



HAL
open science

Modular Analysis of Tree-Topology Models

Jaime Arias, Michal Knapik, Wojciech Penczek, Laure Petrucci

► **To cite this version:**

Jaime Arias, Michal Knapik, Wojciech Penczek, Laure Petrucci. Modular Analysis of Tree-Topology Models. Formal Methods and Software Engineering, ICFEM, Oct 2022, Madrid, Spain. pp.36-53, 10.1007/978-3-031-17244-1_3. hal-03811772

HAL Id: hal-03811772

<https://hal.science/hal-03811772>

Submitted on 22 Oct 2022

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.

Modular Analysis of Tree-Topology Models^{*}

Jaime Arias¹[0000-0003-3019-4902], Michał Knapik²[0000-0003-3259-9786],
Wojciech Penczek²[0000-0001-6477-4863], and Laure
Petrucci¹[0000-0003-3154-5268]

¹ LIPN, CNRS UMR 7030, Université Sorbonne Paris Nord, 99 av. J-B. Clément,
93430 Villetaneuse, France {first.last}@lipn.univ-paris13.fr

² Institute of Computer Science, Polish Academy of Sciences, Jana Kazimierza 5,
01-248 Warsaw, Poland {first.last}@ipipan.waw.pl

Abstract. Networks of automata that synchronise over shared actions are organised according to a graph synchronisation topology. In this topology two automata are connected if they can jointly execute some action. We present a very effective reduction for networks with tree-like synchronisation topologies such that all automata after synchronising with their parents can execute only local (non-synchronising) actions: forever or until *resetting*, i.e. entering the initial state. We show that the reduction preserves reachability, but not liveness. This construction is extended to tree-like topologies of arbitrary automata and investigated experimentally.

Keywords: Model Checking · Networks of Synchronising Automata · State Space Reduction · Tree-Like Synchronisation Topologies.

1 Introduction

Networks of various types of finite automata (or labelled transition systems, LTSs for short) [10,14] are a popular choice of formalism when modelling complex systems such as protocols. In this approach, the components of the system under investigation are abstracted as state machines and the behaviour of the entire system is captured by their synchronised product. However, the cost of computing the synchronised product can be prohibitive: in practice the size of the state space grows exponentially with the number of the sub-modules. This observation led to an extensive research into how to analyse and represent models that consist of interacting components without building the entire state space.

In this paper, we tackle the problem of computing a compact representation of the state space of the entire product of networks of LTSs that exhibit tree-like synchronisation topologies. Intuitively, this means a graph, whose edges depict synchronisation between LTSs over common actions of a system, is a

^{*} The authors acknowledge the support of CNRS and PAN, under the IEA project MoSART, and of NCBR Poland and FNR Luxembourg, under the PolLux/FNR-CORE project STV (POLLUX-VII/1/2019).

tree. Examples of systems having such a synchronisation topology are e.g. attack-defense trees (ADT) [17,7,6,9], hierarchical [4], broadcast [12], multimedia [8], and workflow models [3].

This structure of communication is quite natural also in the system design where circular dependencies are sometimes treated as anti-patterns and cause problems such as deadlocks.

We design algorithms that compute a compact representation of the synchronised product bottom-up, focusing at each level only on the synchronisations between the pairs consisting of a parent and one of its children. As it turns out, the notion of memory is crucial to the size of the constructed LTS (called here the *sum-of-squares product*). Namely, if every component resets (i.e. returns to its initial state) after synchronising with its parent or enters a deadlock, then no information about its post-synchronisation states needs to be preserved. In this case, the sum-of-squares product is quite small and can be computed efficiently. In case some LTSs do not reset after synchronising with their parents, an additional memory gadget is needed. Then, the general sum-of-squares product also preserves reachability, but a reduction or good performance is not guaranteed. The burden of recording the post-synchronisation state of the components can even lead to substantial blowup of the state space.

Outline of the Paper

In Section 2 we provide the basic definitions of synchronising LTSs, live-reset and sync-deadlock LTSs, and synchronisation topologies. A bottom-up reduction for the networks of live-reset and sync-deadlock LTSs is presented in Section 3. We show how to transform these networks in a reachability-preserving way into a new, smaller model called the sum-of-squares product. The key idea is that the new model traces the interactions between the components and their parents, followed by upstream synchronisations. This bottom-up reduction is extended in Section 4 to a wider case of tree-like synchronisation topologies with general LTSs. Section 5 investigates experimentally the effectiveness of the reductions on several scalable examples and many random cases using a novel open-source tool [2]. Section 6 draws conclusions and discusses future work.

2 Tree Synchronisation Systems

In this section we recall the notions of networks of Labelled Transition Systems and their synchronisation topologies. We also introduce and explain the restrictions on the models assumed in the next section. In what follows \mathcal{PV} denotes the set of propositions.

Definition 1 (Labelled Transition System). A Labelled Transition System (LTS) is a tuple $\mathcal{M} = \langle \mathcal{S}, s^I, Acts, \rightarrow, \mathcal{L} \rangle$ where:

1. \mathcal{S} is a finite set of states and $s^I \in \mathcal{S}$ the initial state;

2. $Acts$ is a finite set of action names;
3. $\rightarrow \subseteq \mathcal{S} \times Acts \times \mathcal{S}$ is a transition relation;
4. $\mathcal{L}: \mathcal{S} \rightarrow 2^{\mathcal{P}\mathcal{V}}$ is a labelling function.

We usually write $s \xrightarrow{act} s'$ instead of $(s, act, s') \in \rightarrow$. We also denote $acts(\mathcal{M}) = Acts$ and $states(\mathcal{M}) = \mathcal{S}$. A run in \mathcal{LTS} \mathcal{M} is an infinite sequence of states and actions $\rho = s^0 act^0 s^1 act^1 \dots$ s.t. $s^i \xrightarrow{act^i} s^{i+1}$ for all $i \geq 0$. By $Runs(\mathcal{M}, s)$ we denote the set of all the runs starting from state $s \in \mathcal{S}$. If s is the initial state, we write $Runs(\mathcal{M})$ instead of $Runs(\mathcal{M}, s^I)$.

2.1 \mathcal{LTS} Networks and Synchronisation Topologies

Model checkers such as SPIN [13], UPPAAL [11], and IMITATOR [5] usually expect the systems to be described in a form of interacting modules. Synchronisation over common actions [10] (or *channels*) is a popular primitive that enables such an interaction.

