

Join-Lock-Sensitive Forward Reachability Analysis for Concurrent Programs with Dynamic Process Creation ^{*}

Thomas Martin Gawlitza¹, Peter Lammich², Markus Müller-Olm², Helmut Seidl³,
and Alexander Wenner²

¹ CNRS/VERIMAG, France, Thomas.Gawlitza@imag.fr

² Institut für Informatik, Westfälische Wilhelms-Universität Münster, Germany
{[peter.lammich](mailto:peter.lammich@wwu.de), [markus.mueller-olm](mailto:markus.mueller-olm@wwu.de), [alexander.wenner](mailto:alexander.wenner@wwu.de)}@wwu.de

³ Technische Universität München, Germany, seidl@in.tum.de

Abstract. Dynamic Pushdown Networks (DPNs) are a model for parallel programs with (recursive) procedures and dynamic process creation. *Constraints* on the sequences of spawned processes allow to extend the basic model with joining of created processes [2]. Orthogonally DPNs can be extended with nested locking [8]. Reachability of a regular set R of configurations in presence of stable constraints as well as reachability without constraints but with nested locking are based on computing the set of predecessors $\text{pre}^*(R)$. In the present paper, we present a *forward*-propagating algorithm for deciding reachability for DPNs. We represent sets of executions by sets of *execution trees* and show that the set of all execution trees resulting in configurations from R which either allow a lock-sensitive execution or a join-sensitive execution, is *regular*. Here, we rely on basic results about *macro tree transducers*. As a second contribution, we show that reachability is decidable also for DPNs with both nested locking and joins.

1 Introduction

Bouajjani et al. [2] introduced (Constrained) Dynamic Pushdown Networks ((C)DPNs) for modeling parallel programs with (potentially recursive) procedures and dynamic process creation. They propose an algorithm inspired by the pre^* -algorithm for Pushdown Systems that, under the mild assumption that all constraints are *stable*, computes the set of all predecessors of a given regular set of configurations. Stable constraints are, for instance, capable of modeling *fork-join* parallelism as considered by Esparza and Podelski [4], Müller-Olm and Seidl [9], Seidl and Steffen [10].

Only weak forms of synchronization can be expressed in CDPNs. In order to deal with *nested* locking, Lammich et al. [8] introduced another mechanism for restricting executions. For unconstrained DPNs, they introduce *action trees*⁴ to represent sets of related executions. Generalizing the acquisition history techniques of Kahlon et al. [5, 6], they construct a finite tree automaton to check whether an action tree represents a *lock-sensitive* execution. This allows them to perform a lock-sensitive backward reachability analysis for unconstrained DPNs with nested locking.

^{*} This work was partially funded by the DFG project OpIAT (MU 1508/1-1 and SE 551/13-1) and the ANR project ASOPT.

⁴ Action trees are called execution trees in Lammich et al. [8].

The techniques of Lammich et al. [8] cannot directly be adapted to perform a lock-sensitive *forward* reachability analysis, since the set of all possible executions cannot always be represented by a regular set of action trees. In this paper, we therefore propose an augmented version of action trees, which we call *execution trees* throughout the present paper. Our construction should be seen in analogy to contextfree grammars where the linear derivation process can be represented in treelike form by representing subderivations referring to distinct sections of sentential forms in distinct subtrees. We show that the set of executions can be represented by a *regular* set of execution trees. This enables us to use tree automata techniques for forward reachability analysis. Moreover, by applying known results for *macro tree transducers*, we show that every tree regular property of action trees can be translated into a tree regular property of execution trees. This allows us to re-use the results of Lammich et al. [8] to perform a *forward* reachability analysis for DPNs with joins as well as a lock-sensitive forward reachability analysis for DPNs with nested locking. The algorithms we obtain in this way are both natural and not less efficient than the corresponding algorithms for backward reachability analysis of Bouajjani et al. [2] and Lammich et al. [8], respectively. These algorithms for *forward* reachability are our first contribution.

In programming languages which support multithreaded programming such as JAVA and C with POSIX threads, joins and locks may occur simultaneously in programs. Therefore, as our second contribution, we provide an algorithm for performing a *forward* reachability analysis that is simultaneously join- and lock-sensitive. Following the approach of Lammich et al. [8], we construct a finite tree automaton which accepts all action trees for which a schedule exists that is simultaneously lock-sensitive and join-sensitive. This tree automaton can be used with both, the forward analysis techniques of this paper, and the backward analysis techniques of Lammich et al. [8]. Note that such an automaton cannot be constructed by simply intersecting a tree automaton for lock-sensitive action trees with a tree automaton for join-sensitive action trees, since action trees may represent multiple executions. Hence, an action tree may represent a lock-sensitive execution and a different join-sensitive execution, but no execution that is simultaneously join- and lock-sensitive.

Before we proceed with the technical constructions, we consider an introductory example written in a real-world programming language:

Example 1. We consider the following C program that uses POSIX threads:

```
pthread_mutex_t m = PTHREAD_MUTEX_INITIALIZER;
void* f(void* x) {
    pthread_mutex_lock(&m); printf("A"); printf("B\n");
    pthread_mutex_unlock(&m); }
int main() { pthread_t t[10]; int i;
    for (i = 0; i < 10; i++) pthread_create(&t[i], NULL, f, NULL);
    for (i = 0; i < 10; i++) pthread_join(t[i], NULL);
    printf("End!\n"); return 0; }
```

Here, a main process creates 10 child-processes, waits until all these processes terminate, prints *End!*, and terminates. Each child-process prints *AB*. The lock *m* guarantees that the critical section of *f* is not used by two different processes at the same time. The main process does not need to acquire the lock *m*, since it only may access the

printer when no child-process is alive. A lock-sensitive reachability analysis which does not take joining into account, may flag a spurious warning that exclusive access to the printer might be violated. Likewise, a join-sensitive reachability analysis which ignores the locking cannot exclude potential violations. In order to exclude such false alarms, a reachability analysis is required which is simultaneously join- and lock-sensitive. \square

2 Dynamic Pushdown Networks

Let us first introduce some notation: The Boolean lattice $\{\perp, \top\}$ is denoted by \mathbb{B} . The natural numbers are denoted by \mathbb{N} . We assume $0 \in \mathbb{N}$ and set $\mathbb{N}_+ := \mathbb{N} \setminus \{0\}$. The domain of a partial function $f : A \rightsquigarrow B$ is denoted by $\text{dom}(f)$. \oplus denotes the update operator for partial functions, i.e., for $f, g : A \rightsquigarrow B$ and $x \in A$, $(f \oplus g)(x)$ equals $g(x)$ if $x \in \text{dom}(g)$ and $f(x)$ otherwise. We write $\{x_1 \mapsto y_1, \dots, x_n \mapsto y_n\}$ for the partial function that maps x_i to y_i for all $i \in \{1, \dots, n\}$ and is undefined otherwise.

Let Act be a non-empty set of *actions*. The specific action $\text{sp} \notin \text{Act}$ represents the spawning of a new process. We set $\overline{\text{Act}} := \text{Act} \cup \{\text{sp}\}$. A *Dynamic Pushdown Network* (DPN) M is a tuple $(\text{Act}, P, \Gamma, \Delta)$ where P is a finite set of *control states*, Γ is a finite set of *stack symbols*, and Δ is a finite set of *transition rules* of the following forms:

$$\begin{array}{ll} \text{(Base)} \quad p\gamma \xrightarrow{a} p'\gamma' & \text{(Push)} \quad p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \\ \text{(Pop)} \quad p\gamma \xrightarrow{a} p' & \text{(Spawn)} \quad p\gamma \xrightarrow{\text{sp}} p'\gamma' \triangleright p_s\gamma_s \end{array}$$

Here, $p, p', p_s \in P$, $a \in \text{Act}$, and $\gamma, \gamma', \gamma_1, \gamma_2 \in \Gamma$. Intuitively, DPNs extend pushdown systems by (Spawn)-rules, that create a new process as a side effect.

A *configuration* of a DPN $M = (\text{Act}, P, \Gamma, \Delta)$ is an *unranked tree* where each node in the tree is labeled with an element from $P\Gamma^*$. Intuitively, each node of the tree denotes a process and is labeled with the current configuration of the process. The children of a process are the processes that were spawned by the parent process. Formally, we represent nodes in trees by sequences of natural numbers: A set $T \subseteq \mathbb{N}_+^*$ is called a *tree domain* iff (1) T is finite, and (2) prefix-closed, i.e., $l \in T$ implies $l' \in T$ for all prefixes l' of l , and (3) if $li \in T$ with $i > 1$ then also $l(i-1) \in T$. The set of all tree domains is denoted by $\mathcal{T}^{\mathbb{N}_+}$. Accordingly, a configuration c is a partial mapping $c : \mathbb{N}_+^* \rightsquigarrow P\Gamma^*$ where the domain $\text{dom}(c)$ is a tree domain. Let \mathcal{C} denote the set of all configurations of M . Within a configuration c , processes are organized hierarchically. Each element $l \in \text{dom}(c)$ denotes a process. The set $\{li \in \text{dom}(c) \mid i \in \mathbb{N}_+\}$ contains all child-processes of l within the configuration c . For $c(l) = pw$ with $p \in P, w \in \Gamma^*$, p is the control-state and w the stack of the process l within the configuration c . We call a configuration c_0 an *initial configuration* iff $c_0 = \{\epsilon \mapsto p\gamma\}$ for some $p \in P$ and some $\gamma \in \Gamma$, i.e. it contains a single process with a single symbol on its stack.

A single execution step of a DPN on a configuration applies a local rewrite step to the control-state and stack of one of the processes of the configuration. For all transition rules $r \in \Delta$ and all processes $l \in \mathbb{N}_+^*$, we define the partial function $\mapsto_r^l : \mathcal{C} \rightsquigarrow \mathcal{C}$ that applies r to the process l as follows:

$$\begin{array}{ll} \mapsto_{p\gamma \xrightarrow{a} p'\gamma'}^l(c) := c \oplus \{l \mapsto p'w'w\} & \text{if } c(l) = p\gamma w \\ \mapsto_{p\gamma \xrightarrow{\text{sp}} p'\gamma' \triangleright p_s\gamma_s}^l(c) := c \oplus \{l \mapsto p'\gamma'w, \text{nc}(l, \text{dom}(c)) \mapsto p_s\gamma_s\} & \text{if } c(l) = p\gamma w \end{array}$$

Here, $\text{nc}(l, T) := l \cdot \min\{i \in \mathbb{N}_+ \mid l \cdot i \notin T\}$ is the sequence of natural numbers denoting the next child of process l which can be added to the tree domain T preserving the tree domain property.

Constraint Structures. In this paragraph we define *constraint structures* that allow us to restrict sequences of execution steps of DPNs. The idea is to disable undesired execution steps depending on a global state, that is computed in parallel to the actual configuration of the DPN. Constraint structures will later be used to define join-lock sensitive executions. Formally a constraint structure is a tuple $\mathbb{C} = (\mathbb{S}, \mathbb{T}, s_0)$ where \mathbb{S} is a non-empty set of *global states*, $s_0 \in \mathbb{S}$ is an *initial state*, and $\mathbb{T} : \mathbb{N}_+^* \times \overline{\text{Act}} \rightarrow \mathbb{S} \rightsquigarrow \mathbb{S}$ is a *global state transformer*. For $l \in \mathbb{N}_+^*$, $a \in \overline{\text{Act}}$, and $s \in \mathbb{S}$, $\mathbb{T}(l, a)(s)$ denotes the global state resulting from applying action a to process l in global state s , or is undefined. If $\mathbb{T}(l, a)(s)$ is undefined, we consider the constraint to be violated. In this case, the corresponding action cannot be performed, i.e. the corresponding transition rule is not applicable. In order to capture this, we define the partial function $\overset{\mathbb{C}}{\mapsto}_r^l : \mathbb{S} \times \mathcal{C} \rightsquigarrow \mathbb{S} \times \mathcal{C}$ by

$$\overset{\mathbb{C}}{\mapsto}_r^l((s, c)) := (\mathbb{T}(l, \overline{\text{Act}}(r))(s), \mapsto_r^l(c)) \quad \text{if } \mathbb{T}(l, \overline{\text{Act}}(r))(s) \text{ and } \mapsto_r^l(c) \text{ are defined}$$

Here, $\overline{\text{Act}}(r)$ denotes the action of the transition rule r . The transition rule r is applied to the process l as before and the global state is transformed according to the transformer of the constraint structure. Thus, the function is only defined if the transformation of the global state is defined. Whenever the next global state is undefined, the constraint structure \mathbb{C} prevents the application of r to l .

