

Pairwise Reachability Analysis for Higher Order Concurrent Programs by Higher-Order Model Checking

Kazuhide Yasukata, Naoki Kobayashi, and Kazutaka Matsuda

The University of Tokyo, Hongo, Japan

Abstract. We propose a sound, complete, and automatic method for pairwise reachability analysis of higher-order concurrent programs with recursion, nested locks, joins, and dynamic thread creation. The method is based on a reduction to higher-order model checking (i.e., model checking of trees generated by higher-order recursion schemes). It can be considered an extension of Gawlitz et al.’s work on the join-lock-sensitive reachability analysis for dynamic pushdown networks (DPN) to higher-order programs. To our knowledge, this is the first application of higher-order model checking to sound and complete verification of (reasonably expressive models of) concurrent programs.

1 Introduction

Verification of concurrent programs is important but fundamentally difficult, especially in the presence of recursion. Ramalingam [19] has shown that the reachability problem for two pushdown systems with rendezvous-style synchronization primitives is undecidable. There are two major approaches to cope with this limitation. One is to give up the soundness, and underapproximate the actual behavior of a concurrent program by bounding the number of context switches [18], etc. The other approach is to restrict the synchronization primitives. Kahlon et al. [7] have shown that the pairwise reachability problem (“Given two pushdown systems P_1 and P_2 and control locations ℓ_1 and ℓ_2 , is there a reachable global state where P_1 is at ℓ_1 and P_2 is at ℓ_2 ?”) is decidable if the two pushdown systems synchronize only via nested locking. Lammich et al. [14] later extended it to dynamic pushdown networks (DPN), where processes may be dynamically created. Gawlitz et al. [4] have further extended the result to allow synchronization via joins in addition to nested locking.

In the present paper, we follow the latter line of work and extend it to deal with *higher-order* concurrent programs with recursion, dynamic process creation, joins, and nested locking. We consider the pairwise reachability problem: “Given a program P and two control locations ℓ_1 and ℓ_2 , may the program reach a state where one process is at ℓ_1 and another is at ℓ_2 ?”. We show that this problem can be reduced to *higher-order model checking* [15,9], hence it is decidable. The main idea is to transform a given program to a non-deterministic higher-order

recursion scheme (a kind of tree grammar where non-terminals may take higher-order functions as arguments) \mathcal{G} that generates a tree language $\mathcal{L}(\mathcal{G})$ consisting of all the possible execution histories (called *action trees* [14,4]) of the program ignoring the synchronization constraints imposed by joins and nested locking. Let L_1 be the set of action trees that respect the synchronization constraints imposed by joins and nested locking, and L_2 be the set of action trees that represent histories that end with a state where two of the processes are at ℓ_1 and ℓ_2 . Then, ℓ_1 and ℓ_2 are pairwise reachable if and only if $\mathcal{L}(\mathcal{G}) \cap L_1 \cap L_2 \neq \emptyset$. Since both L_1 and L_2 are regular (where the regularity of L_1 is due to [4]), the latter condition can be decided by using higher-order model checking [15,9]. We formalize the reduction and prove its correctness. We also report preliminary experimental results, which confirm that the approach is feasible at least for small programs, despite the extremely high worst-case complexity of higher-order model checking (k -EXPTIME complete for order- k higher-order recursion schemes [15,12]).

To our knowledge, this is the first application of higher-order model checking to sound and *complete* verification of higher-order *concurrent* programs. The previous applications of higher-order model checking were mainly for higher-order *functional* programs [9,13,16]. For concurrent programs, the previous applications [10,5] were based on the underapproximation approach. One may think that higher-order functions are exotic features for multi-threaded programs. As demonstrated in [9,20], however, even if higher-order functions are not so often used explicitly in source programs, they are required for precisely modelling control/data structures such as exceptions and lists.

The rest of the paper is structured as follows. Section 2 introduces the target language and formally defines the pairwise reachability problem. After providing the necessary backgrounds (such as action trees [14,4] and higher-order model checking [15,9]) Section 3 provides the reduction of the pairwise reachability problem to higher-order model checking. Section 4 reports preliminary experiments. Section 5 discusses related work and Section 6 concludes the paper. Omitted proofs are found in the longer version, available from the authors' web page.

2 Target Language and Pairwise Reachability Problem

This section introduces a higher-order concurrent programming language and defines the pairwise reachability problem for it.

