle\,\neg\varphi_0)$ By (Mcl1)
    $= \neg\mathsf{E}(\mathsf{true} \,{}_{\mathsf{true}}\mathsf{U}_{\alpha_a}\, \neg\varphi_0)$ By Lemma 3 and $\tau \notin [\![\alpha_a]\!]_A$
    which is in $L_\mu^{strong}(A_s)$ following Lemmas 1-1 and 1-3. This is also a consequence of $\mathsf{AG}(\varphi_1 \vee [\alpha_a]\,\varphi_2) \in L_\mu^{strong}(A_s)$ (shown below), replacing φ_1 by false and φ_2 by φ_0.

12. $\mathsf{EF}(\langle \alpha_a \rangle \varphi_0)$
    $= \mathsf{E}(\mathsf{true}\, \mathsf{U}\, \langle \alpha_a \rangle \varphi_0)$ By (Ctl3)
    $= \mathsf{E}(\mathsf{true}\, \mathsf{U}_{\alpha_a}\, \varphi_0)$ By Lemma 3 and $\tau \notin [\![\alpha_a]\!]_A$
    which is in $L_\mu^{strong}(A_s)$ following Lemmas 1-1 and 1-3. This is also a consequence of $\mathsf{EF}(\varphi_1 \wedge \langle \alpha_a \rangle \varphi_2) \in L_\mu^{strong}(A_s)$ (shown below), replacing φ_1 by true and φ_2 by φ_0.
13. $\mathsf{AG}(\varphi_1 \vee [\alpha_a] \varphi_2) = \neg \mathsf{EF}(\neg \varphi_1 \wedge \langle \alpha_a \rangle \neg \varphi_2)$, which is shown to belong to $L_\mu^{strong}(A_s)$ below.
14. $\mathsf{EF}(\varphi_1 \wedge \langle \alpha_a \rangle \varphi_2)$ is semantically equivalent to $\mathsf{EF}(\langle (\varphi_1?.\mathsf{true})^*.\varphi_1?.\alpha_a \rangle \varphi_2)$, which belongs to $L_\mu^{strong}(A_s)$. Indeed, if $\mathsf{EF}(\varphi_1 \wedge \langle \alpha_a \rangle \varphi_2)$ is true, then there is a reachable state satisfying $\varphi_1 \wedge \langle \alpha_a \rangle \varphi_2$, by definition of EF. This state also satisfies $\langle (\varphi_1?.\mathsf{true})^*.\varphi_1?.\alpha_a \rangle \varphi_2$, after an empty sequence of steps matching $(\varphi_1?.\mathsf{true})^*$. On the other hand, if there is a reachable state satisfying $\langle (\varphi_1?.\mathsf{true})^*.\varphi_1?.\alpha_a \rangle \varphi_2$, then there is a reachable state satisfying both φ_1 and $\langle \alpha_a \rangle \varphi_2$, by definition of the weak modality, i.e, the LTS satisfies $\mathsf{EF}(\varphi_1 \wedge \langle \alpha_a \rangle \varphi_2)$.

□

C Detailed performance data of the TFTP case-study

Table 5 presents the detailed performance data obtained on the TFTP case-study, which served as input to generate the curves of Figure 2 (page 12). Parameters P, S, and I correspond respectively to the property number, the scenario, and the TFTP instance. LTS sizes are given in kilostates, memory in megabytes, and time in seconds. Time does not include LTS generation of the component processes from their LNT specification, which is quite fast (a few seconds) and shared by both the mono-bisimulation approach and the combined bisimulations approach. The column "Ratio Strong/Combined" shows the gain obtained by applying the combined bisimulations approach rather than the mono-bisimulation approach with respect to largest LTS size, memory peak, and time.

			Strong bisimulation				Combined bisimulations				Ratio Strong/Combined			
			Kstates		verif.		Kstates		verif.		Kstates		verif.	
P	S	I	largest	final	MB	sec.	largest	final	MB	sec.	largest	final	MB	sec.
08	A	A	1783	1197	136	66	746	503	60	32	2.4	2.4	2.3	2.1
14	A	A	1783	1495	138	83	770	639	62	31	2.3	2.3	2.2	2.7
08	A	B	1783	1295	143	91	347	281	30	17	5.1	4.6	4.8	5.4
17	A	B	1783	1191	138	76	1000	694	79	44	1.8	1.7	1.7	1.7
08	B	A	792	636	62	38	347	289	30	18	2.3	2.2	2.1	2.1
08	B	B	792	707	62	40	174	130	17	13	4.6	5.4	3.6	3.1
14	B	B	792	720	65	47	174	134	17	14	4.6	5.4	3.8	3.4
16	B	B	792	613	63	35	456	356	38	22	1.7	1.7	1.7	1.6
08	C	A	30926	16924	2393	2274	8448	4471	660	575	3.7	3.8	3.4	4.0
09	C	A	30965	16935	2540	2305	7940	7516	636	666	3.9	2.3	4.0	3.5
14	C	A	30965	22772	2417	3033	8679	6141	688	706	3.6	3.7	3.5	4.3
17	C	A	30992	16935	2527	1873	19269	10833	1502	1164	1.6	1.6	1.7	1.6
08	C	B	30926	16853	2391	2257	7995	4176	617	417	3.9	4.0	3.9	5.4
09	C	B	30953	16998	2538	2271	7326	6931	589	502	4.2	2.5	4.3	4.5
14	C	B	30953	22710	2410	2635	8146	5683	637	601	3.8	4.0	3.8	4.4
17	C	B	30992	16998	2529	1913	18033	10173	1406	1155	1.7	1.7	1.8	1.7
08	D	A	36447	19136	2934	2703	9655	5824	783	623	3.8	3.3	3.7	4.3
09	D	A	36447	18731	3152	2903	9636	8607	808	865	3.8	2.1	3.9	3.4
17	D	A	36447	18731	2971	2262	22203	11926	1793	1545	1.6	1.6	1.7	1.5
08	D	B	36447	19062	2933	2787	7405	3498	568	464	4.9	5.4	5.2	6.0
09	D	B	36478	19243	3110	2650	5495	5153	443	412	6.6	3.7	7.0	6.4
14	D	B	36478	22237	3002	3264	7533	3986	588	589	4.8	5.6	5.1	5.5
16	D	B	36478	19243	3161	2303	15576	7830	1204	1232	2.3	2.5	2.6	1.9
08	E	A	17035	10646	1315	1046	4970	3257	388	287	3.4	3.3	3.4	3.6
09	E	A	17035	10375	1397	1312	4856	4620	392	310	3.5	2.2	3.6	4.2
14	E	A	17035	12759	1327	1346	5081	3983	400	384	3.4	3.2	3.3	3.5
16	E	A	17035	10376	1399	897	10723	6643	832	736	1.6	1.6	1.7	1.2
08	E	B	17035	10529	1318	1177	4735	3033	369	253	3.6	3.5	3.6	4.7
09	E	B	17035	10401	1402	1130	4471	4267	363	320	3.8	2.4	3.9	3.5
14	E	B	17035	12725	1329	1342	4810	3714	377	283	3.5	3.4	3.5	4.7
16	E	B	17035	10402	1396	1033	9987	6202	773	793	1.7	1.7	1.8	1.3

Table 5. Results of the TFTP case-study