The set of configurations reachable from an initial configuration c_0 under a constraint structure \mathbb{C} is defined as $\text{post}_{\mathbb{C}, M}^*(c_0) = \{c_0\} \cup \{c \mid \exists k \in \mathbb{N}_+, l_1, \dots, l_k \in \mathbb{N}_+^*, r_1, \dots, r_k \in \Delta, s \in \mathbb{S}. (s, c) = \overset{\mathbb{C}}{\mapsto}_{r_k}^{l_k}(\dots \overset{\mathbb{C}}{\mapsto}_{r_1}^{l_1}(s_0, c_0) \dots)\}$. One goal of this paper is to present an algorithm to check whether a set R of configurations is reachable from an initial configuration c_0 , given one of the constraint structures $\mathbb{C} \in \{\mathbb{C}^u, \mathbb{C}^j, \mathbb{C}^l, \mathbb{C}^{jl}\}$ defined in the following. Formally, this is equivalent to checking whether

$$\text{post}_{\mathbb{C}, M}^*(c_0) \cap R \neq \emptyset.$$

Join/Lock-Constraints. In the present paper we consider joins together with nested locking. For that, let $L = \{1, \dots, n\}$ be a fixed, finite set of locks. We assume a fixed set of actions $\text{Act} = \{\text{jo}, \$, \epsilon\} \cup \{\text{acq}(i), \text{rel}(i) \mid i \in L\}$. The action jo blocks the current process until all child-processes of the current process are stopped. The action $\$$ stops a process. The action ϵ represents actions unrelated to joining and locking. $\text{acq}(i)$ acquires and $\text{rel}(i)$ releases the lock i .

Unconstrained. For dealing with *unconstrained* DPNs, we define the constraint structure $\mathbb{C}^u := (\{\top\}, \mathbb{T}^u, \top)$ by $\mathbb{T}^u(l, a)(\top) := \top$ for all $l \in \mathbb{N}_+^*$, $a \in \overline{\text{Act}}$. This constraint structure never disables any execution steps.

Joins. For dealing with joins we define a constraint structure that keeps track of whether existing processes have stopped or not. Processes can only make steps if they have not

stopped and jo-steps can only be performed if all child processes have stopped. Hence we define $\mathbb{C}^j := (\mathbb{S}^j, \mathbb{T}^j, s_0^j)$ by $\mathbb{S}^j := \mathbb{N}_+^* \rightsquigarrow \mathbb{B}$, $s_0^j(\epsilon) = \perp$, $s_0^j(l)$ is undefined for all $l \in \mathbb{N}_+^* \setminus \{\epsilon\}$, and

$$\mathbb{T}^j(l, a)(s) := \begin{cases} s & \text{if } s(l) = \perp, a = \text{jo}, \\ & \forall i \in \mathbb{N}_+. li \in \text{dom}(s) \Rightarrow s(li) = \top \\ s \oplus \{l \mapsto \top\} & \text{if } s(l) = \perp, a = \$ \\ s \oplus \{\text{nc}(l, \text{dom}(s)) \mapsto \perp\} & \text{if } s(l) = \perp, a = \text{sp} \\ s & \text{if } s(l) = \perp, a \notin \{\text{jo}, \$, \text{sp}\} \end{cases}$$

for all $s \in \mathbb{S}^j$, all $l \in \text{dom}(s)$, and all $a \in \overline{\text{Act}}$. For all $s \in \mathbb{S}^j$ and all $l \in \mathbb{N}_+^*$, $s(l) = \top$ iff the process l exists and it is stopped.

Locks. For dealing with locks we define a constraint structure that keeps track of the set of locks held by each existing process. acq-steps can only be performed if the corresponding lock is held by no process and rel-steps require the corresponding lock to be held by the process. Formally, we define $\mathbb{C}^l := (\mathbb{S}^l, \mathbb{T}^l, s_0^l)$ by $\mathbb{S}^l := \mathbb{N}_+^* \rightsquigarrow 2^L$, $s_0^l(\epsilon) = \emptyset$, $s_0^l(l)$ is undefined for all $l \in \mathbb{N}_+^* \setminus \{\epsilon\}$, and

$$\mathbb{T}^l(l, a)(s) := \begin{cases} s \oplus \{l \mapsto (s(l) \cup \{i\})\} & \text{if } a = \text{acq}(i), i \notin \bigcup_{l' \in \text{dom}(s)} s(l') \\ s \oplus \{l \mapsto (s(l) \setminus \{i\})\} & \text{if } a = \text{rel}(i), i \in s(l) \\ s \oplus \{\text{nc}(l, \text{dom}(s)) \mapsto \emptyset\} & \text{if } a = \text{sp} \\ s & \text{if } a \notin \{\text{acq}(i), \text{rel}(i), \text{sp} \mid i \in L\} \end{cases}$$

for all $s \in \mathbb{S}^l$, all $l \in \text{dom}(s)$, and all $a \in \overline{\text{Act}}$. For all $s \in \mathbb{S}^l$ and all $l \in \mathbb{N}_+^*$, $s(l)$ is the set of locks that are acquired by the process l , provided that the process l exists. Here, we only consider non-reentrant locks, i.e., a process that already owns a lock i is not allowed to acquire i again. However, under certain conditions, DPNs with reentrant locks can be simulated by DPNs with non-reentrant locks at a cost that is exponential in the maximal nesting depth of locks (see Kidd et al. [7], Lammich et al. [8]).

Products of Constraint Structures. For all constraint structures $\mathbb{C}_1 = (\mathbb{S}_1, \mathbb{T}_1, s_0^{(1)})$ and $\mathbb{C}_2 = (\mathbb{S}_2, \mathbb{T}_2, s_0^{(2)})$, we define the *product* $\mathbb{C}_1 \times \mathbb{C}_2$ of \mathbb{C}_1 and \mathbb{C}_2 by

$$\mathbb{C}_1 \times \mathbb{C}_2 := (\mathbb{S}_1 \times \mathbb{S}_2, \mathbb{T}_1 \times \mathbb{T}_2, (s_0^{(1)}, s_0^{(2)})),$$

where $(\mathbb{T}_1 \times \mathbb{T}_2)(l, a)((s_1, s_2)) := (\mathbb{T}_1(l, a)(s_1), \mathbb{T}_2(l, a)(s_2))$ for all $l \in \mathbb{N}_+^*$, all $a \in \overline{\text{Act}}$, and all $(s_1, s_2) \in \mathbb{S}_1 \times \mathbb{S}_2$ iff $\mathbb{T}_1(l, a)(s_1)$ and $\mathbb{T}_2(l, a)(s_2)$ are defined.

Joins and Locks. For dealing with joins and locks simultaneously, we define the constraint structure $\mathbb{C}^{jl} := \mathbb{C}^j \times \mathbb{C}^l$. A DPN together with the constraint structure \mathbb{C}^{jl} is called a *join-lock-DPN* for short.

Example 2. The behavior of the C program of Example 1 can be safely over-approximated by the join-lock-DPN with the following transition rules:

$$p_m m \xrightarrow{\text{sp}} p_m m \triangleright p_f f \quad p_m m \xrightarrow{\text{jo}} \$_m m \quad p_f f \xrightarrow{\text{acq}(1)} q_f f \quad q_f f \xrightarrow{\text{rel}(1)} r_f f \quad r_f f \xrightarrow{\$} \$_f f$$

The main process $p_m m$ may create arbitrary many child-processes $p_f f$ and may terminate (go into the state $\$m$) as soon as all child-processes are stopped. Each child-process acquires the lock 1, releases the lock 1 and then stops. Because of the semantics of locks, two different child-processes may never be in state q_f simultaneously. Moreover, because of the semantics of joins, all child-processes must be in state $\$f$, before the main-process can reach state $\$m$. \square

3 Forward Reachability

In this section we present our approach for checking forward reachability using only operations on regular sets of trees.

Configurations. In order to explicitly describe sets of configurations of a DPN $M = (\text{Act}, P, \Gamma, \Delta)$ using regular tree languages, i.e. tree automata, we represent the unranked trees c from the set \mathcal{C} by *ranked* trees. For that, we use *lists* constructed with `cons` and `nil`. As a shorthand we later write lists `cons($a_1, \dots, \text{cons}(a_k, \text{nil}) \dots$)` as $[a_1; \dots; a_k]$. The set \mathcal{C} of ranked trees which we will use is specified by the tree grammar:

$$LG ::= \text{nil} \mid \text{cons}(\Gamma, LG) \quad LC ::= \text{nil} \mid \text{cons}(C, LC) \quad C ::= P(LG, LC)$$

Hence a configuration is a ranked tree where the root node is labeled with the control state of the initial process. The root node has two branches. One is the list of stack symbols of the initial process. The other is the list of processes spawned by the initial process, where the last process spawned is the first in the list. For each spawned process the list contains a configuration with the process as initial process. The function $f : \mathcal{C} \rightarrow \mathcal{C}$ captures the relation between the unranked and ranked model of configurations. It is defined by $f(c) := f(c, \epsilon)$ for all $c \in \mathcal{C}$, where, for all $c \in \mathcal{C}$ and all $l \in \text{dom}(c)$, $f(c, l) := p([a_1; \dots; a_m], [c'_k; \dots; c'_1])$ if $c(l) = pa_1 \dots a_m$, $p \in P$, $a_1, \dots, a_m \in \Gamma$, $\{i \in \mathbb{N}_+ \mid l \cdot i \in \text{dom}(c)\} = \{1, \dots, k\}$, and $c'_1 = f(c, l1), \dots, c'_k = f(c, lk)$. Recall that $l1$ is the first and lk the last process spawned by process l within configuration c . Since f is a bijection, all functions and concepts defined in Section 2 can be transferred to the new representation and in the following we redefine the set \mathcal{C} to be the set \mathcal{C} of ranked trees. Hence a configuration $c \in \mathcal{C}$ is now a ranked tree and an initial configuration is represented by a term of the form $p([\gamma], \text{nil})$ with $p \in P$ and $\gamma \in \Gamma$.