Definition 2 (Asynchronous Product). Let $\mathcal{M}_i = \langle \mathcal{S}_i, s_i^I, Acts_i, \rightarrow_i, \mathcal{L}_i \rangle$ be \mathcal{LTS} , for $i \in \{1, 2\}$. The asynchronous product of \mathcal{M}_1 and \mathcal{M}_2 is the \mathcal{LTS} $\mathcal{M}_1 || \mathcal{M}_2 = \langle \mathcal{S}_1 \times \mathcal{S}_2, (s_1^I, s_2^I), Acts_1 \cup Acts_2, \rightarrow, \mathcal{L}_{1,2} \rangle$ s.t. $\mathcal{L}_{1,2}((s_1, s_2)) = \mathcal{L}_1(s_1) \cup \mathcal{L}_2(s_2)$ for all $(s_1, s_2) \in \mathcal{S}_1 \times \mathcal{S}_2$ and the transition rule is defined as follows:

$$\frac{act \in Acts_1 \setminus Acts_2 \wedge s_1 \xrightarrow{act}_1 s'_1}{(s_1, s_2) \xrightarrow{act} (s'_1, s_2)} \quad \frac{act \in Acts_2 \setminus Acts_1 \wedge s_2 \xrightarrow{act}_2 s'_2}{(s_1, s_2) \xrightarrow{act} (s_1, s'_2)}$$

$$\frac{act \in Acts_1 \cap Acts_2 \wedge s_1 \xrightarrow{act}_1 s'_1 \wedge s_2 \xrightarrow{act}_2 s'_2}{(s_1, s_2) \xrightarrow{act} (s'_1, s'_2)}$$

The above definition is naturally extended to an arbitrary number of components. We sometimes write $||_{i=1}^n \mathcal{M}_i$ instead of $\mathcal{M}_1 || \dots || \mathcal{M}_n$. If s is a state of $\mathcal{M}_1 || \dots || \mathcal{M}_n$, then by $s_{\mathcal{M}_i}$ we denote its component corresponding to \mathcal{M}_i .

The synchronisation topology [17] is an undirected graph that records how LTSs synchronise with one another.

Definition 3 (Synchronisation Topology). A synchronisation topology (\mathcal{ST}) is a tuple $\mathcal{G} = \langle Net, \mathcal{T} \rangle$, where $Net = \{\mathcal{M}_i\}_{i=1}^n$ is a set of \mathcal{LTS} s for $1 \leq i \leq n$, and $\mathcal{T} \subseteq Net \times Net$ is s.t. $(\mathcal{M}_i, \mathcal{M}_j) \in \mathcal{T}$ iff $i \neq j$ and $Acts_i \cap Acts_j \neq \emptyset$.

Note that \mathcal{T} is induced by Net . Thus, with a slight notational abuse we sometimes treat \mathcal{G} as Net . Moreover, we write $acts(\mathcal{G}) = \bigcup_{i=1}^n acts(\mathcal{M}_i)$.

Definition 4 (Tree Synchronisation Topology). A synchronisation topology \mathcal{G} s.t. \mathcal{T} is a tree rooted in $root(\mathcal{G})$ is called tree synchronisation topology.

It should be noted that $root(\mathcal{G})$ is not always uniquely induced by \mathcal{G} , the root is thus a part of the signature of tree synchronisation topology.

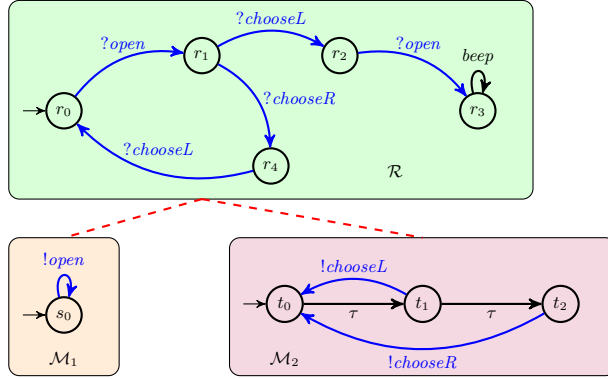


Fig. 1. A simple tree synchronisation topology \mathcal{G}_x .

Let us fix a tree synchronisation topology $\mathcal{G} = \langle \text{Net}, \mathcal{T}, \text{root}(\mathcal{G}) \rangle$. For each $\mathcal{M} \in \text{Net}$ by $\text{parent}(\mathcal{M})$ we denote the parent of \mathcal{M} in \mathcal{T} (we assume that $\text{parent}(\text{root}(\mathcal{G})) = \emptyset$). By $\text{children}(\mathcal{M})$ we mean the set of the children of \mathcal{M} . By $\text{upacts}(\mathcal{M})$ (resp. $\text{downacts}(\mathcal{M})$) we denote the set of actions via which \mathcal{M} synchronises with its parent (resp. children). For each $\text{act} \in \text{downacts}(\mathcal{M})$, $\text{snd}(\mathcal{M}, \text{act})$ denotes the component $\mathcal{M}' \in \text{children}(\mathcal{M})$ s.t. $\text{act} \in \text{upacts}(\mathcal{M}')$. Thus, $\text{snd}(\mathcal{M}, \text{act})$ is the child of \mathcal{M} that synchronises with \mathcal{M} over act . If \mathcal{M} is clear from the context, we simply write $\text{snd}(\text{act})$. The *local* unsynchronised actions of \mathcal{M} are defined as $\text{locacts}(\mathcal{M}) = \text{acts}(\mathcal{M}) \setminus (\text{downacts}(\mathcal{M}) \cup \text{upacts}(\mathcal{M}))$.

For brevity, whenever we refer to a state, action label, transition, or state space of \mathcal{G} we mean a state, action label, transition, or state space of $\prod_{i=1}^n \mathcal{M}_i$. We also extend the notion of runs to synchronisation topologies: $\text{Runs}(\mathcal{G}, s) = \text{Runs}(\prod_{i=1}^n \mathcal{M}_i, s)$ for each $s \in \text{states}(\prod_{i=1}^n \mathcal{M}_i)$. Moreover, for each $F \subseteq \text{acts}(\mathcal{G})$ we introduce the notion of *one-shot* F runs, $\text{Runs}_F(\mathcal{G}, s)$, as the set of all the runs in $\text{Runs}(\mathcal{G}, s)$ along which each action from F appears at most once.

Example 1. Figure 1 presents a small tree $\mathcal{ST} \mathcal{G}_x$ with the root \mathcal{R} and two children \mathcal{M}_1 and \mathcal{M}_2 . The auxiliary symbols $!/?$ are syntactic sugar, used to distinguish between upacts and downacts . Here, $\text{upacts}(\mathcal{R}) = \emptyset$, $\text{downacts}(\mathcal{R}) = \{\text{open}, \text{chooseL}, \text{chooseR}\}$, and $\text{locacts}(\mathcal{R}) = \{\text{beep}\}$. Similarly, $\text{upacts}(\mathcal{M}_1) = \{\text{open}\}$, $\text{upacts}(\mathcal{M}_2) = \{\text{chooseL}, \text{chooseR}\}$, $\text{downacts}(\mathcal{M}_1) = \text{locacts}(\mathcal{M}_1) = \text{downacts}(\mathcal{M}_2) = \emptyset$, and $\text{locacts}(\mathcal{M}_2) = \tau$.

We start by dealing with networks whose components all share a similar, simple structure.

Definition 5 (Live-Reset LTS). A LTS \mathcal{M} with initial state $s_{\mathcal{M}}^I$ is live-reset if for each run $\rho = s^0 \text{act}^0 s^1 \text{act}^1 \dots : \forall i \in \mathbb{N}$ if $\text{act}^i \in \text{upacts}(\mathcal{M})$, then $s^{i+1} = s_{\mathcal{M}}^I$.