Definition 1. A *program* is a finite set of function definitions

$$\{F_1 \tilde{x}_1 = e_1, \dots, F_n \tilde{x}_n = e_n\},$$

where F_i denotes a defined function symbol, and e ranges over the set Exp of expressions, defined by:

$$e ::= \$ \mid x \mid F \mid \mathbf{if}_e \mid e_1 \mid e_2 \mid \mathbf{join}; e \mid \mathbf{acq}_i; e \mid \mathbf{rel}_i; e \mid \mathbf{spawn}(e_c); e \mid e^\ell$$

Here, i ranges over a finite set **Lock** of (non-reentrant) locks, and ℓ ranges over a finite set **Label** of program point labels. Note that the arity of each function may be 0. We require that the function symbols F_1, \dots, F_n are different from each other, and that any program p contains exactly one definition of a “main” function S , of the form $S = e$.

We explain the intuitive meaning of each expression; the formal operational semantics is given later. The expression $\$$ represents the termination of the current process. The expression **if** $_{-} e_1 e_2$ executes either e_1 or e_2 non-deterministically, and the expression $e_1 e_2$ applies the function e_1 to e_2 . As defined later, function calls are based on the call-by-name semantics; call-by-value programs can be transformed to call-by-name programs by using the CPS transformation [17]. The expression **spawn** (e_c) ; e spawns a new child process that executes e_c , and continues to execute e without waiting for the child process. The expression **join**; e waits for the termination of all the processes that the current process has created, and then executes e . The expression **acq** $_i$; e waits to acquire the lock i , and then executes e . The expression **rel** $_i$; e releases the lock i and executes e . The label ℓ in the expression e^ℓ is used for specifying the pairwise reachability problem, and does not affect the operational semantics.

In this paper, we consider only “well-typed” programs, as defined below.

Definition 2 (types). *The set of types is inductively defined by:*

$$\tau ::= \mathbf{unit} \mid \tau_1 \rightarrow \tau_2$$

Here, **unit** is the type of the unit value $\$$, and $\tau_1 \rightarrow \tau_2$ is the type of functions from τ_1 to τ_2 . The **order** and the **arity** of types are inductively defined by:

$$\begin{aligned} \text{order}(\mathbf{unit}) &= 0 & \text{order}(\tau_1 \rightarrow \tau_2) &= \max(\text{order}(\tau_1) + 1, \text{order}(\tau_2)) \\ \text{arity}(\mathbf{unit}) &= 0 & \text{arity}(\tau_1 \rightarrow \tau_2) &= \text{arity}(\tau_2) + 1 \end{aligned}$$

A **type environment** Γ is a map from a finite set of variables to types. The **type judgment relation** $\Gamma \vdash e : \tau$ for expressions is the least relation closed under the following rules:

$$\begin{array}{c} \frac{}{\emptyset \vdash \$: \mathbf{unit}} \quad \frac{}{\Gamma\{x \mapsto \tau\} \vdash x : \tau} \quad \frac{\Gamma \vdash e_1 : \mathbf{unit} \quad \Gamma \vdash e_2 : \mathbf{unit}}{\Gamma \vdash \mathbf{if}_{-} e_1 e_2 : \mathbf{unit}} \\ \frac{\Gamma \vdash e_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma \vdash e_2 : \tau_1}{\Gamma \vdash e_1 e_2 : \tau_2} \quad \frac{\Gamma \vdash e : \mathbf{unit}}{\Gamma \vdash \mathbf{join}; e : \mathbf{unit}} \quad \frac{\Gamma \vdash e : \mathbf{unit}}{\Gamma \vdash \mathbf{acq}_i; e : \mathbf{unit}} \\ \frac{\Gamma \vdash e : \mathbf{unit}}{\Gamma \vdash \mathbf{rel}_i; e : \mathbf{unit}} \quad \frac{\Gamma \vdash e : \mathbf{unit} \quad \Gamma \vdash e_c : \mathbf{unit}}{\Gamma \vdash \mathbf{spawn}(e_c); e : \mathbf{unit}} \quad \frac{\Gamma \vdash e : \mathbf{unit}}{\Gamma \vdash e^\ell : \mathbf{unit}} \end{array}$$

A program $p = \{F_1 x_{11}, \dots, x_{1k_1} = e_1, \dots, F_n x_{n1}, \dots, x_{nk_n} = e_n\}$ is **well-typed** under Γ if $\Gamma = \{F_i \mapsto \tau_{i1} \rightarrow \dots \rightarrow \tau_{ik_i} \rightarrow \mathbf{unit} \mid i \in \{1, \dots, n\}\}$ and $\Gamma \cup \{x_{j1} \mapsto \tau_{j1}, \dots, x_{jk_j} \mapsto \tau_{jk_j}\} \vdash e_j : \mathbf{unit}$ holds for each $j \in \{1, \dots, n\}$. The **order** of p (well-typed under Γ) is $\max(\{\text{order}(\Gamma(F)) \mid F \in \text{dom}(\Gamma)\})$.