Known Results. For all pushdown systems M , $\text{pre}^*(R)$ and $\text{post}^*(R)$ are regular, whenever $R \subseteq \mathcal{C}$ is a regular set of configurations. Bouajjani et al. [1] provide a construction which, starting from a finite automaton for R , successively adds transitions until an automaton for $\text{pre}^*(R)$ is obtained. Bouajjani et al. [2] applied a generalization of this construction to implement pre^* -operators for DPNs with *stable* constraints. The post^* -operator for DPNs, however, does *not* preserve regularity as shown in Bouajjani et al. [2] for a word model of configurations. This also holds for the tree model of this paper:

Example 3 (A Non-Regular post^ -Image).* We consider the DPN $M = (\text{Act}, P, \Gamma, \Delta)$ that is defined by $P = \{p, q, p', q', r, r', \$\}$, $\Gamma = \{a\}$, and

$$\Delta = \{pa \xrightarrow{\text{sp}} qa \triangleright ra, ra \xrightarrow{\$} \$, qa \xrightarrow{\epsilon} paa, pa \xrightarrow{\text{jo}} p', p'a \xrightarrow{\text{sp}} q'a \triangleright r'a, q'a \xrightarrow{\epsilon} p'\}.$$

Here, $\text{post}_{\mathcal{C}^u, M}^*(p([a], \text{nil}))$ is not regular, since the process can spawn at most k processes of type $r'a$ after spawning k processes of type ra . \square

Execution Trees. We avoid the problem of non-regularity by introducing *execution trees*. Execution trees can be regarded as configuration trees with a history of rule applications that lead from an initial configuration to the reached configuration. Inner nodes of execution trees are annotated with rules of the DPN and leaf nodes contain information about the current control state and, if existent, the current top-most stack symbol.

Starting from the root node, the leftmost branch of an execution tree contains the rules applied to the initial process in the order of application. Whenever a new process was spawned, the corresponding *sp*-rule has a separate branch containing the rules applied to the new process. Additionally, if a stack symbol is pushed that is later popped, the push-node has two branches: The left one describes the execution up to popping the stack symbol and the right one describes the rest of the execution.

To this end, we consider the set Δ of all transition rules of the DPN M as a ranked alphabet where *Push*-rules have rank 1 or 2, *Spawn*-rules have rank 2, *Base*-rules have rank 1, and *Pop*-rules have ranks 0. Additionally, for $p \in P$ and $\gamma \in \Gamma$, we introduce symbols $\langle p, \gamma \rangle$ of rank 0. The set X of execution trees is then described by the following regular tree grammar:

$$\begin{aligned} \text{XT} &::= \langle \text{Base} \rangle(\text{XT}) \mid \langle \text{Push} \rangle(\text{XT}, \text{XT}) \mid \langle \text{Pop} \rangle \mid \langle \text{Spawn} \rangle(\text{XT}, X) \\ \text{XN} &::= \langle \text{Base} \rangle(\text{XN}) \mid \langle \text{Push} \rangle(\text{XN}) \mid \langle \text{Push} \rangle(\text{XT}, \text{XN}) \mid \langle \text{Spawn} \rangle(\text{XN}, X) \mid \langle P \times \Gamma \rangle \\ X &::= \text{XT} \mid \text{XN} \end{aligned}$$

Here, *Base* is the set of *Base*-rules, *Push* is the set of *Push*-rules, *Pop* is the set of *Pop*-rules, and *Spawn* is the set of *Spawn*-rules. The elements of the set X are the execution trees, where *XT* (*terminated execution trees*) contains the execution trees that end with popping the stack-symbol of the current stacklevel, and *XN* (*not yet terminated execution trees*) contains the execution trees that do not pop the stack-symbol of the current stacklevel.

An execution tree t reaches a configuration $c(t)$. We can extract this configuration by $c(t) := c(t, \text{nil}, \text{nil})$ with:

$$\begin{aligned} c(\langle p, \gamma \rangle, w, \bar{c}) &:= p(\text{cons}(\gamma, w), \bar{c}) \\ c(\langle p\gamma \xrightarrow{a} p'\gamma' \rangle(t), w, \bar{c}) &:= c(t, w, \bar{c}) \\ c(\langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle(t), w, \bar{c}) &:= c(t, \text{cons}(\gamma_2, w), \bar{c}) \\ c(\langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle(t_1, t_2), w, \bar{c}) &:= c(t_2, w, \text{ch}(t_1, \bar{c})) \\ c(\langle p\gamma \xrightarrow{a} p' \rangle, w, \bar{c}) &:= p'(w, \bar{c}) \\ c(\langle p\gamma \xrightarrow{\text{sp}} p\gamma \triangleright p_s\gamma_s \rangle(t, t_s), w, \bar{c}) &:= c(t, w, \text{cons}(c(t_s), \bar{c})) \\ \text{ch}(\langle p, \gamma \rangle, \bar{c}) &:= \bar{c} \\ \text{ch}(\langle p\gamma \xrightarrow{a} p'\gamma' \rangle(t), \bar{c}) &:= \text{ch}(t, \bar{c}) \\ \text{ch}(\langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle(t), \bar{c}) &:= \text{ch}(t, \bar{c}) \\ \text{ch}(\langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle(t_1, t_2), \bar{c}) &:= \text{ch}(t_2, \text{ch}(t_1, \bar{c})) \\ \text{ch}(\langle p\gamma \xrightarrow{a} p' \rangle, \bar{c}) &:= \bar{c} \end{aligned}$$

$$\text{ch}(\langle p\gamma \xrightarrow{\text{sp}} p\gamma \triangleright p_s\gamma_s \rangle(t, t_s), \bar{c}) := \text{ch}(t, \text{cons}(c(t_s), \bar{c}))$$

Here ch is an auxiliary function that extracts the list of configurations of all spawned processes in an execution tree. The function is used to collect the configurations of processes which were spawned in left branches of push-nodes. Note that initial configurations $p([\gamma], \text{nil})$ are represented by execution trees consisting of a single node $\langle p, \gamma \rangle$.

The mappings c and ch for extracting the configuration can be interpreted as total deterministic *macro tree transducers* (see, e.g., Engelfriet and Vogler [3] for definitions and fundamental results), which use each of their input and output parameters exactly once. The pre-image of a regular set of trees w.r.t. the translation of any macro tree transducer is again regular [3]. By carefully inspecting the macro tree transducers realizing the mappings c and ch , we obtain:

Lemma 1. *For every regular set C of configurations, $c^{-1}(C)$ is a regular set of execution trees. If C is given by a nondeterministic finite tree automaton with n states and m transitions, a finite tree automaton for $c^{-1}(C)$ can be constructed in time $\mathcal{O}(m^2 + |\Delta| \cdot n^4 + |\Delta| \cdot m \cdot n^2 + |P| \cdot |\Gamma| \cdot n)$.*

Due to lack of space, the proof of this and some other lemmas is deferred to the appendix.

We now define a semantics $\mathbb{C} \mapsto_r^l$ on execution trees that is compatible with the semantics $\mathbb{C} \mapsto_r^l$ on configurations. For that, for all $r \in \Delta$, we define the partial function $\mapsto_r : X \rightsquigarrow X$ that applies r to the root process of an execution tree as follows:

$$\begin{aligned} \mapsto_{p\gamma \xrightarrow{a} p'\gamma'}(\langle p, \gamma \rangle) &:= \langle p\gamma \xrightarrow{a} p'\gamma' \rangle(\langle p', \gamma' \rangle) \\ \mapsto_{p\gamma \xrightarrow{a} p'\gamma_1\gamma_2}(\langle p, \gamma \rangle) &:= \langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle(\langle p', \gamma_1 \rangle) \\ \mapsto_{p\gamma \xrightarrow{a} p'}(\langle p, \gamma \rangle) &:= \langle p\gamma \xrightarrow{a} p' \rangle \\ \mapsto_{p\gamma \xrightarrow{\text{sp}} p'\gamma' \triangleright p_s\gamma_s}(\langle p, \gamma \rangle) &:= \langle p\gamma \xrightarrow{\text{sp}} p'\gamma' \triangleright p_s\gamma_s \rangle(\langle p', \gamma' \rangle, \langle p_s, \gamma_s \rangle) \\ \mapsto_r(\langle p\gamma \xrightarrow{a} p'\gamma' \rangle(t)) &:= \langle p\gamma \xrightarrow{a} p'\gamma' \rangle(\mapsto_r(t)) \\ \mapsto_r(\langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle(t)) &:= \begin{cases} \langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle(t') & \text{if } t' \in \text{XN} \\ \langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle(t', \langle s(c(t')), \gamma_2 \rangle) & \text{if } t' \in \text{XT} \\ \text{where } t' = \mapsto_r(t) \end{cases} \\ \mapsto_r(\langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle(t, t')) &:= \langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle(t, \mapsto_r(t')) \\ \mapsto_r(\langle p\gamma \xrightarrow{\text{sp}} p'\gamma' \triangleright p_s\gamma_s \rangle(t, t_s)) &:= \langle p\gamma \xrightarrow{\text{sp}} p'\gamma' \triangleright p_s\gamma_s \rangle(\mapsto_r(t), t_s) \end{aligned}$$

Here, $s(p(w, \bar{c})) := p$ for all $p(w, \bar{c}) \in C$ extracts the control-state of the root process of an execution tree. It is used to extract the control state at the end of a left branch of a push-node. This control-state is the initial control-state for the continued execution on the right branch.

We define $\mapsto_r^l(t)$ to be the application of \mapsto_r to the subtree of the execution tree which has the process identified by $l \in \mathbb{N}_+^*$ as root process. One can find the specific subtree by counting the number of sp -rules along the leftmost branch starting from the root node. If the correct sp -rule was found, we descend into the right branch

and continue. If we arrive at the end of a left branch of a push-node, we continue at the right branch of that node. Analogously to Section 2 we define $\mathbb{C} \mapsto_r^l$ and the set $\text{tpost}_{\mathbb{C},M}^*(\langle p, \gamma \rangle) = \{\langle p, \gamma \rangle\} \cup \{t \mid \exists k \in \mathbb{N}_+, l_1, \dots, l_k \in \mathbb{N}_+^*, r_1, \dots, r_k \in \Delta, s \in \mathbb{S}. (s, t) = \mathbb{C} \mapsto_{r_k}^{l_k} (\dots \mathbb{C} \mapsto_{r_1}^{l_1} (s_0, \langle p, \gamma \rangle) \dots)\}$ of successors of an execution tree representing an initial configuration for execution trees. An execution tree is called *consistent* iff there exists some $p \in P$ and some $\gamma \in \Gamma$ such that $t \in \text{tpost}_{\mathbb{C}^u, M}^*(\langle p, \gamma \rangle)$, i.e. $\text{tpost}_{\mathbb{C}^u, M}^*(\langle p, \gamma \rangle)$ is the set of consistent execution trees starting in $p\gamma$. We then have:

Lemma 2. $c(\mathbb{C} \mapsto_r^l(t)) = \mathbb{C} \mapsto_r^l(c(t))$ holds for all execution trees $t \in X$, all transition rules $r \in \Delta$, all positions $l \in \mathbb{N}_+^*$ and constraint structures \mathbb{C} .

Thus, to check reachability we can check $c(\text{tpost}_{\mathbb{C}, M}^*(\langle p, \gamma \rangle)) \cap R \neq \emptyset$ or equivalently

$$\text{tpost}_{\mathbb{C}, M}^*(\langle p, \gamma \rangle) \cap c^{-1}(R) \neq \emptyset.$$

We have already mentionend, that $c^{-1}(R)$ is regular if R is regular. In order to obtain an algorithm which allows to check reachability solely based on operations on regular sets of trees, it remains to show that $\text{tpost}_{\mathbb{C}, M}^*(\langle p, \gamma \rangle)$ is regular as well. We first show this for unconstrained DPNs, i.e. for $\mathbb{C} = \mathbb{C}^u$.

Regular characterization of $\text{tpost}_{\mathbb{C}, M}^(\langle p, \gamma \rangle)$.* Let $p \in P$ and $\gamma \in \Gamma$. Our goal is to compute the set $\text{tpost}_{\mathbb{C}^u, M}^*(\langle p, \gamma \rangle)$ of all execution trees that can be reached starting from the execution tree $\langle p, \gamma \rangle$ without any restrictions due to the constraint structure. For that we construct a finite tree automaton which recognizes $\text{tpost}_{\mathbb{C}^u, M}^*(\langle p, \gamma \rangle)$.