Definition 6 (Sync-Deadlock LTS). A LTS \mathcal{M} is sync-deadlock if for each run $\rho = s^0 act^0 s^1 act^1 \dots: \forall i \in \mathbb{N}$ if $act^i \in upacts(\mathcal{M})$, then for each $j > i$, $act^j \in locacts(\mathcal{M})$.

LTSs that are either live-reset or sync-deadlock are said to be *sync-memoryless*³. Intuitively, after synchronising with its parent a sync-memoryless LTS either immediately enters its initial state (if live-reset) or executes only local actions (if sync-deadlock).

If every LTS of \mathcal{ST} \mathcal{G} is live-reset (resp. sync-deadlock), then we say that \mathcal{G} is live-reset (resp. sync-deadlock). It is easy to see that the tree \mathcal{ST} in Figure 1 is live-reset.

Definition 7 (Reachability and liveness). For each $p \in \mathcal{PV}$ we write $\mathcal{G} \models EFP$ (resp. $\mathcal{G} \models EGp$) iff there exists $\rho \in Runs(\mathcal{G})$ s.t. $\rho = s^0 act^0 s^1 act^1 \dots$ and $p \in \mathcal{L}(s^i)$ for some (resp. for all) $i \in \mathbb{N}$.

By replacing \models with \models_F and $Runs$ with $Runs_F$ in Definition 7 we obtain the notion of *one-shot F-reachability* and the dual of liveness. Both EF and EG are Computation Tree Logic (CTL) modalities [10]. If $\mathcal{G} \models EFP$, then we say that p is reachable in \mathcal{G} from the initial state.

Definition 8. Let $\mathcal{N} \subseteq Net$ and ρ^* be a prefix of some $\rho \in Runs(\mathcal{G})$ s.t. $\rho^* = s^0 act^0 s^1 act^1 \dots$. By $\rho^* \downarrow (\mathcal{N})$ we denote the projection of ρ^* to the asynchronous product of LTSs in \mathcal{N} , i.e. the result of transforming ρ^* by (1) firstly, projecting each s^i on the LTSs in \mathcal{N} ; (2) secondly, removing the actions that do not belong to $\bigcup_{\mathcal{M} \in \mathcal{N}} acts(\mathcal{M})$, together with their destinations.

Intuitively, $\rho^* \downarrow (\mathcal{N})$ contains the parts of the global states of ρ^* that belong to $\|\mathcal{M} \in \mathcal{N}\mathcal{M}$ and actions executed by some $\mathcal{M} \in \mathcal{N}$. It is not difficult to show that $\rho^* \downarrow (\mathcal{N})$ does not need to be a valid run of $\|\mathcal{M} \in \mathcal{N}\mathcal{M}$.

Example 2. Consider a sequence: $\eta = (r_0, s_0, t_0) \tau (r_0, s_0, t_1) \tau (r_0, s_0, t_2) open (r_1, s_0, t_2) chooseR (r_4, s_0, t_0) \tau (r_4, s_0, t_1) chooseL (r_0, s_0, t_0)$. Here, we have $\eta \downarrow (\{\mathcal{R}, \mathcal{M}_1\}) = (r_0, s_0) open(r_1, s_0) chooseR(r_4, s_0) chooseL(r_0, s_0)$.

3 Compact Representations of State Spaces of Live-Reset and Sync-Deadlock Trees

In this section we show how to generate compact representations of state spaces of sync-memoryless tree topologies preserving reachability. The procedure is presented in two steps. We start with the case of two-level trees. Then, we modify the construction to deal with trees of arbitrary height in a bottom-up manner.

³ The family of sync-memoryless LTSs can in the future be extended beyond these two classes.

3.1 Constructions for Two-level Trees

Throughout this subsection let \mathcal{G} be a sync-memoryless tree \mathcal{ST} with components $Net = \{\mathcal{R}, \mathcal{M}_1, \dots, \mathcal{M}_n\}$ s.t. $root(\mathcal{G}) = \mathcal{R}$ and $children(\mathcal{R}) = \{\mathcal{M}_1, \dots, \mathcal{M}_n\}$. Moreover, let $\mathcal{R} = \langle \mathcal{S}_{\mathcal{R}}, s_{\mathcal{R}}^I, Acts_{\mathcal{R}}, \rightarrow_{\mathcal{R}}, \mathcal{L}_{\mathcal{R}} \rangle$ and $\mathcal{M}_i = \langle \mathcal{S}_i, s_i^I, Acts_i, \rightarrow_i, \mathcal{L}_i \rangle$, for $i \in \{1, \dots, n\}$. We employ the observations on the nature of synchronisations with sync-memoryless components in the following definition.

Definition 9 (Sum-of-squares Product). Define $\mathcal{SQ}^u(\mathcal{G}) = \langle \mathcal{S}_{sq}^u, s_{sq}^I, Acts_{sq}, \rightarrow_{sq}, \mathcal{L}_{sq} \rangle$ as an \mathcal{LTS} s.t.:

1. $\mathcal{S}_{sq}^u = \bigcup_{i=1}^n (\mathcal{S}_i \times \mathcal{S}_{\mathcal{R}}) \cup \{s_{sq}^I\}$.
2. $s_{sq}^I \notin \mathcal{S}_{sq}^u$ is a fresh initial state.
3. $Acts_{sq} = acts(\mathcal{G}) \cup \{\epsilon\}$, where $\epsilon \notin acts(\mathcal{G})$ is a fresh, silent action.
4. The transition relation \rightarrow_{sq} is defined as follows:
 - (a) $s_{sq}^I \xrightarrow{\epsilon}_{sq} (s_i^I, s_{\mathcal{R}}^I)$, for all $i \in \{1, \dots, n\}$;
 - (b) if $(s_i, s_{\mathcal{R}}) \xrightarrow{act}_{sq} (s'_i, s'_{\mathcal{R}})$ is a transition in $\mathcal{M}_i || \mathcal{R}$, then also $(s_i, s_{\mathcal{R}}) \xrightarrow{act}_{sq} (s'_i, s'_{\mathcal{R}})$, for all $i \in \{1, \dots, n\}$;
 - (c) if \mathcal{G} is live-reset and $(s_i, s_{\mathcal{R}}) \xrightarrow{act}_{sq} (s_i^I, s'_{\mathcal{R}})$ is a transition in $\mathcal{M}_i || \mathcal{R}$, then $(s_i, s_{\mathcal{R}}) \xrightarrow{act}_{sq} (s_j^I, s'_{\mathcal{R}})$ for all $j \in \{1, \dots, n\} \setminus \{i\}$;
 - (d) if \mathcal{G} is sync-deadlock and $(s_i, s_{\mathcal{R}}) \xrightarrow{act}_{sq} (s'_i, s'_{\mathcal{R}})$ is a synchronised transition in $\mathcal{M}_i || \mathcal{R}$, then $(s_i, s_{\mathcal{R}}) \xrightarrow{act}_{sq} (s_j^I, s'_{\mathcal{R}})$, for all $j \in \{1, \dots, n\} \setminus \{i\}$.
5. $\mathcal{L}_{sq}(s_i, s_{\mathcal{R}}) = \mathcal{L}_i(s_i) \cup \mathcal{L}_{\mathcal{R}}(s_{\mathcal{R}}) \cup \bigcup_{j \neq i} \mathcal{L}_j(s_j^I)$, for each $(s_i, s_{\mathcal{R}}) \in \mathcal{S}_{sq}^u$.