Remark 1. In the language above, we have only the unit-value as a base value. As we have higher-order functions, however, we can encode booleans and conditionals by using the Church encoding. Values in infinite data domains (such

as unbounded integers) can be dealt with (soundly but incompletely) by using predicate abstraction [9,13].

Example 1. Here is an example of an order-2 program, well-typed under $\Gamma = \{S \mapsto \mathbf{unit}, F \mapsto (\mathbf{unit} \rightarrow \mathbf{unit}) \rightarrow \mathbf{unit} \rightarrow \mathbf{unit}, G \mapsto \mathbf{unit} \rightarrow \mathbf{unit}, H \mapsto (\mathbf{unit} \rightarrow \mathbf{unit}) \rightarrow \mathbf{unit} \rightarrow \mathbf{unit}\}$.

$$p = \left\{ \begin{array}{ll} S = F G \$ & F g t = \mathbf{if}_-(\mathbf{spawn}(H g \$); F g t) (\mathbf{join}; t) \\ G t = t^\ell & H g t = \mathbf{acq}_1; (g(\mathbf{rel}_1; t)) \end{array} \right\}$$

This program is obtained by CPS-transforming the following C-like code:

```

main() { f(g); }
f(g) {
  if(*) { spawn { h(g); }; f(g); }
  else { join(); }
}
g() {
  L: <critical section>;
  return;
}
h(g) { acq(1); g(); rel(1); }
    
```

The root process spawns non-deterministically many child processes, and then waits for the child processes by the join operation. Each child acquires the lock 1 and then release 1 at the program point ℓ .

We now define the formal semantics of programs. A **configuration** c of a program is a map from a finite set consisting of sequences of natural numbers (where each sequence serves as a process identifier) to the set of triples (e, L, s) consisting of an expression e , a sequence L of locks, and a natural number s . Intuitively, $c(\pi) = (e, i_1 \dots i_k, s)$ means that the process π is executing e , that it holds locks i_1, \dots, i_k that have been acquired in this order, and that it has created s child processes so far. The **reduction relation** $c \longrightarrow_p c'$ on configurations is defined by the following rules.

$$\begin{array}{c}
 \frac{F x_1 \dots x_k \rightarrow e \in p}{c \uplus \{ \pi \mapsto (F e_1 \dots e_k, L, s) \} \longrightarrow_p c \uplus \{ \pi \mapsto ([e_1/x_1, \dots, e_k/x_k]e, L, s) \}} \\
 \\
 c \uplus \{ \pi \mapsto (\mathbf{if}_- e_1 e_2, L, s) \} \longrightarrow_p c \uplus \{ \pi \mapsto (e_i, L, s) \} \quad (i \in \{1, 2\}) \\
 \\
 c \uplus \{ \pi \mapsto (e^\ell, L, s) \} \longrightarrow_p c \uplus \{ \pi \mapsto (e, L, s) \} \\
 \\
 \frac{i \notin \mathbf{locked}(c \uplus \{ \pi \mapsto (\mathbf{acq}_i; e, L, s) \})}{c \uplus \{ \pi \mapsto (\mathbf{acq}_i; e, L, s) \} \longrightarrow_p c \uplus \{ \pi \mapsto (e, L \cdot i, s) \}} \\
 \\
 c \uplus \{ \pi \mapsto (\mathbf{rel}_i; e, L \cdot i, s) \} \longrightarrow_p c \uplus \{ \pi \mapsto (e, L, s) \} \\
 \\
 c \uplus \{ \pi \mapsto (\mathbf{spawn}(e_c); e, L, s) \} \longrightarrow_p c \uplus \{ \pi \mapsto (e, L, s + 1), \pi \cdot s \mapsto (e_c, \epsilon, 0) \} \\
 \\
 \frac{\{ j \mid \pi \cdot j \in \mathit{dom}(c) \} = \emptyset}{c \uplus \{ \pi \mapsto (\mathbf{join}; e, L, s) \} \longrightarrow_p c \uplus \{ \pi \mapsto (e, L, s) \}}
 \end{array}$$

$$c \uplus \{ \pi \mapsto (\$, \epsilon, s) \} \longrightarrow_p c$$

Here, $\mathbf{locked}(c)$ represents the set of all acquired locks, defined by:

$$\mathbf{locked}(c) = \bigcup_{c(\pi)=(e, i_1 \dots i_k, s)} \{ i_1, \dots, i_k \}.$$