For all control states $p, p' \in P$ and all stack symbols $\gamma \in \Sigma$, we inductively define the set $\mathcal{T}^{p, \gamma, p'}$ (resp. $\mathcal{N}^{p, \gamma}$) of *terminated* (resp. *not yet terminated*) execution trees with initial control state p and top-most stack symbol γ and terminating with control state p' (resp. not terminating) as follows:

- Open Leaf:** For all $p \in P$, and all $\gamma \in \Gamma$: $\langle p, \gamma \rangle \in \mathcal{N}^{p, \gamma}$
- Pop:** For all $r = p\gamma \xrightarrow{a} p' \in \Delta$: $\langle r \rangle \in \mathcal{T}^{p, \gamma, p'}$
- Base:** For all $r = p\gamma \xrightarrow{a} p'\gamma' \in \Delta$, and all $p'' \in P$:
 - If $t \in \mathcal{T}^{p', \gamma', p''}$, then $\langle r \rangle(t) \in \mathcal{T}^{p, \gamma, p''}$.
 - If $t \in \mathcal{N}^{p', \gamma'}$, then $\langle r \rangle(t) \in \mathcal{N}^{p, \gamma}$.
- Push:** For all $r = p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \in \Delta$, and all $p'', p''' \in P$:
 - If $t_1 \in \mathcal{T}^{p', \gamma_1, p''}$ and $t_2 \in \mathcal{T}^{p'', \gamma_2, p'''}$, then $\langle r \rangle(t_1, t_2) \in \mathcal{T}^{p, \gamma, p'''}$.
 - If $t_1 \in \mathcal{T}^{p', \gamma_1, p''}$ and $t_2 \in \mathcal{N}^{p'', \gamma_2}$, then $\langle r \rangle(t_1, t_2) \in \mathcal{N}^{p, \gamma}$.
 - If $t_1 \in \mathcal{N}^{p', \gamma_1}$, then $\langle r \rangle(t_1) \in \mathcal{N}^{p, \gamma}$.
- Spawn:** For all $r = p\gamma \xrightarrow{\text{sp}} p'\gamma' \triangleright p_s\gamma_s \in \Delta$, and all $p'', p'_s \in P$:
 - If $t \in \mathcal{T}^{p', \gamma', p''}$ and $t_s \in \mathcal{T}^{p_s, \gamma_s, p'_s} \cup \mathcal{N}^{p_s, \gamma_s}$, then $\langle r \rangle(t, t_s) \in \mathcal{T}^{p, \gamma, p''}$.
 - If $t \in \mathcal{N}^{p', \gamma'}$ and $t_s \in \mathcal{T}^{p_s, \gamma_s, p'_s} \cup \mathcal{N}^{p_s, \gamma_s}$, then $\langle r \rangle(t, t_s) \in \mathcal{N}^{p, \gamma}$.

Thus, we get:

- Theorem 1.** 1. $\text{tpost}_{\mathbb{C}^u, M}^*(\langle p, \gamma \rangle) = \bigcup_{p' \in P} \mathcal{T}^{p, \gamma, p'} \cup \mathcal{N}^{p, \gamma}$.
2. The sets $\mathcal{T}^{p, \gamma, p'}$, $\mathcal{N}^{p, \gamma}$, and finally $\text{tpost}_{\mathbb{C}^u, M}^*(\langle p, \gamma \rangle)$ can be recognized by a finite tree automaton that can be constructed in time $\mathcal{O}(|P|^2 \cdot (|\Gamma| + |\Delta|))$.

Proof (Sketch). Statement 1 can be shown by straight-forward induction. We focus on Statement 2: The states of the finite tree automaton correspond to the sets $\mathcal{T}^{p,\gamma,p'}$ and $\mathcal{N}^{p,\gamma}$. Hence, we have $|P|^2 \cdot |I|$ states. Moreover, there are $\mathcal{O}(|P| \cdot |I| + |P|^2 \cdot |\Delta|)$ transition rules. This gives us the desired complexity estimation. \square

Hence, we can effectively check reachability for unconstrained DPNs using only operations on regular tree languages, i.e. finite tree automata. To close the gap to DPNs with a constraint structure, we observe, that the order in which rules are applied to processes running in parallel in an unconstrained DPN is irrelevant for the finally reached execution tree. However, this is not the case for DPNs with a constraint structure, since the order of rule applications determines how the global state of the constraint structure is transformed. Consequently, there may be orders which are not feasible, since a necessary transformation is undefined. To capture this, we define a *schedule* $(l_1, a_1), \dots, (l_k, a_k)$ of an execution tree t as a sequence of process-labels and actions, such that there exists $p \in P, \gamma \in \Gamma$ and rules $r_1, \dots, r_k \in \Delta$, with $\bar{\text{Act}}(r_i) = a_i$, such that $t = \Rightarrow_{r_k}^{l_k} (\dots \Rightarrow_{r_1}^{l_1} (\langle p, \gamma \rangle) \dots)$. Each consistent execution tree has at least one schedule. A schedule is call *executable* under a constraint structure iff additionally $\mathbb{T}(l_k, a_k)(\dots \mathbb{T}(l_1, a_1)(s_0) \dots)$ is defined. Then, $\text{tpost}_{\mathbb{C}, M}^*(\langle p, \gamma \rangle)$ is exactly the set of consistent execution trees starting in $p\gamma$ that have an executable schedule under \mathbb{C} . As a consequence we can write $\text{tpost}_{\mathbb{C}, M}^*(\langle p, \gamma \rangle) = \text{tpost}_{\mathbb{C}^u, M}^*(\langle p, \gamma \rangle) \cap H_{\mathbb{C}}$, where $H_{\mathbb{C}}$ is a set of execution trees that contains all consistent execution trees that have an executable schedule and no consistent execution tree that has no executable schedule. Then our reachability problem boils down to checking

$$\text{tpost}_{\mathbb{C}^u, M}^*(\langle p, \gamma \rangle) \cap H_{\mathbb{C}} \cap c^{-1}(R) \neq \emptyset.$$

In the remainder of this paper, we construct regular sets $H_{\mathbb{C}}$ for all constraint structures $\mathbb{C} \in \{\mathbb{C}^j, \mathbb{C}^l, \mathbb{C}^{j1}\}$. Thus, checking reachability reduces to checking emptiness of intersections of regular sets of trees.

4 Action Trees

At the end of the last section we have observed that, in order to check reachability under a constraint structure, we need to specify a regular set of execution trees that contains all consistent execution trees with an executable schedule and no consistent execution tree without an executable schedule. We have also observed that executability of a schedule depends only on the actions of the rules and identities of the processes taking part in the execution. Consequently, we do not need the full power of the execution tree model to capture executability. Indeed, it is simpler to characterize executability if the left branches of push-nodes are inlined as in the trees used by Lammich et al. [8]. Therefore, in this section, we introduce *action trees* as an abstraction of execution trees that does this inlining and omits information about control states and stack contents.

Formally, the set of action trees \mathcal{T}^{Act} is defined by the tree grammar $\mathcal{T}^{\text{Act}} ::= \epsilon \mid \langle \text{Act} \rangle(\mathcal{T}^{\text{Act}}) \mid \langle \text{sp} \rangle(\mathcal{T}^{\text{Act}}, \mathcal{T}^{\text{Act}})$. The root node and each right child of an sp-node within an action tree represents a process that performs the actions of the left-most path starting from this node and ending in a leaf.

Schedules. For all $l \in \mathbb{N}_+^*$, $a \in \overline{\text{Act}}$, we define the partial function $\mapsto_a^l : \mathcal{T}^{\text{Act}} \rightsquigarrow \mathcal{T}^{\text{Act}}$, that constructs action trees by

$$\begin{array}{ll}
\mapsto_a^l(\epsilon) := \langle a \rangle(\epsilon) & \text{if } l = \epsilon, a \in \text{Act} \\
\mapsto_{\text{sp}}^l(\epsilon) := \langle \text{sp} \rangle(\epsilon, \epsilon) & \text{if } l = \epsilon \\
\mapsto_a^l(\langle a' \rangle(t)) := \langle a' \rangle(\mapsto_a^l(t)) & \text{if } a' \in \text{Act} \\
\mapsto_a^l(\langle \text{sp} \rangle(t, t_s)) := \langle \text{sp} \rangle(\mapsto_a^l(t), t_s) & \text{if } l = \epsilon \\
\mapsto_a^l(\langle \text{sp} \rangle(t, t_s)) := \langle \text{sp} \rangle(t, \mapsto_a^{l'}(t_s)) & \text{if } l = 1 \cdot l' \\
\mapsto_a^l(\langle \text{sp} \rangle(t, t_s)) := \langle \text{sp} \rangle(\mapsto_a^{(i-1) \cdot l'}(t), t_s) & \text{if } l = i \cdot l', i > 1
\end{array}$$

$\mapsto_a^l(t)$ appends the action a to the process l of the action tree. Identification of processes is similar as in the case of execution trees.

A schedule of an action tree t is a finite sequence $(l_1, a_1), \dots, (l_k, a_k)$, where $l_i \in \mathbb{N}_+^*$ and $a_i \in \overline{\text{Act}}$ for all $i \in \{1, \dots, k\}$, with $t = \mapsto_{a_k}^{l_k}(\dots \mapsto_{a_1}^{l_1}(\epsilon) \dots)$. Action trees always have at least one schedule in contrast to execution trees, where only consistent execution trees have a schedule. Recall that a schedule is called executable under a constraint structure \mathbb{C} iff $\mathbb{T}(l_k, a_k)(\dots \mathbb{T}(l_1, a_1)(s_0) \dots)$ is defined.

Execution Trees and Action Trees. Each execution tree t can be abstracted to an action tree $\mathbf{a}(t)$ by cutting away unnecessary information. We define $\mathbf{a}(t) := \mathbf{a}(t, \epsilon)$, where:

$$\begin{array}{l}
\mathbf{a}(\langle p, \gamma \rangle, x) := x \\
\mathbf{a}(\langle p\gamma \xrightarrow{a} p'\gamma' \rangle(t), x) := \langle a \rangle(\mathbf{a}(t, x)) \\
\mathbf{a}(\langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle(t), x) := \langle a \rangle(\mathbf{a}(t, x)) \\
\mathbf{a}(\langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle(t_1, t_2), x) := \langle a \rangle(\mathbf{a}(t_1, \mathbf{a}(t_2, x))) \\
\mathbf{a}(\langle p\gamma \xrightarrow{a} p' \rangle, x) := \langle a \rangle(x) \\
\mathbf{a}(\langle p\gamma \xrightarrow{\text{sp}} p\gamma \triangleright p_s\gamma_s \rangle(t, t_s), x) := \langle \text{sp} \rangle(\mathbf{a}(t, x), \mathbf{a}(t_s))
\end{array}$$

The mapping \mathbf{a} for extracting action trees can again be interpreted as a total deterministic macro tree transducer. We get:

Lemma 3. *For every regular set T of action trees, $\mathbf{a}^{-1}(T)$ is a regular set of execution trees. If T is given by a nondeterministic finite tree automaton with n states and m transitions, a finite tree automaton for $\mathbf{a}^{-1}(T)$ can be constructed in time $\mathcal{O}(|P| \cdot |G| \cdot n + |\Delta| \cdot m \cdot n^2)$.*

Moreover, scheduling of action trees is compatible with that of execution trees:

Lemma 4. *A consistent execution tree t has an executable schedule under a constraint structure \mathbb{C} iff $\mathbf{a}(t)$ has an executable schedule.*

Consequently, for the set $T_{\mathbb{C}}$ of all action trees that have an executable schedule under the constraint structure \mathbb{C} , each consistent execution tree $t \in \mathbf{a}^{-1}(T_{\mathbb{C}})$ has an executable schedule. Vice versa, each consistent execution tree $t \notin \mathbf{a}^{-1}(T_{\mathbb{C}})$ has no executable schedule. Thus, we define $H_{\mathbb{C}} := \mathbf{a}^{-1}(T_{\mathbb{C}})$. Since \mathbf{a}^{-1} preserves regularity, $H_{\mathbb{C}}$ is regular if $T_{\mathbb{C}}$ is regular.