We call $\mathcal{SQ}^u(\mathcal{G})$ the Sum-of-squares Product of \mathcal{G} .

Intuitively, $\mathcal{SQ}^u(\mathcal{G})$ at any given moment traces only the interactions between the root and one of its children. Item 4a of Definition 9 introduces the new initial state connected via ϵ -transitions with the initial states of each square product $\mathcal{M}_i || \mathcal{R}$. Item 4b ensures that the square product preserves the local and synchronised actions of the root and each child. Item 4c means that after resetting each component \mathcal{M}_i can release control to another module, for live-reset topologies. Item 4d serves a similar purpose for sync-deadlock topologies.

Example 3. Fig. 2 presents the sum-of-squares product $\mathcal{SQ}^u(\mathcal{G}_x)$ for the small tree \mathcal{ST} of Example 1. The fresh initial state s_{sq}^I is coloured green, the yellow box surrounds the square product $\mathcal{M}_1 || \mathcal{R}$, and the pink box surrounds the square product $\mathcal{M}_2 || \mathcal{R}$. The red colour of some states is explained in Section 3.2. The arcs that correspond to the transitions synchronised between a child and the root are coloured red if the control switches between components (according to Item 4c of Definition 9) and blue otherwise. The local transitions are black.

Theorem 1 (Sum-of-squares Product Preserves Reachability). Let \mathcal{G} be a sync-memoryless two-level tree \mathcal{ST} with root \mathcal{R} and $p \in \mathcal{PV}$. If \mathcal{G} is live-reset, then $\mathcal{G} \models EFp$ iff $\mathcal{SQ}^u(\mathcal{G}) \models EFp$. If \mathcal{G} is sync-deadlock and $F = downacts(\mathcal{R})$, then $\mathcal{G} \models_F EFp$ iff $\mathcal{SQ}^u(\mathcal{G}) \models_F EFp$.

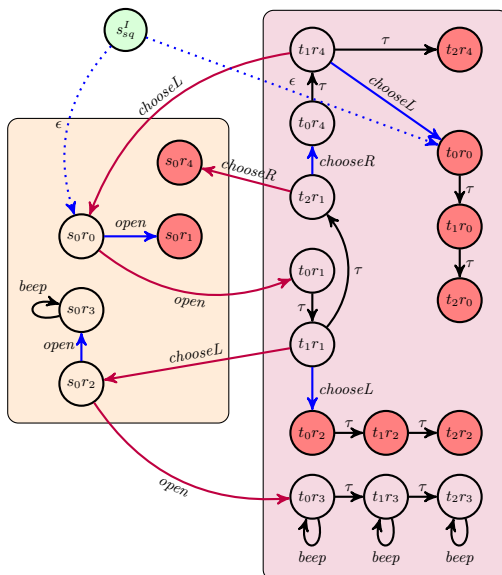


Fig. 2. The sum-of-squares product of the tree \mathcal{ST} of Example 1.

Proof. We only deal with the case of live-reset topology; the case of sync-deadlock follows similarly.

Recall that we assume $Net = \{\mathcal{R}, \mathcal{M}_1, \dots, \mathcal{M}_n\}$ with root \mathcal{R} and children $\{\mathcal{M}_i\}_{i=1}^n$. Let $\mathcal{G} \models EFP$ and $\rho = s^0 act^0 s^1 act^1 \dots$ be a run of \mathcal{G} s.t. $p \in \mathcal{L}(s^i)$ for some $i \in \mathbb{N}$. If ρ contains only local actions, then for any of the components $C \in Net$ each of its local states visited along ρ can be reached by executing in $\mathcal{SQ}^u(\mathcal{G})$ firstly the ϵ -action (cf. Item 4a of Definition 9) and then local actions of C found in ρ (cf. Item 4b of Definition 9). Moreover, consider a situation where a non-root component C synchronises with the root only a finite number of times. Now, each local state s of C that is reachable along ρ only after executing the final synchronising action of C can be reached in $\mathcal{SQ}^u(\mathcal{G})$ using only ϵ - and local actions. Such a run in $\mathcal{SQ}^u(\mathcal{G})$ reaching s can be built as described previously due to the fact that the final synchronisation resets the component. Also, let $\hat{\rho}$ be the result of iterative removal from ρ of all the local actions of any non-root component C that appear after its last synchronising action together with their targets, and replacing along the run all the further local states of C with s^I_C . It is not difficult to see that $\hat{\rho}$ visits the same local states of the root that are visited along ρ and all the local states of each non-root component C visited before the final synchronising action of C is executed.

We can therefore further assume (\star) that ρ contains at least one synchronising action and can be represented as $\rho = \alpha_1 F_1 \alpha_2 F_2 \dots$, where for each $i \in \mathbb{N}$ there exist $j, k \in \mathbb{N}$ such that $\alpha_i = s^j act^j \dots s^k act^k s^{k+1}$ and $act^j, \dots, act^k \in locacts(\mathcal{R}) \cup \bigcup_{i=1}^n locacts(\mathcal{M}_i)$, $F_i \in downacts(\mathcal{R})$, and each local action of any component is eventually followed by some of its synchronising actions.

Actions are never synchronised between children, thus it can be proven by induction on the length of the run that the actions in ρ can be reordered to obtain a run $\rho' \in \text{Runs}(\mathcal{G}, s^0)$ that can be represented as $\rho' = \alpha'_1 F_1 \alpha'_2 F_2 \dots$, such that:

1. For any $i \in \mathbb{N}$ there exist $j, k \in \mathbb{N}$ such that $\alpha'_i = s'^j \text{act}'^j \dots s'^k \text{act}'^k s'^{k+1}$ and $\text{act}'^j, \dots, \text{act}'^k \in \text{locacts}(\mathcal{R}) \cup \text{locacts}(\text{snd}(F_i))$.
2. For each $i \in \mathbb{N}$ we have $\alpha_i \downarrow (\{\mathcal{R}, \text{snd}(F_i)\}) = \alpha'_i \downarrow (\{\mathcal{R}, \text{snd}(F_i)\})$.
3. For each s'^j in α'_i , if 0 is the coordinate of root and k is the coordinate of $\text{snd}(F_i)$, then $s'^j = (s_0, s_1^I, \dots, s_{k-1}^I, s_k, s_{k+1}^I, \dots)$ for some $s_0 \in \text{states}(\mathcal{R})$, $s_k \in \text{states}(\text{snd}(F_i))$.