Definition 3 (pairwise reachability). *Let p be a (well-typed) program and ℓ_1, ℓ_2 be labels. We say that (ℓ_1, ℓ_2) is **pairwise reachable** by p , written $p \models \ell_1 \parallel \ell_2$, if there exist c, π_1, π_2 ($\pi_1 \neq \pi_2$) such that $\{ \epsilon \mapsto (S, \epsilon, 0) \} \longrightarrow_p^* c$ with $c(\pi_1) = (e_1^{\ell_1}, L_1, s_1)$ and $c(\pi_2) = (e_2^{\ell_2}, L_2, s_2)$ for some $e_1, e_2, L_1, L_2, s_1, s_2$. The **pairwise reachability problem** is the decision problem of checking whether $p \models \ell_1 \parallel \ell_2$ holds.*

Example 2. Recall the example program showed in Example 1. The verification problem “can the program point ℓ (i.e., the critical section) be reached by multiple processes simultaneously?” can be reduced to the pairwise reachability for the program p and the pair (ℓ, ℓ) . In this case, the answer to the problem is “no”.

3 From Pairwise Reachability to Higher-Order Model Checking

In this section, we show that the pairwise reachability problem can be reduced to higher-order model checking [15], hence it is decidable. The basic idea of this reduction is to transform a program to a grammar called a higher-order recursion scheme [15], which generates *action trees* [4,14] that represent all the possible executions of the program. Since the set of action trees that represent valid executions (i.e., those that respect synchronization constraints on joins and nested locks) is regular, the pairwise reachability problem can be reduced to an inclusion problem between the tree language generated by the higher-order recursion scheme and the regular language, which can be further reduced to higher-order model checking. We first review action trees and higher-order model checking in Sections 3.1 and 3.2 respectively. We then present the reduction from the pairwise reachability analysis to higher-order model checking.

3.1 Action Trees

An action tree [4,14] is a finite tree that represents a history of executions of a program up to a certain state. It ignores how the executions of multiple processes are interleaved, and expresses only process-wise execution histories and the parent/child relationship between processes. Gawlitza et al. [4] originally introduced action trees to represent execution histories of dynamic pushdown networks, but the notion of action trees is independent of a particular computation model. Here we use them to represent execution histories of higher-order concurrent programs introduced in the previous section.

Definition 4 (action trees). The set of **action trees**, ranged over by γ , is defined inductively by:

$$\gamma ::= \perp \mid \langle \$ \rangle \mid \ell \mid \langle \text{jo} \rangle \gamma \mid \langle \text{sp} \rangle \gamma_1 \gamma_2 \mid \langle \text{Acq}_i \rangle \gamma \mid \langle \text{Rel}_i \rangle \gamma.$$

Here, ℓ ranges over (program) labels and i over locks.

We have used the term representation of (labelled) trees above. Trees can also be considered as map from paths to labels, by: $(a \gamma_1 \cdots \gamma_n)^\# = \{ \epsilon \mapsto a \} \cup \{ i \cdot u \mapsto b \mid \gamma_i^\#(u) = b \}$.

Each non-leaf node represents an action of each process. The tree $\langle \text{jo} \rangle \gamma$ ($\langle \text{Acq}_i \rangle \gamma$ and $\langle \text{Rel}_i \rangle \gamma$ respectively) means that the process performed join (acquires and releases the lock i , respectively) and then behaved like γ . The tree $\langle \text{sp} \rangle \gamma_1 \gamma_2$ means that the process spawned a child process that behaved like γ_2 , and the process itself behaved like γ_1 . Thus, the leftmost path from the root node in an action tree represents a sequence of actions performed by the root process, and each leftmost path from the second child of a $\langle \text{sp} \rangle$ -node represents a sequence of actions performed by the spawned process. Each leaf node of an action tree represents the current state of each process: $\langle \$ \rangle$ means that the process has terminated, ℓ means that the process is at the program point ℓ , and \perp means that the process is at a program point not labeled by any element of Label. Figure 1 shows an example of an action tree. It represents an execution history where the root process spawns a child process, waits for the child, and terminates (represented by $\langle \$ \rangle$), and the child process acquires and releases the lock i , and then terminates. It corresponds to the following execution of the program in Example 1:




Fig. 1. Action tree

$$\begin{aligned} & \{ \epsilon \mapsto (S, \epsilon, 0) \} \longrightarrow \{ \epsilon \mapsto (FG \$, \epsilon, 0) \} \longrightarrow^* \{ \epsilon \mapsto (\text{spawn}(HG \$); FG \$, \epsilon, 0) \} \\ & \longrightarrow \{ \epsilon \mapsto (FG \$, \epsilon, 1), 0 \mapsto (HG \$, \epsilon, 0) \} \\ & \longrightarrow^* \{ \epsilon \mapsto (\text{join}; \$, \epsilon, 1), 0 \mapsto (\text{acq}_1; G(\text{rel}_1); \$), \epsilon, 0 \} \\ & \longrightarrow^* \{ \epsilon \mapsto (\text{join}; \$, \epsilon, 1), 0 \mapsto (G(\text{rel}_1); \$), 1, 0 \} \\ & \longrightarrow^* \{ \epsilon \mapsto (\text{join}; \$, \epsilon, 1), 0 \mapsto ((\text{rel}_1); \$)^\ell, 1, 0 \} \\ & \longrightarrow^* \{ \epsilon \mapsto (\text{join}; \$, \epsilon, 1), 0 \mapsto (\$, \epsilon, 0) \} \\ & \longrightarrow \{ \epsilon \mapsto (\text{join}; \$, \epsilon, 1) \} \longrightarrow \{ \epsilon \mapsto (\$, \epsilon, 1) \} \longrightarrow \emptyset \end{aligned}$$

Note that not every action tree represents a valid execution history. For example, consider the action tree: $\langle \text{sp} \rangle (\langle \text{Acq}_1 \rangle \ell_1) (\langle \text{Acq}_1 \rangle \ell_2)$. It represents a state where the parent and child processes are at program points ℓ_1 and ℓ_2 respectively, after *both* having acquired the lock 1 (and not released it yet), which is obviously impossible. In order to exclude action trees that do not respect synchronization constraints, we introduce a transition system on *abstract configurations*, obtained by removing expressions from configurations introduced in the previous section.

Definition 5. An **abstract configuration** is a map from a finite set consisting of sequences of natural numbers (where each sequence serves as a process identifier) to the set of pairs (L, s) consisting of a sequence L of locks, and a natural number s . The transition relation on abstract configurations is defined by:

$$\begin{array}{c}
 \frac{i \notin \mathbf{locked}(c \uplus \{ \pi \mapsto (L, s) \})}{c \uplus \{ \pi \mapsto (L, s) \} \xrightarrow{\pi, \langle \text{Acq}_i \rangle} c \uplus \{ \pi \mapsto (L \cdot i, s) \}} \\
 \\
 c \uplus \{ \pi \mapsto (L \cdot i, s) \} \xrightarrow{\pi, \langle \text{Rel}_i \rangle} c \uplus \{ \pi \mapsto (L, s) \} \\
 \\
 c \uplus \{ \pi \mapsto (L, s) \} \xrightarrow{\pi, \langle \text{sp} \rangle} c \uplus \{ \pi \mapsto (L, s + 1), \pi \cdot s \mapsto (\epsilon, 0) \} \\
 \\
 \frac{\{ k \mid \pi \cdot k \in \text{dom}(c) \} = \emptyset}{c \uplus \{ \pi \mapsto (L, s) \} \xrightarrow{\pi, \langle \text{jo} \rangle} c \uplus \{ \pi \mapsto (L, s) \}} \\
 \\
 c \uplus \{ \pi \mapsto (\epsilon, s) \} \xrightarrow{\pi, \langle \$ \rangle} c
 \end{array}$$

Here, $\mathbf{locked}(c)$ is defined similarly to that for configurations, as

$$\mathbf{locked}(c) = \bigcup_{c(\pi) = (i_1 \dots i_k, s)} \{ i_1, \dots, i_k \}.$$

Each action tree can be mapped to an abstract configuration as follows.