Fig. 1. Action Tree with join- and lock-sensitive but no join-lock-sensitive schedule.

Lammich et al. [8] showed how to construct a tree automaton check^l that accepts the set $T_{\mathbb{C}^l}$ of all action trees that can be scheduled lock-sensitively, i.e. have an executable schedule under \mathbb{C}^l . A tree automaton check^j that accepts the set $T_{\mathbb{C}^j}$ of all action trees that can be scheduled join-sensitively, i.e. have an executable schedule under \mathbb{C}^j , can be – rather straightforwardly – constructed in polynomial time as well. The tree automaton check^j checks that all child-processes that are spawned before a jo-action are stopped using the action \$ and the action \$ is only performed at the end of branches.

The language $T_{\mathbb{C}^{jl}}$, however, need not to be the intersection of the languages $T_{\mathbb{C}^j}$ and $T_{\mathbb{C}^l}$. The reason is that the set of join-sensitive schedules and the set of lock-sensitive schedules of t might be disjoint. An example for an action tree with a join-sensitive and a lock-sensitive, but no join-lock-sensitive schedule can be seen in Figure 4.

5 Join-Lock-Sensitive Schedules of Action Trees

In this section we construct a deterministic bottom-up tree automaton check^{jl} which recognizes the set of all action trees that can be scheduled join-lock-sensitively. We restrict ourselves to action trees where each single process in isolation uses locks in a well-nested fashion, releases all locks before stopping, and does not try to perform actions after the stop action. We call such action trees join-lock-well-formed, a notion defined more formally in the next paragraph.

Well-Formedness. A sequence of actions $a_1 \cdots a_k$ is called *well-nested* iff its lock operations are a prefix of a sequence from the context-free language $U ::= \epsilon \mid U \cdot U \mid \text{acq}(1) \cdot U \cdot \text{rel}(1) \mid \cdots \mid \text{acq}(n) \cdot U \cdot \text{rel}(n)$. It is called *non-reentrant* iff each acquired lock is released before it is acquired again, i.e., for each $a_i = \text{acq}(j)$, $a_{i'} = \text{acq}(j)$, $i < i'$, there is an i'' with $i < i'' < i'$ and $a_{i''} = \text{rel}(j)$, and, symmetrically, each released lock is re-acquired before it is released again, i.e., for each $a_i = \text{rel}(j)$, $a_{i'} = \text{rel}(j)$, $i < i'$ there is an i'' with $i < i'' < i'$ and $a_{i''} = \text{acq}(j)$. A sequence of actions is called *well-formed* iff it is both well-nested and non-reentrant. It is called *join-lock-well-formed* iff

(1) it is well-formed, (2) there is no action after the action $\$,$ and (3) if the sequence ends with the action $\$,$ then its sequence of lock operations is from $U,$ i.e., all acquired locks are released before terminating with the action $\$.$

An action tree is called *join-lock-well-formed* iff the actions of each process form a join-lock-well-formed sequence. A join-lock-DPN M is called *join-lock-well-formed* iff all action trees that correspond to consistent execution trees of M are join-lock-well-formed. Note that the set of join-lock-well-formed action trees is regular. Thus, we can use the techniques presented above to check whether a DPN is join-lock-well-formed.

Acquisition-Structures. For defining the tree automaton check^{jl} , we use a modified version of *acquisition structures* as introduced by Lammich et al. [8].

An acquisition of a lock i without a matching release is called a *final acquisition* of $i.$ A matching release means, that the same process releases the lock, i.e., that there is a release of i on the leftmost path starting at the acquisition node. Symmetrically, a release of i without a matching acquisition is called an *initial release.* Acquisitions and releases that are not final acquisitions or initial releases are called *usages.*

Let t be a join-lock-well-formed action tree. We first check the following properties, that are obviously necessary for the existence of a join-lock-sensitive schedule: (1) Every lock is finally acquired at most once. (2) For every process, the last action of all child-processes that are spawned before some join is $\$.$ We write $\text{ok}(t)$ iff t fulfills the requirements (1) and (2). The evaluation of the predicate ok can easily be implemented by a deterministic bottom-up tree automaton where the complexity of the construction is exponential in the number of locks.

To each subtree of a join-lock-well-formed action tree, we assign an acquisition structure that can be used for constructing a join-lock-sensitive schedule. In this paper we define an acquisition structure either to be \perp or a tuple $(A, R, R_{\text{jo}}, U, \rightarrow_A, J)$ where:

- A is the set of acquired locks (final acquisitions and usages).
- R is the set of initially released locks. Note that, due to well-nestedness, only the root process of the subtree may contain initial releases.
- R_{jo} is the set of locks that are initially released after the next join. If the root process contains no join operation, we have $R_{\text{jo}} = \emptyset.$ Moreover, we have $R_{\text{jo}} \subseteq R.$
- U is the set of locks whose usages are required for termination of the root process, including locks whose usages are required for termination of spawned processes the root process joins with.
- $\rightarrow_A \subseteq L^2$ is the *acquisition graph.* We have $i \rightarrow_A j$ iff there is an acquisition of j after a final acquisition of $i.$
- $J \in \mathbb{B}$ is \top iff the root process performs at least one jo.

The acquisition structure $\text{as}(t)$ of an action tree t is defined by the following rules, if the subtrees of t have non-bottom acquisition structures and the side conditions of the

rules are fulfilled. In any other case, we set $\text{as}(t) := \perp$.

$$\begin{aligned}
\text{as}(\epsilon) &:= \text{as}(\langle \$ \rangle(\epsilon)) := (\emptyset, \emptyset, \emptyset, \emptyset, \emptyset, \perp) \\
\text{as}(\langle \epsilon \rangle(t)) &:= \text{as}(t) \\
\text{as}(\langle \text{acq}(i) \rangle(t)) &:= (A \cup \{i\}, R \setminus \{i\}, R_{\text{jo}} \setminus \{i\}, U \cup \{i\}, \rightarrow_A, J), \quad \text{if } i \in R \\
\text{as}(\langle \text{acq}(i) \rangle(t)) &:= (A \cup \{i\}, R, R_{\text{jo}}, U, \rightarrow_{A \cup \{i\}} \times A, J), \quad \text{if } i \notin R \\
\text{as}(\langle \text{rel}(i) \rangle(t)) &:= (A, R \cup \{i\}, R_{\text{jo}}, U, \rightarrow_A, J) \\
\text{as}(\langle \text{jo} \rangle(t)) &:= (A, R, R, U, \rightarrow_A, \top) \\
\text{as}(\langle \text{sp} \rangle(t, t')) &:= (A \cup A', R, R_{\text{jo}}, U, \rightarrow_{A \cup A'}, \perp), \quad \text{if } J = \perp \\
\text{as}(\langle \text{sp} \rangle(t, t')) &:= (A \cup A', R, R_{\text{jo}}, U \cup U', \rightarrow_{A \cup A'}, \top), \\
&\quad \text{if } J = \top \text{ and } R_{\text{jo}} \cap U' = \emptyset
\end{aligned}$$

where $\text{as}(t) = (A, R, R_{\text{jo}}, U, \rightarrow_A, J)$ and $\text{as}(t') = (A', R', R'_{\text{jo}}, U', \rightarrow_{A'}, J')$. We prove that a join-lock-well-formed action tree t can be scheduled join-lock-sensitively iff $\text{ok}(t)$ and $\text{as}(t) = (A, R, R_{\text{jo}}, U, \rightarrow_A, J)$ with acyclic \rightarrow_A . Since the number of acquisition structures is finite, a deterministic bottom-up tree automaton check^{jl} can be constructed that accepts the action trees that satisfy the above property:

Theorem 2. *There is a deterministic bottom-up tree automaton check^{jl} which accepts a join-lock-well-formed action tree t iff t can be scheduled join-lock-sensitively; check^{jl} can be constructed in time at most exponential in the number of locks.*

Proof (Sketch). Due to lack of space only a sketch is presented here. The complete proof can be found in the appendix.

The complexity estimate for check^{jl} follows from the fact that the number of acquisition structures is exponentially bounded in the number of locks.

Assume that check^{jl} does not accept t . Hence, not $\text{ok}(t)$ or $\text{as}(t) = \perp$ or $\text{as}(t) = (A, R, R_{\text{jo}}, U, \rightarrow_A, J)$ with cyclic \rightarrow_A . First of all, it is easy to see that there is no join-lock sensitive schedule if $\text{ok}(t)$ does not hold. If $\text{as}(t) = \perp$, this results from a failed $R_{\text{jo}} \cap U' = \emptyset$ check in the second sp -rule of m . This means that there is a process that spawns another process and then joins, while it holds a lock that is required by the spawned process for termination. Hence, no lock-sensitive schedule is possible. The edges in \rightarrow_A capture ordering constraints on final acquisitions of locks: $i \rightarrow_A j$ means, that there is a final acquisition of i that must be scheduled before the final acquisition of j , if any. If \rightarrow_A is cyclic, these constraints obviously cannot be fulfilled.

If check^{jl} accepts t , then $\text{ok}(t)$ holds and $\text{as}(t) = (A, R, R_{\text{jo}}, U, \rightarrow_A, J)$ with an acyclic \rightarrow_A . We first choose a total ordering $<$ on the locks with $\rightarrow_A \subseteq <$. Then, we construct inductively a schedule $s(t')$ for each subtree t' of t with the following properties: (1) The actions in $s(t')$ form a suffix of a join-lock-well-formed action sequence. (2) $s(t')$ is join-sensitive, i.e. steps of spawned processes are scheduled before the respective join operations. (3) Acquisitions in $s(t')$ respect the ordering $<$ in the sense that only locks $i > j$ are acquired after an unmatched acquisition of j . (4) Only required locks are used before termination, i.e. before termination of the root process, all acquired locks are released and are from U' (where $\text{as}(t') = (A', R', R'_{\text{jo}}, U', \rightarrow_{A'}, J')$).

The construction of this schedule is straightforward except for subtrees of the form $t' = \langle \text{sp} \rangle(t_1, t_2)$. In this case, $s(t')$ is constructed as a specific shuffling of $s(t_1)$ and

$s(t_2)$. Then, by (1) and (2), the schedule $s(t)$ constructed for the entire tree t is join-lock sensitive. \square

We are now prepared to put all parts together:

Theorem 3. *Let $M = (\text{Act}, P, \Gamma, \Delta)$ be a join-lock-well-formed join-lock-DPN, $p \in P$, $\gamma \in \Gamma$, and R a regular set of configurations. It can be decided whether a configuration from R can be reached by M , starting from the configuration $p([\gamma], \text{nil})$ through an execution that respects joins and locks simultaneously, i.e., whether or not*

$$\text{post}_{\mathbb{C}^{\text{jl}}, M}^*(p([\gamma], \text{nil})) \cap R \neq \emptyset$$

holds, by checking the following intersection of regular sets:

$$\text{tpost}_{\mathbb{C}^{\text{u}}, M}^*(\langle p, \gamma \rangle) \cap a^{-1}(T) \cap c^{-1}(R) \neq \emptyset$$

where T is the language of check^{jl} . The complexity is polynomial in the size of M and in that of a tree automaton recognizing R and exponential only in the number of locks.