Intuitively, ρ' is built from ρ in such a way that firstly only the root and the component that synchronises with the root over F_1 are allowed to execute their local actions while all the other components stay in their initial states; then F_1 is fired; and then this scheme is repeated for F_2, F_3 , etc. We can now project ρ' on spaces of squares of the root and components active in a given interval, to obtain $\rho'' = \alpha'_1 \downarrow (\{\mathcal{R}, \text{snd}(F_1)\}) F_1 \alpha'_2 \downarrow (\{\mathcal{R}, \text{snd}(F_2)\}) F_2 \dots$. Notice that $\rho'' \in \mathcal{SQ}^u(\mathcal{G})$. We now show that ρ'' visits all the local states that appear along ρ . To this end firstly observe that ρ and ρ' contain the same local states. Moreover, if a local state of the root is visited in ρ' then it is also visited in ρ'' , as we always project on squares of the root and some other component. If a local state of a non-root component is visited along ρ' before executing some synchronising action F_i then it will also be present before executing F_i along ρ'' . This follows from the construction of ρ'' from ρ' and our assumption (\star) of the structure of ρ . As we have shown that ρ'' visits each local state that appears along ρ , this part of the proof is concluded.

Let $\mathcal{SQ}^u(\mathcal{G}) \models EFP$ and $\rho \in \text{Runs}(\mathcal{SQ}^u(\mathcal{G}))$ visits a state labelled with p . Now, it suffices to replace in ρ each state (s_k, s_0) that belongs to the square $\mathcal{M}_k \times \mathcal{R}$ with the global state $(s_0, s_1^I, \dots, s_{k-1}^I, s_k, s_{k+1}^I, \dots)$ of \mathcal{G} . The result of this substitution is a run of \mathcal{G} that visits p . \square

The next proposition shows that the sum-of-squares does not preserve EG .

Proposition 1 (Sum-of-squares Does Not Preserve EG). *There is a live-reset two-level tree $\mathcal{ST} \mathcal{G}$ s.t. for some $p \in \mathcal{PV}$, $\mathcal{G} \models EGP$ and $\mathcal{SQ}^u(\mathcal{G}) \not\models EGP$.*

Proof. Consider the tree $\mathcal{ST} \mathcal{G}_y$ in Fig. 3. Here, we have $\mathcal{G}_y \models EGP$, but each path ρ along which p holds globally starts with \mathcal{M}_1^y executing τ followed by \mathcal{M}_2^y executing τ and, consecutively, *chooseR*. Thus, it is not possible to partition ρ into intervals where one child and the root fire local actions until they synchronise and possibly release the control to another child. Hence, $\mathcal{SQ}^u(\mathcal{G}_y) \not\models EGP$. \square

As illustrated in Example 3, the size of the state space of the sum-of-squares product of a sync-memoryless $\mathcal{ST} \mathcal{G}$ can be equal to or greater than the size of the state space of \mathcal{G} . On the other hand, the size of a representation of a state will be smaller in the sum-of-squares product, as it records only local states of at most two components of the network. However, in less degenerate cases than




Fig. 3. The sum-of-squares does not preserve EG.

our toy model we can expect significant reductions. Let us consider a tree with a root and n children, each with a state space of size m . The number of states in the asynchronous product is m^{n+1} . In contrast, the size of the sum-of-squares product of such a topology is $n \cdot m^2 + 1$. Indeed we record pairs of states from a child and its root, i.e. m^2 possible states. This is done for all n children, and there is an added fresh initial state. To summarise, the sum-of-squares product has a size in $O(n \cdot m^2)$.

3.2 Reduced Sum-of-Squares for Any Tree Height

We now adapt the sum-of-squares product of two-level live-reset tree topologies to the general case. To this end we introduce the auxiliary operation $cmpl$ that transforms $\mathcal{S}Q^u(\mathcal{G})$ into a live-reset LTS. Intuitively, $cmpl$ redirects each transition that enters the initial state of the root to the fresh initial state s_{sq}^I . No additional operations are needed for sync-deadlock topologies.

Definition 10. Let \mathcal{G} be a live-reset two-level tree \mathcal{ST} with components $Net = \{\mathcal{R}, \mathcal{M}_1, \dots, \mathcal{M}_n\}$. By $cmpl(\mathcal{S}Q^u(\mathcal{G}))$ we denote the result of replacing in $\mathcal{S}Q^u(\mathcal{G})$ every transition $((s_{\mathcal{M}_i}, s_{\mathcal{R}}), act, (s'_{\mathcal{M}_i}, s'_{\mathcal{R}}))$ with $((s_{\mathcal{M}_i}, s_{\mathcal{R}}), act, s_{sq}^I)$.

Algorithm 1 recursively performs reduction for two-level trees in a bottom-up manner. If the topology is live-reset, then each reduction is followed by applying $cmpl(\mathcal{S}Q^u(\cdot))$ to the computed sum-of-squares to ensure that the output is also live-reset. Note that \mathcal{G}_{child} denotes the subtree of \mathcal{G} rooted in $child$. The next theorem states soundness and correctness of Algorithm 1.

Theorem 2 (reduceNet(\mathcal{G}) Preserves Reachability). Let $\mathcal{G} = \langle Net, \mathcal{T} \rangle$, be a sync-memoryless tree \mathcal{ST} , $F = \bigcup_{\mathcal{M} \in Net} \text{downacts}(\mathcal{M})$, and $p \in \mathcal{PV}$. If \mathcal{G} is live-reset, then $\mathcal{G} \models EFp$ iff $\text{reduceNet}(\mathcal{G}) \models EFp$. If \mathcal{G} is sync-deadlock, then $\mathcal{G} \models_F EFp$ iff $\text{reduceNet}(\mathcal{G}) \models_F EFp$.

Proof. (Sketch) The proof follows via induction on the height of the tree \mathcal{G} . As we have Theorem 1, it suffices to prove that $cmpl(\mathcal{S}Q^u(\mathcal{G}))$ preserves reachability

for any two-level live-reset ST \mathcal{G} . This, however, can be done in a way very similar to the proof of Theorem 1 and is omitted.

Algorithm 1 $reduceNet(\mathcal{G})$

Input: sync-memoryless tree sync. topology \mathcal{G}

Output: \mathcal{LTS} \mathcal{M} that preserves reachability of each proposition p in \mathcal{G} .

```

1: if  $\mathcal{G}$  consists of a single LTS then
2:   return  $\mathcal{G}$  (*  $\mathcal{G}$  is a leaf *)
3: end if
4: let  $redChdn := \emptyset$ 
5: for  $chld \in children(root(\mathcal{G}))$  do
6:    $redChdn = redChdn \cup \{reduceNet(\mathcal{G}_{chld})\}$ 
7: end for
8: let  $\mathcal{G}' := \{root(\mathcal{G})\} \cup redChdn$ 
9: if  $\mathcal{G}$  is live-reset then
10:  return  $cmpl(\mathcal{SQ}^u(\mathcal{G}'))$ 
11: else
12:  return  $\mathcal{SQ}^u(\mathcal{G}')$ 
13: end if

```

In certain hierarchical systems such as attack-defense trees [17] we are interested only in the reachability of the root's locations. This enables for additional optimisations related to removing deadlocks and livelocks that halt the root's evolution.