$$\begin{array}{l}
 \theta_{\pi, L, s}(\langle \$ \rangle) = \emptyset \quad \theta_{\pi, L, s}(\langle \text{sp} \rangle \gamma_1 \gamma_2) = \theta_{\pi, L, s+1}(\gamma_1) \cup \theta_{\pi \cdot s, \epsilon, 0}(\gamma_2) \\
 \theta_{\pi, L, s}(\langle \text{jo} \rangle \gamma) = \theta_{\pi, L, s}(\gamma) \quad \theta_{\pi, L, s}(\langle \text{Acq}_i \rangle \gamma) = \theta_{\pi, L \cdot i, s}(\gamma) \\
 \theta_{\pi, L \cdot i, s}(\langle \text{Rel}_i \rangle \gamma) = \theta_{\pi, L, s}(\gamma) \quad \theta_{\pi, L, s}(\gamma) = \{ \pi \mapsto (L, s) \} \text{ (if } \gamma \in \{ \perp \} \cup \text{Label})
 \end{array}$$

We write $\theta(t)$ for $\theta_{\epsilon, \epsilon, s}(t)$, and write $\gamma_1 \xrightarrow{u} \gamma_2$ if $\gamma_1(u) = \perp$ and γ_2 is obtained from γ_1 by replacing \perp at u with a tree of the form $a \perp \dots \perp$ with $a = \langle \$ \rangle, \langle \text{sp} \rangle, \langle \text{jo} \rangle, \langle \text{Acq}_i \rangle, \langle \text{Rel}_i \rangle$.

We can now define “valid” action trees as follows.

Definition 6. An action tree γ is **join-lock sensitive** if there is a sequence $\perp = \gamma_0 \xrightarrow{u_1} \gamma_1 \xrightarrow{u_2} \dots \xrightarrow{u_n} \gamma_n = \gamma'$ such that $\theta(\gamma_0) \xrightarrow{\text{pn}(u_1, \gamma), \gamma_1(u_1)} \theta(\gamma_1) \xrightarrow{\text{pn}(u_2, \gamma), \gamma_2(u_2)} \dots \xrightarrow{\text{pn}(u_n, \gamma), \gamma_n} \theta(\gamma_n)$, where γ' is the action tree obtained from γ by replacing all $\ell \in \text{Label}$ with \perp . Here, $\text{pn}(u, \gamma)$ represents the identifier of the process that executes the action of $\gamma(u)$. It is defined by:

$$\begin{array}{ll}
 \text{pn}(\epsilon, \gamma) & = \epsilon & \text{cn}(\epsilon, \gamma) & = 0 \\
 \text{pn}(u \cdot 1, \gamma) & = \text{pn}(u, \gamma) & \text{cn}(u \cdot 1, \gamma) & = \text{cn}(u, \gamma) \quad \text{if } \gamma(u) \neq \langle \text{sp} \rangle \\
 \text{pn}(u \cdot 2, \gamma) & = \text{pn}(u, \gamma) \cdot \text{cn}(u, \gamma) & \text{cn}(u \cdot 1, \gamma) & = \text{cn}(u, \gamma) + 1 \quad \text{if } \gamma(u) = \langle \text{sp} \rangle \\
 & & \text{cn}(u \cdot 2, \gamma) & = 0
 \end{array}$$

Note that $\text{pn}(u_i, \gamma_i) = \text{pn}(u_i, \gamma)$.

The following is the key property of action trees, which we use in our reduction from pairwise reachability analysis to higher-order model checking.

Lemma 1 (Gawlitza et al. 2011, Section 5 [4]). *The set $L_{\text{sensitive}}$ of join-lock-sensitive action trees is a regular tree language.*

Remark 2. We have modified the original definition of join-lock sensitive (schedulable) action trees. Our notion of join-lock sensitive action trees corresponds to join-lock-well-formed, join-lock sensitive schedulable action trees [4].

3.2 Higher-Order Model Checking

In this subsection, we review *higher-order recursion schemes* and (a variation of) *higher-order model checking* [15].

The set of **sorts** is given by the grammar: $\kappa ::= \circ \mid \kappa_1 \rightarrow \kappa_2$. Intuitively \circ describes trees, and $\kappa_1 \rightarrow \kappa_2$ describes functions from κ_1 to κ_2 . The **order** of sorts is defined by: $\text{order}(\circ) = 0$ and $\text{order}(\kappa_1 \rightarrow \kappa_2) = \max(\text{order}(\kappa_1) + 1, \text{order}(\kappa_2))$.