6 Conclusion

We have developed a *forward* propagating algorithm for analyzing reachability in Dynamic Pushdown Networks. The key idea was to represent all possible executions by means of a regular set of execution trees. We then showed how to restrict the set of all possible executions to executions that respect joins and locks simultaneously. By that, we obtained an algorithm for analyzing reachability w.r.t. join-lock-sensitive executions. Our algorithms are polynomial in the size of the DPN as well as in the size of the finite tree automaton describing the set of configurations to be reached and exponential only in the number of locks used by the DPN.

It would be interesting to explore to what extent nested locking can be combined with concepts that are more general than joins, e.g. *stable* constraints as considered by Bouajjani et al. [2]. It might also be challenging to see which other properties can be decided for the different classes of DPNs beyond reachability. Finally, it is left to future work to practically evaluate the provided techniques on real world examples.

Bibliography

- [1] A. Bouajjani, J. Esparza, and O. Maler. Reachability analysis of pushdown automata: Application to model-checking. In A. W. Mazurkiewicz and J. Winkowski, editors, *CONCUR*, volume 1243 of *Lecture Notes in Computer Science*, pages 135–150. Springer, 1997. ISBN 3-540-63141-0.
- [2] A. Bouajjani, M. Müller-Olm, and T. Touili. Regular symbolic analysis of dynamic networks of pushdown systems. In *Concurrency Theory. 16th Int. Conf. (CONCUR)*, pages 473–487. LNCS 3653, Springer, 2005.
- [3] J. Engelfriet and H. Vogler. Macro tree transducers. *J. Comput. Syst. Sci.*, 31(1): 71–146, 1985.

- [4] J. Esparza and A. Podelski. Efficient algorithms for pre* and post* on interprocedural parallel flow graphs. In *POPL*, pages 1–11, 2000.
- [5] V. Kahlon and A. Gupta. An automata-theoretic approach for model checking threads for LTL properties. In *Proc. of LICS 2006*, pages 101–110. IEEE Computer Society, 2006.
- [6] V. Kahlon, F. Ivancic, and A. Gupta. Reasoning about threads communicating via locks. In *Proc. of CAV 2005*, volume 3576 of *LNCS*. Springer, 2005.
- [7] N. Kidd, A. Lal, and T. W. Reps. Language strength reduction. In M. Alpuente and G. Vidal, editors, *SAS*, volume 5079 of *Lecture Notes in Computer Science*, pages 283–298. Springer, 2008. ISBN 978-3-540-69163-1.
- [8] P. Lammich, M. Müller-Olm, and A. Wenner. Predecessor sets of dynamic push-down networks with tree-regular constraints. In A. Bouajjani and O. Maler, editors, *CAV*, volume 5643 of *Lecture Notes in Computer Science*, pages 525–539. Springer, 2009. ISBN 978-3-642-02657-7.
- [9] M. Müller-Olm and H. Seidl. On optimal slicing of parallel programs. In *STOC*, pages 647–656, 2001.
- [10] H. Seidl and B. Steffen. Constraint-based inter-procedural analysis of parallel programs. *Nord. J. Comput.*, 7(4):375–, 2000.

A Omitted Proofs

Proof (of Lemma 1). For a proof, we provide an explicit construction for the tree automaton in question. Let S be a regular set of (representations of) configurations. Assume that we are given a nondeterministic finite automaton A that recognizes S , i.e., $\mathcal{L}(A) = S$. W.l.o.g. the set of states of A is $\overline{Q} = Q \dot{\cup} W \dot{\cup} C$, the transition relation of A is $\hookrightarrow = \hookrightarrow_W \dot{\cup} \hookrightarrow_C \dot{\cup} \hookrightarrow_Q$, the set of accepting states F of A is a subset of Q , and the following holds:

- W is a finite set of states for stack contents together with a designated state s_W for the empty stack;
- $\hookrightarrow_W \subseteq \text{cons}(\Gamma, W) \times W$;
- C is a finite set of states for sequences of processes together with a designated state s_C for the empty sequence of processes;
- $\hookrightarrow_C \subseteq \text{cons}(Q, C) \times C$.
- Q is a finite set of states for processes;
- $\hookrightarrow_Q \subseteq P(W, C) \times Q$

The finite automaton A' that accepts $c^{-1}(S)$ then is given by

- the set of states $Q \times W \times C \cup C \times C$;
- the set of accepting states $F \times \{s_W\} \times \{s_C\}$;
- the transition relation \hookrightarrow' of A' which is defined as follows:
 1. $\langle p, \gamma \rangle \hookrightarrow' (q, w, c)$ for all $p(w_1, c) \hookrightarrow_Q q$, $\text{cons}(\gamma, w) \hookrightarrow_W w_1$.
($\mathcal{O}(|\hookrightarrow|^2)$ transitions)
 2. $\langle p\gamma \xrightarrow{a} p'\gamma' \rangle ((q, w, c)) \hookrightarrow' (q, w, c)$ for all $p\gamma \xrightarrow{a} p'\gamma' \in \Delta$, $q \in Q$, $w \in W$, $c \in C$.
($\mathcal{O}(|\Delta| \cdot |\overline{Q}|^3)$ transitions)
 3. $\langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle ((q, w_1, c)) \hookrightarrow' (q, w, c)$ for all $p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \in \Delta$, $\text{cons}(\gamma_2, w_1) \hookrightarrow_W w$, $q \in Q$, $c \in C$.
($\mathcal{O}(|\Delta| \cdot |\hookrightarrow| \cdot |\overline{Q}|^2)$ transitions)
 4. $\langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle ((c_2, c), (q, w, c_2)) \hookrightarrow' (q, w, c)$ for all $p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \in \Delta$, $q \in Q$, $w \in W$, $c, c_2 \in C$.
($\mathcal{O}(|\Delta| \cdot |\overline{Q}|^4)$ transitions)
 5. $\langle p\gamma \xrightarrow{a} p' \rangle \hookrightarrow' (q, w, c)$ for all $p\gamma \xrightarrow{a} p' \in \Delta$, $p'(w, c) \hookrightarrow_Q q$.
($\mathcal{O}(|\Delta| \cdot |\hookrightarrow|)$ transitions)
 6. $\langle p\gamma \xrightarrow{\text{sp}} p\gamma \triangleright p_s\gamma_s \rangle ((q, w, c_1), (q_1, s_W, s_C)) \hookrightarrow' (q, w, c)$ for all $p\gamma \xrightarrow{\text{sp}} p\gamma \triangleright p_s\gamma_s \in \Delta$, $\text{cons}(q_1, c_1) \hookrightarrow_C c$, $q \in Q$, $w \in W$.
($\mathcal{O}(|\Delta| \cdot |\hookrightarrow| \cdot |\overline{Q}|^2)$ transitions)
 7. $\langle p, \gamma \rangle \hookrightarrow' (c, c)$ for all $p \in P$, $\gamma \in \Gamma$, $c \in C$
($\mathcal{O}(|P| \cdot |\Gamma| \cdot |\overline{Q}|)$ transitions)
 8. $\langle p\gamma \xrightarrow{a} p'\gamma' \rangle ((c, c_1)) \hookrightarrow' (c, c_1)$ for all $p\gamma \xrightarrow{a} p'\gamma' \in \Delta$, $c, c_1 \in C$.
($\mathcal{O}(|\Delta| \cdot |\overline{Q}|^2)$ transitions)

9. $\langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle((c, c_1)) \hookrightarrow' (c, c_1)$ for all $p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \in \Delta, c, c_1 \in C$.
($\mathcal{O}(|\Delta| \cdot |\overline{Q}|^2)$ transitions)
10. $\langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle((c, c_2), (c_2, c_1)) \hookrightarrow' (c, c_1)$ for all $p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \in \Delta, c, c_1, c_2 \in C$.
($\mathcal{O}(|\Delta| \cdot |\overline{Q}|^3)$ transitions)
11. $\langle p\gamma \xrightarrow{a} p' \rangle \hookrightarrow' (c, c)$ for all $p\gamma \xrightarrow{a} p' \in \Delta, c \in C$.
($\mathcal{O}(|\Delta| \cdot |\overline{Q}|)$ transitions)
12. $\langle p\gamma \xrightarrow{\text{sp}} \text{sp}\gamma \triangleright p_s\gamma_s \rangle((c, c_2), (q, s_W, s_C)) \hookrightarrow' (c, c_1)$ for all $p\gamma \xrightarrow{\text{sp}} p\gamma \triangleright p_s\gamma_s \in \Delta, \text{cons}(q, c_2) \hookrightarrow_C c_1, c \in C$.
($\mathcal{O}(|\Delta| \cdot |\hookrightarrow| \cdot |\overline{Q}|)$ transitions)

The tree automaton A' can be constructed in time

$$\mathcal{O}(|\hookrightarrow|^2 + |\Delta| \cdot |\overline{Q}|^4 + |\Delta| \cdot |\hookrightarrow| \cdot |\overline{Q}|^2 + |P| \cdot |T| \cdot |\overline{Q}|).$$

□

Proof (of Lemma 2). The statement can be shown by a straight-forward induction.

Proof (of Theorem 1). The statements can be shown by a straight-forward induction.

Proof (of Lemma 3). For a proof, we provide an explicit construction for the tree automaton in question. Let T be a regular set of action trees. Assume that we are given a nondeterministic finite automaton A that recognizes T , i.e., $\mathcal{L}(A) = T$. Let Q be the set of states of A , and \hookrightarrow be the transition relation of A , and F be the accepting states of A . Assume that the state q_ϵ accepts exactly the set $\{\epsilon\}$ of action trees.

The set of states of The finite automaton A' that accepts $a^{-1}(T)$ is given by

- the set of states Q^2 ;
- the set of accepting states $F \times \{q_\epsilon\}$; and
- the transition relation \hookrightarrow' which is defined as follows:
 1. $\langle p, \gamma \rangle \hookrightarrow' (q, q)$ for all $p \in P, \gamma \in T, q \in Q$.
($\mathcal{O}(|P| \cdot |T| \cdot |Q|)$ transitions)
 2. $\langle p\gamma \xrightarrow{a} p'\gamma' \rangle((q_1, q')) \hookrightarrow' (q, q')$ for all $p\gamma \xrightarrow{a} p'\gamma' \in \Delta, \langle a \rangle(q_1) \hookrightarrow q, q' \in Q$.
($\mathcal{O}(|\Delta| \cdot |\hookrightarrow| \cdot |Q|)$ transitions)
 3. $\langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle((q_1, q')) \hookrightarrow' (q, q')$ for all $p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \in \Delta, \langle a \rangle(q_1) \hookrightarrow q, q' \in Q$.
($\mathcal{O}(|\Delta| \cdot |\hookrightarrow| \cdot |Q|)$ transitions)
 4. $\langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle((q_1, q_2), (q_2, q'_2)) \hookrightarrow' (q, q'_2)$ for all $p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \in \Delta, \langle a \rangle(q_1) \hookrightarrow q, q_2, q'_2 \in Q$.
($\mathcal{O}(|\Delta| \cdot |\hookrightarrow| \cdot |Q|^2)$ transitions)
 5. $\langle p\gamma \xrightarrow{a} p' \rangle \hookrightarrow' (q, q_1)$ for all $p\gamma \xrightarrow{a} p' \in \Delta, \langle a \rangle(q_1) \hookrightarrow q$.
($\mathcal{O}(|\Delta| \cdot |\hookrightarrow|)$ transitions)

6. $\langle p\gamma \xrightarrow{\text{sp}} p'\gamma' \triangleright p_s\gamma_s \rangle((q_1, q'_1), (q_2, q_\epsilon)) \hookrightarrow' (q, q'_1)$ for all $p\gamma \xrightarrow{\text{sp}} p'\gamma' \triangleright p_s\gamma_s \in \Delta$, $\langle \text{sp} \rangle(q_1, q_2) \hookrightarrow q, q'_1 \in Q$.
 ($\mathcal{O}(|\Delta| \cdot |\hookrightarrow| \cdot |Q|)$ transitions)

The tree automaton A' can be constructed in time

$$\mathcal{O}(|P| \cdot |\Gamma| \cdot |Q| + |\Delta| \cdot |\hookrightarrow| \cdot |Q|^2).$$

□

Proof (of Lemma 4). Let the set \overline{X} be defined by the following regular tree grammar:

$$\begin{aligned} \overline{X}' &::= \epsilon \mid \langle \text{Base} \cup \text{Push} \cup \text{Pop} \rangle(\overline{X}') \mid \langle \text{Spawn} \rangle(\overline{X}', \overline{X}) \\ \overline{X} &::= \langle P \times \Gamma \rangle(\overline{X}') \end{aligned}$$

The trees in \overline{X} are inlined execution trees, where the initial configuration is annotated at the root, all additional branches of *Push*-rules are inlined and leaves are replaced by ϵ .