Definition 11 (Root-deadlock). Let \mathcal{G} be a ST with $Net = \{\mathcal{R}, \mathcal{M}_1, \dots, \mathcal{M}_n\}$, $root(\mathcal{G}) = \mathcal{R}$ and $children(\mathcal{R}) = \{\mathcal{M}_1, \dots, \mathcal{M}_n\}$. We say that a state s of $\mathcal{M}_i \in children(\mathcal{R})$ is in a root-deadlock iff there is no run $\rho \in Runs(\mathcal{M}_i, s)$ s.t. $\rho = s^0 act^0 s^1 act^1 \dots$ with $act^i \in acts(\mathcal{R})$, where $s^0 = s$, for some $i \in \mathbb{N}$.

The set of root-deadlocked states of an \mathcal{LTS} can be computed in polynomial time using either a model checker or conventional graph algorithms. These states can be removed without affecting the reachability of a location of the root.

Example 4. Let us consider the LTS $\mathcal{SQ}^u(\mathcal{G}_x)$ in Fig. 2. If it is a child of another LTS $parent(\mathcal{SQ}^u(\mathcal{G}_x))$ in a live-reset topology and $beep \in upacts(\mathcal{SQ}^u(\mathcal{G}_x))$, then all the states coloured red are root-deadlocked. Fig. 4 displays the reduced sum-of-squares product of the topology.

Let us now evaluate the size of the state spaces for a tree of height h , considering each node has n children, each with a state space of size m . Such a tree has $\sum_{i=0}^h n^i$ nodes. Hence, the number of states in the asynchronous product is $m^{\sum_{i=0}^h n^i}$. The size of the sum-of-squares product of such a topology is $n^h \cdot m^{h+1} + \sum_{i=0}^{h-1} n^i \cdot m^i$. Indeed, for $h = 1$, this is exactly the size obtained for two-level trees in the previous section. The proof for any arbitrary




Fig. 4. The reduced sum-of-squares product of the tree \mathcal{ST} of Example 4.

height is easily done by induction. To compute the increase from height h to height $h + 1$, the same arguments as for the computation in a two-level tree hold. This leads to the following maximal number of states for height $h + 1$: $1 + n \cdot (n^h \cdot m^{h+1} + \sum_{i=0}^{h-1} n^i \cdot m^i) \cdot m = n^{h+1} \cdot m^{h+2} + \sum_{i=0}^h n^i \cdot m^i$. To summarise, the sum-of-squares product has a size in $O(n^h \cdot m^{h+1})$.

4 The General Case and Local Products

The case of live-reset LTS tree networks entails substantial state space reductions due to the fact that children reset after synchronising with their root. Before resetting, only the local states of the root and one of its children need to be recorded; it can be assumed that all the children wait in their initial states. After single-level synchronisation, only the information about the root's state is relevant, as the synchronising child resets. In the general case of tree-like networks of synchronising LTSs, both the root's state and post-synchronisation state of its children have to be preserved.

In what follows let \mathcal{G} be a two-level tree \mathcal{ST} s.t. $Net = \{\mathcal{R}, \mathcal{M}_1, \dots, \mathcal{M}_n\}$, where $root(\mathcal{G}) = \mathcal{R}$ and $children(\mathcal{R}) = \{\mathcal{M}_1, \dots, \mathcal{M}_n\}$. Let $\mathcal{R} = \langle \mathcal{S}_{\mathcal{R}}, s_{\mathcal{R}}^I, Acts_{\mathcal{R}}, \rightarrow_{\mathcal{R}}, \mathcal{L}_{\mathcal{R}} \rangle$ and $\mathcal{M}_i = \langle \mathcal{S}_i, s_i^I, Acts_i, \rightarrow_i, \mathcal{L}_i \rangle$, for $i \in \{1, \dots, n\}$. For each $i \in \{1, \dots, n\}$ let $postSync(\mathcal{M}_i)$ denote the set of all the local states of \mathcal{M}_i visited immediately after synchronising with \mathcal{R} . Formally:

$$postSync(\mathcal{M}_i) = \{s \in \mathcal{S}_i \mid \exists act \in Acts_{\mathcal{R}} \cap Acts_i \text{ s.t. } (s', r') \xrightarrow{act} (s, r) \text{ in } \mathcal{S}_i \times \mathcal{S}_{\mathcal{R}}\}.$$

The *memory unit*, defined as follows, is a collection of all the vectors of states of the components preserved after synchronisation.

Definition 12 (Memory Unit). *The post-synchronisation memory unit of \mathcal{G} is defined as: $mem(\mathcal{M}_1, \dots, \mathcal{M}_n) = \prod_{i=1}^n (\{s_i^I\} \cup postSync(\mathcal{M}_i))$.*

Intuitively, Definition 13 generalises Definition 9 by extending the states of square products with the memory unit that is updated and consulted whenever a

synchronisation takes place. This memory gadget can be implemented efficiently, but here we present it in an explicit way for the sake of readability.

Let $i \in \{1, \dots, n\}$, $s_i \in \mathcal{S}_i$, and $m \in \text{mem}(\mathcal{M}_1, \dots, \mathcal{M}_n)$. By $m[i/s_i]$ we denote the *memory update* of m defined as follows: $m[j] = m[j]$ for all $i \neq j$ and $m[i] = s_i$.

Definition 13 (General Sum-of-squares Product). Let $\mathcal{GSQ}(\mathcal{G}) = \langle \mathcal{S}_{gsq}, s_{gsq}^I, \text{Acts}_{gsq}, \rightarrow_{gsq}, \mathcal{L}_{gsq} \rangle$ be an \mathcal{LTS} s.t.:

1. $\mathcal{S}_{gsq} = (\bigcup_{i=1}^n (\mathcal{S}_i \times \mathcal{S}_{\mathcal{R}}) \times \text{mem}(\mathcal{M}_1, \dots, \mathcal{M}_n)) \cup \{s_{gsq}^I\}$.
2. $s_{gsq}^I \notin \mathcal{S}_{gsq}$ is a fresh initial state.
3. $\text{Acts}_{gsq} = \text{acts}(\mathcal{G}) \cup \{\epsilon\}$, where $\epsilon \notin \text{acts}(\mathcal{G})$ is a fresh, silent action.
4. The transition relation \rightarrow_{gsq} is defined as follows:
 - (a) For each $i \in \{1, \dots, n\}$, $s_{gsq}^I \xrightarrow{\epsilon}_{gsq} (s_i^I, s_{\mathcal{R}}^I, m_0)$, where $\forall_{i=1}^n m_0[i] = s_i^I$.
 - (b) If $s_i \xrightarrow{act}_i s'_i$ and $act \in \text{locacts}(\mathcal{M}_i)$, then $(s_i, s_{\mathcal{R}}, m) \xrightarrow{act}_{sq} (s'_i, s_{\mathcal{R}}, m)$, for each $s_{\mathcal{R}} \in \mathcal{S}_{\mathcal{R}}$ and $m \in \text{mem}(\mathcal{M}_1, \dots, \mathcal{M}_n)$; similarly, if $s_{\mathcal{R}} \xrightarrow{act}_{\mathcal{R}} s'_{\mathcal{R}}$ and $act \in \text{locacts}(\mathcal{R})$, then $(s_i, s_{\mathcal{R}}, m) \xrightarrow{act}_{sq} (s_i, s'_{\mathcal{R}}, m)$, for each $s_i \in \mathcal{S}_i$ and $m \in \text{mem}(\mathcal{M}_1, \dots, \mathcal{M}_n)$.
 - (c) If $act \in \text{upacts}(\mathcal{M}_i)$, $s_i \xrightarrow{act}_i s'_i$, and $s_{\mathcal{R}} \xrightarrow{act}_{\mathcal{R}} s'_{\mathcal{R}}$, then $(s_i, s_{\mathcal{R}}, m) \xrightarrow{act}_{sq} (s_j, s'_{\mathcal{R}}, m')$, where $m' = m[i/s'_i]$ and $s_j = m[j]$ for some $j \in \{1, \dots, n\}$.
5. $\mathcal{L}_{gsq}(s_i, s_{\mathcal{R}}, m) = \mathcal{L}_i(s_i) \cup \mathcal{L}_{\mathcal{R}}(s_{\mathcal{R}}) \cup \bigcup_{j \neq i} \mathcal{L}_j(m[j])$, for each $(s_i, s_{\mathcal{R}}, m) \in \mathcal{S}_{gsq}$.

We call $\mathcal{GSQ}(\mathcal{G})$ the *General Sum-of-squares Product* of \mathcal{G} .

Item 4a of Definition 13 expresses that the new initial state enables ϵ -transitions to the initial state of any square product $\mathcal{M}_i || \mathcal{R}$ with the memory unit set to the starting values. Similarly to the corresponding case in Definition 9, Item 4b ensures that the local actions are fully asynchronous and do not affect memory unit. The idea behind Item 4c is that in a state $(s_i, s_{\mathcal{R}}, m)$ executing a synchronised action act will require updating the outcome states, saving them into memory, and then moving all the values from memory to the current state tracker. The latter may mean switching to another component. The transition rule, as defined here, concatenates these three steps into one.

Example 5. Fig. 5 presents a three-component $\mathcal{ST} \mathcal{G}_z$. Note that the child \mathcal{M}_2^z is not live-reset, as it does not reset after synchronising with \mathcal{R}^z via *chooseF*. Fig. 6 presents the general sum-of-squares product of \mathcal{G}_z .

The following analogues of Theorem 1 and Theorem 2 show that the general sum-of-squares preserves reachability. We omit the proofs of Theorem 3 and Theorem 4, as they follow via almost exactly the same techniques.

Theorem 3 (General sum-of-squares Preserves Reachability). Let \mathcal{G} be a two-level tree \mathcal{ST} . For each $p \in \mathcal{PV}$ we have $\mathcal{G} \models \text{EF}p$ iff $\mathcal{GSQ}(\mathcal{G}) \models \text{EF}p$.




Fig. 5. A simple non live-reset tree synchronisation topology \mathcal{G}_z

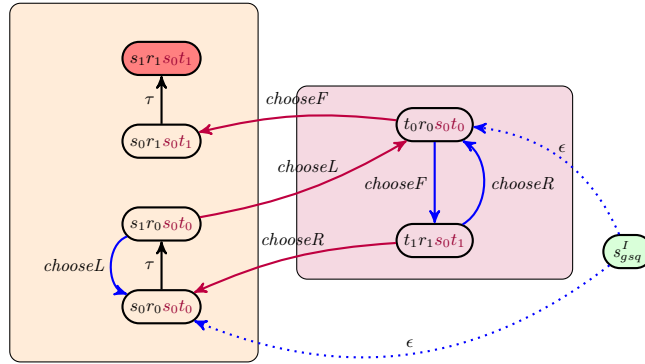


Fig. 6. The general sum-of-squares product of the non live-reset tree \mathcal{ST} in Fig. 5

A two-level general sum-of-squares can be adapted to deal with tree topologies \mathcal{G} of any height by using recursive construction. The appropriate algorithm, denoted by $reduceGenNet(\mathcal{G})$, is a slight modification of Algorithm 1 with recursive call in Line 6 replaced with $reduceGenNet(\mathcal{G}_{child})$ and the if-else conditional in Line 9 substituted with **return** $\mathcal{GSQ}(\mathcal{G}')$.

Theorem 4 (*$reduceGenNet(\mathcal{G})$ Preserves Reachability*). *Let \mathcal{G} be a tree \mathcal{ST} . For each $p \in \mathcal{PV}$, we have $\mathcal{G} \models EFP$ iff $reduceGenNet(\mathcal{G}) \models EFP$.*

It can be also easily observed that $\mathcal{GSQ}(\cdot)$ does not preserve EG , by the same argument as in Proposition 1 and Fig. 3. Moreover, as previously, root-deadlocked states (e.g. the red location in Fig. 6) can be removed if only the reachability of a location in the root is to be preserved.

The general sum-of-squares does not guarantee a reduction of the statespace size. If \mathcal{G} is a two-level tree \mathcal{ST} with $Net = \{\mathcal{R}, \mathcal{M}_1, \dots, \mathcal{M}_n\}$, then the size of the statespace of $reduceGenNet(\mathcal{G})$ can reach $\sum_{i=1}^n |\mathcal{S}_i| \cdot |\mathcal{S}_{\mathcal{R}}| \cdot \prod_{i=1}^n |\mathcal{S}_i|$, thus it

crucially depends on the size of the memory unit. On the other hand, we conjecture that the memory needed to preserve the general sum-of-squares product is often much smaller than the memory needed to hold the asynchronous product of the entire network. This conjecture is based on the observation that the states of $\mathcal{GSQ}(\mathcal{G})$ are composed of two parts: the pair of local states of two interacting modules and the memory unit which can be shared when implemented efficiently.

5 Experiments

In this section we evaluate the implementation of reductions for sync-memoryless tree topologies. The files to reproduce our tests and figures can be found at <https://depot.lipn.univ-paris13.fr/parties/publications/live-trees>.

The theory presented in Section 3 has been implemented in the open-source tool LTR [2], written in C. LTR accepts LTSs networks in a *modgraph* format [16]. The size of the fully synchronised product is computed using a Binary Decision Diagrams-based open-source Python tool *DD-Net-Checker* [1].

Model Generators: Attack-Defence Trees Attack-Defence Trees (*ADTs*) [15] are graphical models for representing possible scenarios of incoming risks and methods for their mitigation for complex systems. While descending from informal models, *ADTs* have been extended with various semantics. Here, we re-use the semantics based on translating the *ADTs* to networks of communicating LTSs [17,7]. These networks form tree-like synchronisation topologies with all the LTSs being sync-deadlock. An attack is deemed a success if a special location in the root of a network is reachable. For the purpose of this paper we implemented a simple translator from *ADTs* to *modgraph* format.