Definition 7 (Higher-Order Recursion Scheme). *A (non-deterministic) higher-order recursion scheme (HORS, for short) is a quadruple: $\mathcal{G} = (\Sigma, \mathcal{N}, \mathcal{R}, S)$ where Σ is a ranked alphabet (i.e., a map from a finite set of symbols called **terminals** to their arities); \mathcal{N} is a map from a finite set of symbols called **non-terminals** to sorts; $S \in \text{dom}(\mathcal{N})$ is the start symbol of sort \circ ; and \mathcal{R} is a finite set of **transition rules** of the form $A x_1 \cdots x_\ell \rightarrow t$, where t ranges over the set of applicative terms defined by $t ::= \circ \mid a \mid A \mid t_1 t_2$. Here, a ranges over $\text{dom}(\Sigma)$ and A ranges over $\text{dom}(\mathcal{N})$. If $A x_1 \cdots x_\ell \rightarrow t \in \mathcal{R}$, then $\mathcal{N}(A)$ must be of the form $\kappa_1 \rightarrow \cdots \rightarrow \kappa_\ell \rightarrow \circ$ and $\mathcal{N} \cup \{x_1 : \kappa_1, \dots, x_\ell : \kappa_\ell\} \vdash_\Sigma t : \circ$ must be derivable by using the following rules (where non-terminals are treated as variables).*

$$\begin{array}{c} \mathcal{K} \vdash_\Sigma a : \underbrace{\circ \rightarrow \cdots \rightarrow \circ}_{\Sigma(a)} \rightarrow \circ \qquad \mathcal{K} \vdash_\Sigma x : \mathcal{K}(x) \\ \\ \frac{\mathcal{K} \vdash_\Sigma t_1 : \kappa_1 \rightarrow \kappa_2 \quad \mathcal{K} \vdash_\Sigma t_2 : \kappa_1}{\mathcal{K} \vdash_\Sigma t_1 t_2 : \kappa_2} \end{array}$$

The order of a HORS \mathcal{G} , written $\text{order}(\mathcal{G})$ is $\max(\{\text{order}(A) \mid A \in \text{dom}(\mathcal{N})\})$.

Note that unlike *deterministic* HORS, there may be an arbitrary number of rewriting rules for each non-terminal. We omit the adjective ‘non-deterministic’ in the rest of this paper.

To define the rewriting relation, we define the notion of the reduction context.

Definition 8. *The set of **reduction contexts** is defined by:*

$$C ::= [] \mid a t_1 \dots t_{i-1} C t_{i+1} \dots t_n$$

For a reduction context C , we write $C[t]$ for the term obtained from C by replacing $[]$ with t .

Then, the rewriting relation $\longrightarrow_{\mathcal{G}}$ on terms is defined by:

$$C[A t_1 \dots t_m] \longrightarrow_{\mathcal{G}} C[[t_1/x_1, \dots, t_m/x_m]t] \quad (\text{if } A x_1 \dots x_m \rightarrow t)$$

In this paper, we consider a HORS as a generator of a language of finite trees, rather than that of an infinite tree [15].

Definition 9 (Tree Languages of Recursion Schemes). *Let $\mathcal{G} = (\Sigma, \mathcal{N}, \mathcal{R}, S)$ be a HORS. The language generated by \mathcal{G} is defined by:*

$$\mathcal{L}(\mathcal{G}) = \{t \mid t \text{ is a ranked } \Sigma\text{-labeled tree and } S \longrightarrow_{\mathcal{G}}^* t\}$$

Example 3. Consider an order-2 HORS $\mathcal{G} = (\Sigma, \mathcal{N}, \mathcal{R}, S)$ such that

$$\begin{aligned} \Sigma &= \{a \mapsto 2, b \mapsto 1, c \mapsto 0\} \\ \mathcal{N} &= \{S \mapsto \circ, F \mapsto (\circ \rightarrow \circ) \rightarrow \circ \rightarrow \circ, G \mapsto \circ \rightarrow \circ, T \mapsto (\circ \rightarrow \circ) \rightarrow \circ \rightarrow \circ\} \\ \mathcal{R} &= \{S \rightarrow F G c, F g x \rightarrow a(g x)(F(T g)x), F g x \rightarrow g x, \\ &\quad G x \rightarrow b x, T g x \rightarrow g(g x)\} \end{aligned}$$

Here is an example of reduction from S to $a(bc)(b^2c)$.

$$\begin{aligned} S &\longrightarrow_{\mathcal{G}} F G c \longrightarrow_{\mathcal{G}} a(Gc)(F(TG)c) \longrightarrow_{\mathcal{G}} a(bc)(F(TG)c) \longrightarrow_{\mathcal{G}} a(bc)(TGc) \\ &\longrightarrow_{\mathcal{G}} a(bc)(G(Gc)) \longrightarrow_{\mathcal{G}}^* a(bc)(b(bc)) \end{aligned}$$

The language $\mathcal{L}(\mathcal{G})$ is $\{a(bc)(a(b^2c)(a \dots (a(b^{2^{n-1}}c)(b^{2^n}c)) \dots)) \mid n \in \mathbb{N}_+\}$.