$t \in \overline{X}$ represents a set of executions that in particular lead to the same configuration $c'(t)$ that is defined by $c'(\langle p, \gamma \rangle(t)) := c'(t, p([\gamma], \text{nil}))$, where

$$\begin{aligned} c'(\epsilon, c) &:= c \\ c'(\langle p\gamma \xrightarrow{a} p'\gamma' \rangle(t), p(\gamma w, \bar{c})) &:= c'(t, p'(\text{cons}(\gamma', w), \bar{c})) \\ c'(\langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle(t), p(\text{cons}(\gamma, w), \bar{c})) &:= c'(t, p'(\text{cons}(\gamma_1, \text{cons}(\gamma_2, w)), \bar{c})) \\ c'(\langle p\gamma \xrightarrow{a} p' \rangle(t), p(\text{cons}(\gamma, w), \bar{c})) &:= c'(t, p'(w, \bar{c})) \\ c'(\langle p\gamma \xrightarrow{\text{sp}} p'\gamma' \triangleright p_s\gamma_s \rangle(t_1, t_2), p(\text{cons}(\gamma, w), \bar{c})) &:= c'(t_1, p'(\text{cons}(\gamma', w), \text{cons}(c'(t_2), \bar{c}))) \end{aligned}$$

We are now going to define a semantics \rightarrow_r^l on \overline{X} that is compatible with the semantics \mapsto_r^l on configurations. For that we first define $\text{app}_r : \overline{X}' \rightarrow \overline{X}'$ for all $r \in \Delta$ as follows:

$$\begin{aligned} \text{app}_r(\epsilon) &= \begin{cases} \langle r \rangle(\epsilon, \langle p_s, \gamma_s \rangle(\epsilon)) & \text{if } r = p\gamma \xrightarrow{\text{sp}} p'\gamma' \triangleright p_s\gamma_s \\ \langle r \rangle(\epsilon) & \text{otherwise} \end{cases} \\ \text{app}_r(\langle r' \rangle(t)) &= \langle r' \rangle(\text{app}_r(t)) \\ \text{app}_r(\langle r' \rangle(t_1, t_2)) &= \langle r' \rangle(\text{app}_r(t_1), \text{app}_r(t_2)) \end{aligned}$$

Then, for every $r \in \Delta$, $\rightarrow_r(\langle p, \gamma \rangle(t)) := \langle p, \gamma \rangle(\text{app}_r(t))$ iff $\mapsto_r^\epsilon(c'(t))$ is defined.

Similarly to Section 3 we extend \rightarrow_r to \rightarrow_r^l , by applying \rightarrow_r to the subtree of an inlined execution tree representing process $l \in \mathbb{N}_+^*$ of the execution.

An inlined execution tree t is called consistent, if there exists $p \in P, \gamma \in \Gamma, k \in \mathbb{N}, l_1, \dots, l_k \in \mathbb{N}_+^*, r_1, \dots, r_k \in \Delta$ with $t = \rightarrow_{r_k}^{l_k}(\dots \rightarrow_{r_1}^{l_1}(\langle p, \gamma \rangle(\epsilon)) \dots)$.

We now decompose the mapping a from execution trees to action trees into two parts. The first part converts the execution tree into an inlined execution tree and the second part converts the inlined execution tree into an action tree.

The inlined execution tree $\mathbf{a}_1(t)$ of an execution tree t is defined by:

$$\begin{aligned}
\mathbf{a}_1(\langle p, \gamma \rangle) &:= \langle p, \gamma \rangle(\epsilon) \\
\mathbf{a}_1(\langle p\gamma \xrightarrow{a} p'w' \rangle(t)) &:= \langle p, \gamma \rangle(\mathbf{a}_1(\langle p\gamma \xrightarrow{a} p'w' \rangle(t), \epsilon)) \\
\mathbf{a}_1(\langle p\gamma \xrightarrow{a} p'w' \triangleright p_s\gamma_s \rangle(t_1, t_2)) &:= \langle p, \gamma \rangle(\mathbf{a}_1(\langle p\gamma \xrightarrow{a} p'w' \triangleright p_s\gamma_s \rangle(t_1, t_2), \epsilon)) \\
\mathbf{a}_1(\langle p, \gamma \rangle, x) &:= x \\
\mathbf{a}_1(\langle p\gamma \xrightarrow{a} p'\gamma' \rangle(t), x) &:= \langle p\gamma \xrightarrow{a} p'\gamma' \rangle(\mathbf{a}_1(t), x) \\
\mathbf{a}_1(\langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle(t), x) &:= \langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle(\mathbf{a}_1(t), x) \\
\mathbf{a}_1(\langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle(t_1, t_2), x) &:= \langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle(\mathbf{a}_1(t_1), \mathbf{a}_1(t_2), x)) \\
\mathbf{a}_1(\langle p\gamma \xrightarrow{a} p' \rangle, x) &:= \langle p\gamma \xrightarrow{a} p' \rangle(x) \\
\mathbf{a}_1(\langle p\gamma \xrightarrow{\text{sp}} p\gamma \triangleright p_s\gamma_s \rangle(t, t_s), x) &:= \langle p\gamma \xrightarrow{\text{sp}} p\gamma \triangleright p_s\gamma_s \rangle(\mathbf{a}_1(t), x, \mathbf{a}_1(t_s))
\end{aligned}$$

The action tree $\mathbf{a}_2(t)$ of an inlined execution tree $t \in \overline{X}$ is defined as follows:

$$\begin{aligned}
\mathbf{a}_2(\langle p, \gamma \rangle(t)) &:= \mathbf{a}_2(t) \\
\mathbf{a}_2(\langle p\gamma \xrightarrow{a} p'\gamma' \rangle(t)) &:= \langle a \rangle(\mathbf{a}_2(t)) \\
\mathbf{a}_2(\langle p\gamma \xrightarrow{a} p'\gamma_1\gamma_2 \rangle(t)) &:= \langle a \rangle(\mathbf{a}_2(t)) \\
\mathbf{a}_2(\langle p\gamma \xrightarrow{a} p' \rangle(t)) &:= \langle a \rangle(\mathbf{a}_2(t)) \\
\mathbf{a}_2(\langle p\gamma \xrightarrow{\text{sp}} p\gamma \triangleright p_s\gamma_s \rangle(t, t_s)) &:= \langle \text{sp} \rangle(\mathbf{a}_2(t), \mathbf{a}_2(t_s))
\end{aligned}$$

The following statements can be shown by straight-forward inductions:

Lemma 5. 1. $\mathbf{a} = \mathbf{a}_2 \circ \mathbf{a}_1$.

2. $\mathbf{a}_1(\mapsto_r^l(t)) = \mapsto_r^l(\mathbf{a}_1(t))$ for all $t \in X$, all transition rules $r \in \Delta$, and all $l \in \mathbb{N}^*$. Hence, for all $t \in X$, the set of executable schedules of t equals the set of executable schedules of $\mathbf{a}_1(t)$. And $t \in X$ is consistent iff $\mathbf{a}_1(t)$ is consistent.
3. The set of executable schedules of a consistent $t \in \overline{X}$ equals the set of executable schedules of $\mathbf{a}_2(t)$.

The original statement is then a corollary of the above lemma. \square

Proof (of Theorem 2). Assume that check^{jl} does not accept t . Hence, not $\text{ok}(t)$ or $\text{as}(t) = \perp$ or $\text{as}(t) = (A, R, R_{\text{jo}}, U, \rightarrow_A, J)$ with cyclic \rightarrow_A . First of all, if not $\text{ok}(t)$, there is obviously no join-lock sensitive schedule.

Secondly, if $\text{as}(t) = \perp$, we choose a minimal subtree t' of t with $\text{as}(t') = \perp$. This subtree must be of the form $t' = \langle \text{sp} \rangle(t_1, t_2)$ with $\text{as}(t_1) = (A, R, R_{\text{jo}}, U, \rightarrow_A, \top)$ and $\text{as}(t_2) = (A', R', R'_{\text{jo}}, U', \rightarrow_{A'}, J')$ and $R_{\text{jo}} \cap U' \neq \emptyset$ as all other cases yield non- \perp values. $R_{\text{jo}} \cap U' \neq \emptyset$ means that there is a lock x required by the spawned process for termination, but held by the spawning process and released only after the join, such that no join-lock sensitive schedule is possible.

Finally, assume $\text{as}(t) = (A, R, R_{\text{jo}}, U, \rightarrow_A, J)$ with cyclic \rightarrow_A . The edges in the acquisition graph \rightarrow_A capture ordering constraints on final acquisitions of locks. More specifically, $i \rightarrow_A j$ means, that there is a final acquisition of i that must be scheduled before any final acquisition of j . If there is a cycle $i_1 \rightarrow_A \dots \rightarrow_A i_k \rightarrow_A i_1$, $k > 1$, with pairwise disjoint locks i_1, \dots, i_k , the ordering constraints on the final acquisitions of

i_1, \dots, i_k obviously cannot be fulfilled. If there is a cycle $i \rightarrow_A i$ involving just a single lock i , we have to schedule an acquisition of i after a final acquisition of i , which also contradicts lock-sensitivity.

Assume now that check^{jl} accepts t . This implies that $\text{ok}(t)$ holds and that $\text{as}(t) = (A, R, R_{\text{jo}}, U, \rightarrow_A, J)$ with an acyclic \rightarrow_A . For the proof, we first annotate each node of t with its process $l \in \mathbb{N}_+^*$. In addition, spawn nodes are annotated with the spawned process. This results in a tree of the form:

$$\bar{T}^{\text{Act}} ::= \mathbb{N}_+^* \mid \langle \mathbb{N}_+^* \times \text{Act} \rangle (\bar{T}^{\text{Act}}) \mid \langle \mathbb{N}_+^* \times \{\text{sp}\} \times \mathbb{N}_+^* \rangle (\bar{T}^{\text{Act}}, \bar{T}^{\text{Act}})$$

where in a tree of the form $\langle l, \text{sp}, l' \rangle (t, t')$, l is the spawning process and l' is the spawned process. The translation of the action tree t to its annotated version $\bar{t} \in \bar{T}^{\text{Act}}$ is straightforward. A *schedule* of an annotated action tree is a topological ordering of the annotations of its inner nodes. It is easy to see that, from a schedule of \bar{t} a schedule of t is obtained by removing the additional annotations at the spawn actions.