Comparing with *ADT* Reductions In [17] we presented techniques for simplifying tree-like networks using *pattern-* and *layered* reductions. The former are similar to partial order reductions and the crux of latter is in the observation that it is sufficient to consider only runs where each level of a tree fully synchronises with its children before the execution proceeds to a higher level. In [17] these techniques are implemented using time parameter injection into LTSs networks and translation to timed LTSs networks.

The comparison with the results of reductions from [17] is included for reference, as the cited work is aimed at the full timed LTS-based semantics of *ADTs* which involves numeric attributes such as the time and cost of attack. Under this semantics the networks produce considerably larger fully synchronised models. Therefore, the comparison may be slightly unfair, favouring our approach.

***ADT* Experiments** Table 1 shows the evaluation of the experiments on scalable models from [17]. Table 2 presents the results of running the experiments on security case studies, also taken from [17]. It should be noted that for the

latter we used slightly simpler models than in [17], as contrary to the former paper our tool does not handle data variables such as cost, etc. The goal in each of these scenarios is reachability of a certain location in the root node and all the networks are sync-deadlock. The timeout was set to 30 minutes (displayed as TO in the tables).

The model signature in the first column of Table 1 consists of the branching factor per an *ADT* node, the total number of nodes, the depth and the width of an *ADT* [17]. In both tables the second pair of columns collects the details of unreduced models, the third collects the results of applying the sum-of-squares (here, abbreviated to *sos*) construction to the unreduced models, the fourth of applying only pattern reduction, the fifth of pipelining the pattern and layer reductions, and the sixth of pipelining the pattern reduction and the sum-of-squares. The remaining columns collect the relative reduction/blowup rates.

The experimental data suggests that the reductions for sync-memoryless networks, proposed in this paper, are often comparable or exceeding [17]. This is especially evident when the sum-of-squares is applied to pattern-reduced models. The nature of the sum-of-squares construction can lead to statespace blowup, but this seems observable only for smaller networks. Moreover, larger networks enable more relative reductions.

To assess this further we conducted an independent series of scalable experiments on random live-reset tree networks. We generated 210 live-reset tree networks of depths 1–3 and computed their (reduced) sum-of-squares products. Fig. 7 presents the scatterplot of the results. The red line denotes no reduction.

The same phenomenon as for sync-deadlock trees can be observed: the degree of reduction increases with the size of the network.




Fig. 7. Statespace sizes of sum-of-squares for live-reset tree networks.

6 Conclusion

In this paper we outlined how to simplify large tree networks of LTSs in which the components reset or deadlock after synchronising with their parents. We

also proposed and investigated a similar construction for the general tree-like synchronisation topologies. It is shown that the constructions preserve a certain form of reachability, but do not preserve liveness. An experimental evaluation shows that the method yields extremely effective reductions for sync-memoryless networks.

We raise several questions to be explored as future work. Firstly, we only have a very rough theoretical estimate of the size of reductions for sync-memoryless networks. Stronger estimations can be obtained. Secondly, for general networks we only put a hypothesis that the state space of the general sum-of-squares may substantially grow as compared to the full asynchronous product (albeit with possibly smaller memory usage). This can be investigated experimentally. Moreover, the “vanilla” technique of general sum-of-squares is straightforward, thus surely enables many optimisations. Thirdly, the class of live-reset tree networks is probably one of many that do not need tracing what happens after synchronisation (sync-deadlock networks are slightly different, as one-shot transitions still need to be traced). Other such classes and topologies could probably be identified. Finally, the sum-of-squares for sync-deadlock LTS networks guarantees that a state is reachable before reduction iff it is *one-shot F*-reachable after (see Definition 7 and Theorem 1). We however do not know the complexity of verifying *one-shot F*-reachability; here we conjecture that it is NP-hard, which may make the reductions for this class of models less impressive than suggested by experimental results. All of these concerns can be addressed in further research.

References

1. DD-Net-Checker. <https://github.com/MichalKnapik/dd-net-checker> (2021)
2. LTR. <https://github.com/MichalKnapik/automata-net-reduction-tool> (2021)
3. van der Aalst, W.M.P., van Hee, K.M.: Workflow Management: Models, Methods, and Systems. Cooperative information systems, MIT Press (2002)
4. Aminof, B., Kupferman, O., Murano, A.: Improved model checking of hierarchical systems. *Inf. Comput.* **210**, 68–86 (2012)
5. André, É.: IMITATOR 3: Synthesis of timing parameters beyond decidability. In: CAV. LNCS, vol. 12759, pp. 552–565. Springer (2021)
6. André, É., Lime, D., Ramparison, M., Stoelinga, M.: Parametric analyses of attack-fault trees. *Fundam. Informaticae* **182**(1), 69–94 (2021)
7. Arias, J., Budde, C.E., Penczek, W., Petrucci, L., Sidoruk, T., Stoelinga, M.: Hackers vs. security: Attack-defence trees as asynchronous multi-agent systems. In: ICFEM. LNCS, vol. 12531, pp. 3–19. Springer (2020)
8. Arias, J., Celerier, J.M., Desiante-Catherine, M.: Authoring and automatic verification of interactive multimedia scores. *Journal of New Music Research* (2016)
9. Arias, J., Petrucci, L., Masko, L., Penczek, W., Sidoruk, T.: Minimal schedule with minimal number of agents in attack-defence trees. In: ICECCS. pp. 1–10. IEEE (2022)
10. Baier, C., Katoen, J.: Principles of model checking. MIT Press (2008)
11. Behrmann, G., David, A., Larsen, K.G., Håkansson, J., Pettersson, P., Yi, W., Hendriks, M.: UPPAAL 4.0. In: QEST. pp. 125–126. IEEE Computer Society (2006)

12. Belardinelli, F., Lomuscio, A., Murano, A., Rubin, S.: Verification of broadcasting multi-agent systems against an epistemic strategy logic. In: IJCAI. pp. 91–97 (2017)
13. Holzmann, G.J.: The SPIN Model Checker - primer and reference manual. Addison-Wesley (2004)
14. Knapik, M., Meski, A., Penczek, W.: Action synthesis for branching time logic: Theory and applications. *ACM Trans. Embed. Comput. Syst.* **14**(4), 64:1–64:23 (2015)
15. Kordy, B., Mauw, S., Radomirovic, S., Schweitzer, P.: Attack-defense trees. *J. Log. Comput.* **24**(1), 55–87 (2014)
16. Lakos, C., Petrucci, L.: Modular analysis of systems composed of semiautonomous subsystems. In: ACSD. pp. 185–196. IEEE Computer Society (2004)
17. Petrucci, L., Knapik, M., Penczek, W., Sidoruk, T.: Squeezing state spaces of (attack-defence) trees. In: ICECCS. pp. 71–80. IEEE (2019)