We call a schedule a *u-schedule*, if its lock actions are in the context free language $U ::= \epsilon \mid U \cdot U \mid \text{acq}(1) \cdot U \cdot \text{rel}(1) \mid \dots \mid \text{acq}(n) \cdot U \cdot \text{rel}(n)$. We call a schedule a *p-schedule*, if it is a prefix of a u-schedule. Similar, a schedule is called *s-schedule*, if it is a suffix of a u-schedule, and *sp-schedule* if it is a suffix of a prefix of a u-schedule. Intuitively, u-schedules describe schedules with only matched lock and unlock operations, p-schedules may additionally contain unmatched acquisitions, s-schedules may contain unmatched releases and sp-schedules may contain both, unmatched acquisitions and unmatched releases.

We call a schedule *s non-reentrant*, if its sequence of actions is non-reentrant. Note that non-reentrant p-schedules correspond to well-formed action sequences.

As \rightarrow_A is acyclic, we can choose a total order $<$ on locks s.t. $\rightarrow_A \subseteq <$, i.e. a topological ordering of the locks w.r.t. \rightarrow_A . We now construct inductively for each subtree t' of \bar{t} with $\text{as}(t') = (A', R', R'_{\text{jo}}, U', \rightarrow_{A'}, J')$ a schedule $s(t')$ with the following properties:

1. $s(t')$ is a non-reentrant sp-schedule (i.e. its actions form a suffix of a well-formed sequence).
2. $s(t')$ is join-sensitive, i.e. $s(t') = s_1 \langle l, \text{sp}, l' \rangle s_2 \langle l, \text{jo} \rangle s_3 \implies \langle l', \$ \rangle \in s_2$
3. The order of acquisitions respects $<$ in the following sense:

$$s(t') = s_1 \langle l, \text{acq}(i) \rangle s_2 \langle l', \text{acq}(j) \rangle s_3 \wedge \langle l, \text{rel}(i) \rangle \notin s_2 s_3 \implies i < j$$

i.e. a final acquisition is only followed by acquisitions of greater locks w.r.t. $<$.

4. Only required locks are used before termination of the root process, i.e.

$$s(t') = s_1 \langle l(t'), \$ \rangle s_2 \implies s_1 \text{ is an s-schedule acquiring only locks from } U'$$

where $l(t')$ is the process of subtree t' , i.e. the annotation of its root node.

The schedule constructed for the entire tree \bar{t} then corresponds to a join-lock sensitive schedule of t .

The schedule of an atomic tree $\langle l \rangle$ is the empty sequence, that trivially satisfies properties 1-4.

The schedule of a tree of the form $t' = \langle l, a \rangle(t'')$ is $s(t') = \langle l, a \rangle s(t'')$. Let $\text{as}(t'') = (A'', R'', R''_{\text{jo}}, U'', \rightarrow_{A''}, J'')$. If $a = \text{jo}$, properties 1-4 are trivially preserved. For $a = \text{rel}(i)$, properties 2 and 3 are trivially preserved. Properties 1 and 4 are preserved as sp-schedules as well as s-schedules are closed under prepending of release actions, and non-reentrance of the constructed schedule follows from the well-formedness assumption on t . For $a = \text{acq}(i)$, we distinguish whether $i \in R''$ or not.

Case $i \in R''$ As t is join-lock-well-formed, all unmatched releases in t'' done by process l , and the first unmatched release action must be $\langle l, \text{rel}(i) \rangle$. This implies that $s(t')$ is an sp-schedule. Its non-reentrance follows from non-reentrance of $s(t'')$. Preservation of Property 4 is shown analogously, and preservation of Properties 2 and 3 is trivial.

Case $i \notin R''$ As t is join-lock-well-formed, $s(t'')$ cannot contain unmatched releases. Hence, $s(t')$ is an sp-schedule. If $s(t'')$ would contain an acquisition of i , we would have $(i, i) \in \rightarrow_{A'} \subseteq \rightarrow_A$ in contradiction to acyclicity of \rightarrow_A . Hence, $s(t')$ stays non-reentrant. Preservation of Property 2 is trivial. Preservation of Property 3 follows because for each action $\langle l', \text{acq}(j) \rangle \in s(t'')$, the pair (i, j) is included in $\rightarrow_{A'} \subseteq \rightarrow_A \subseteq <$. Property 4 holds in this case, as $l(t')$ cannot be stopped due to the assumption that processes release all their locks before stopping.

Let us now consider subtrees of the form $t' = \langle l_1, \text{sp}, l_2 \rangle(t_1, t_2)$ with $\text{as}(t_1) = (A_1, R_1, R_{\text{jo}}^1, U_1, \rightarrow_{A_1}^1, J_1)$ and $\text{as}(t_2) = (A_2, R_2, R_{\text{jo}}^2, U_2, \rightarrow_{A_2}^2, J_2)$.

For the proof, a schedule of the form $\langle l, \text{acq}(i) \rangle w$ with a u-schedule w is called a *macro-step*. For a macro-step, we define $\text{acq}(\langle l, \text{acq}(i) \rangle w) := i$. Note that any sp-schedule s can be written as $s = s_1 s_2$ where s_1 is an s-schedule, and s_2 is a sequence of macro-steps. Analogously, every p-schedule s can be written as $s = s_1 s_2$ where s_1 is a u-schedule and s_2 is a sequence of macrosteps.

For sequences of macrosteps, we now define a shuffle-operator \parallel , that constructs from two sequences of macrosteps that respect $<$ an interleaving that also respects $<$:

$$\begin{aligned} \varepsilon \parallel s &:= s \parallel \varepsilon := s \\ s_1 r_1 \parallel s_2 r_2 &:= \begin{cases} s_1(r_1 \parallel s_2 r_2) & \text{if } \text{acq}(s_1) < \text{acq}(s_2) \\ s_2(s_1 r_1 \parallel r_2) & \text{if } \text{acq}(s_2) < \text{acq}(s_1) \\ \text{undefined} & \text{otherwise} \end{cases} \end{aligned}$$

We make a case distinction whether the root process of t_1 , i.e. l_1 , performs a join operation or not.

Case 1: If l_1 does not perform a join operation, we have $J_1 = \perp$ and $\langle l_1, \text{jo} \rangle \notin s(t_1)$. As $s(t_1)$ is an sp-schedule, it can be written as $s(t_1) = s_1 s_2$ for an s-schedule s_1 and a sequence of macrosteps s_2 . Due to well-formedness, t_2 contains no unmatched releases, hence $s(t_2)$ is a p-schedule and can be written as $s(t_2) = s'_1 s'_2$ for an u-schedule s'_1 and a sequence of macrosteps s'_2 .

As $\text{ok}(t)$ holds the unmatched acquisitions in s_2 and s'_2 are disjoint. Hence $s_2 \parallel s'_2$ is defined and we set $s(t') = s_1 s'_1 (s_2 \parallel s'_2)$. It can be checked that this preserves Properties 1-4:

1. Sequences of macro-steps are p-schedules, hence $(s_2 \parallel s'_2)$ is a p-schedule. Prepending the u-schedule s'_1 yields a p-schedule again, and prepending the s-schedule s_1 yields an sp-schedule.

For non-reentrance, we first observe that $s(t')$ can only violate non-reentrance if at least one unmatched operation is involved, as shuffling was done regarding matched operations atomically. Unmatched release operations can occur only in s_1 , hence a violation of non-reentrance would also imply that $s(t_1)$ is reentrant, in contradiction to our assumptions. Unmatched acquisition operations can only occur in $(s_2 \parallel s'_2)$. As $\text{ok}(t)$ holds, the unmatched acquisitions in s_2 and s'_2 are disjoint. Hence, the only possibility to violate non-reentrance is that there is an unmatched acquisition $\text{acq}(i)$ from, say (w.l.o.g.) s_2 conflicting with a matched acquisition from s'_2 , i.e. $s_2 = s_{21}(\langle l, \text{acq}(i) \rangle w) s_{22}$ and $s'_2 = s'_{21}(\langle l', \text{acq}(j) \rangle w') s'_{22}$ with $\langle l', \text{acq}(i) \rangle \in w'$. Due to our inductive assumption on $s(t_2)$, we have $\text{acq}(m') < i$ for each macrostep $m' \in s'_{21}(\langle l', \text{acq}(j) \rangle w')$, hence the shuffle operation will schedule the steps from $s'_{21}(\langle l', \text{acq}(j) \rangle w')$ before the unmatched acquisition of i , preserving non-reentrance.

2. Join-sensitivity is preserved as the root-process contains no join operation, and the relative order of the steps from $s(t_1)$ and $s(t_2)$ are preserved in $s(t')$.
3. This property is preserved by the definition of the \parallel -operation.
4. If the root-process stops, it frees all its locks due to the join-lock-well-formedness assumption. Hence, the action $\langle l(t'), \$ \rangle$ is in s_1 , and Property 4 is obviously preserved.

Case 2: Consider now the case that the root process of t_1 , i.e. l_1 , performs a join operation. Then we have $J_1 = \top$ and $(l_1, \text{jo}) \in s(t_1)$. We can write $s(t_1)$ as $s(t_1) = s_1 s_2$ with an sp-schedule s_1 and a sequence of macrosteps s_2 , and $s(t_2)$ can be written as $s(t_2) = s'_1 \langle l_2, \$ \rangle s'_2 s'_3$ with a u-schedule s'_1 that only uses lock-operations from U_2 , a u-schedule s'_2 and a sequence of macrosteps s'_3 (by the join-lock-well-formedness assumption and Property 4 of $s(t_2)$).

If $\langle l_1, \text{jo} \rangle \notin s_1$, we define $s(t') := s_1 s'_1 \langle l_2, \$ \rangle s'_2 (s_2 \parallel s'_3)$, otherwise, we can write s_1 in the form $s_1 = s_{11} s_{12}$ where s_{11} is an s-schedule of maximal length with $\langle l_1, \text{jo} \rangle \notin s_{11}$, and we define $s(t') := s_{11} s'_1 \langle l_2, \$ \rangle s_{12} s'_2 (s_2 \parallel s'_3)$.

Again, this construction preserves Properties 1-4:

1. That the constructed schedule is an sp-schedule can be proven analogously to Case 1. For preservation of non-reentrance, one has to regard the new case that, for the schedule $s(t') = s_{11} s'_1 \langle l_2, \$ \rangle s_{12} s'_2 (s_2 \parallel s'_3)$, an unmatched release in s_{12} conflicts with a lock operation in s'_1 . However, the unmatched releases in s_{12} are captured by R_{jo}^1 , the lock operations in s'_1 are captured by U_2 and due to $\text{as}(t) \neq \perp$ we have $R_{\text{jo}}^1 \cap U_2 = \emptyset$, hence there are no such conflicts. The other cases are argued analogously to Case 1.
2. For join-operations of the non-root processes, preservation of join-sensitivity is analogous to Case 1. As all steps of the spawned process are scheduled before the first join operation, and the relative ordering of the steps in $s(t_1)$ is preserved, the constructed schedule is join-sensitive.
3. This property is, as in Case 1, preserved by the definition of the \parallel -operation.

4. If the root process l_1 stops, all its actions, in particular the join-operation, are contained in s_1 . Hence we are in the case $s(t') := s_{11}s'_1\langle l_2, \$ \rangle s_{12}s'_2(s_2\parallel s'_3)$, and s_{12} can be written as $s_{121}\langle l_1, \$ \rangle s_{122}$, for s-schedules s_{121} and s_{122} . Hence, $s_{11}s'_1\langle l_2, \$ \rangle s_{121}$ is also an s-schedule, and it only uses locks from $U' = U_1 \cup U_2$. □

Proof (of Theorem 3). Follows from Theorems 1, 2 and Lemma 1, 2, 3, 4.