The following theorem is an easy corollary of Ong's result on the model checking of (deterministic) HORS [15].

Theorem 1. *Given a HORS \mathcal{G} and a regular tree language L , it is decidable whether $\mathcal{L}(\mathcal{G}) \subseteq L$.*

In the present paper, we call the inclusion problem above a higher-order model checking problem. The standard model checking problem for HORS [15] is the problem of deciding whether the infinite tree generated by a deterministic HORS satisfies a given property.

3.3 Reduction from Pairwise Reachability to Higher-Order Model Checking

Now we show how to reduce pairwise reachability to higher-order model checking. First, we define a transformation from a concurrent program to a HORS that generates action trees of the which join-lock sensitive subset represent all and only the possible reachable configurations of the program.

Definition 10. *Let $p = \{F_1 x_{11}, \dots, x_{1k_1} = e_1, \dots, F_n x_{n1}, \dots, x_{nk_n} = e_n\}$ be a well-typed higher-order concurrent program under Γ . A (non-deterministic)*

HORS \mathcal{G}_p is defined by:

$$\begin{aligned}
 \Sigma &= \{ \langle \text{sp} \rangle \mapsto 2, \langle \text{jo} \rangle \mapsto 1, \langle \$ \rangle \mapsto 0, \perp \mapsto 0 \} \\
 &\cup \{ \langle \text{Acq}_i \rangle, \langle \text{Rel}_i \rangle \mid i \in \text{Lock} \} \cup \{ \ell \mapsto 0 \mid \ell \in \text{Label} \} \\
 \mathcal{G}_p &= (\Sigma, \mathcal{N}_\Gamma \cup \mathcal{N}_0, \mathcal{R}_\Gamma \cup \mathcal{R}_0, S) \\
 \mathcal{N}_\Gamma &= \{ F_1 \mapsto (\Gamma(F_1))^\sharp, \dots, F_n \mapsto (\Gamma(F_n))^\sharp \} \\
 \mathcal{N}_0 &= \{ E_\$ \mapsto \circ, E_{\text{if}_-} \mapsto (\circ \rightarrow \circ \rightarrow \circ), E_{\text{join}} \mapsto (\circ \rightarrow \circ), E_{\text{spawn}} \mapsto (\circ \rightarrow \circ \rightarrow \circ) \} \\
 &\cup \{ E_{\text{acq}_i} \mapsto (\circ \rightarrow \circ) \mid i \in \text{Lock} \} \cup \{ E_{\text{rel}_i} \mapsto (\circ \rightarrow \circ) \mid i \in \text{Lock} \} \\
 &\cup \{ E_\ell \mapsto (\circ \rightarrow \circ) \mid \ell \in \text{Label} \} \\
 \mathcal{R}_\Gamma &= \{ F_1 \tilde{x}_1 \rightarrow \mathcal{E}(e_1), \dots, F_n \tilde{x}_n \rightarrow \mathcal{E}(e_n) \} \\
 \mathcal{R}_0 &= \{ E_{\text{if}_-} x y \rightarrow x, E_{\text{if}_-} x y \rightarrow y, E_{\text{join}} x \rightarrow \langle \text{jo} \rangle x, E_{\text{spawn}} x y \rightarrow \langle \text{sp} \rangle x y \} \\
 &\cup \{ E_{\text{acq}_i} x \rightarrow \langle \text{Acq}_i \rangle x \mid i \in \text{Lock} \} \cup \{ E_{\text{rel}_i} x \rightarrow \langle \text{Rel}_i \rangle x \mid i \in \text{Lock} \} \\
 &\cup \{ E_\ell x \rightarrow \ell \mid \ell \in \text{Label} \} \cup \{ E_\ell x \rightarrow x \mid \ell \in \text{Label} \} \cup \{ E_\$ \rightarrow \langle \$ \rangle \} \\
 &\cup \{ E \tilde{x} \rightarrow \perp \mid E \in \text{dom}(\mathcal{N}_\Gamma \cup \mathcal{N}_0) \setminus \{ E_\ell \mid \ell \in \text{Label} \} \}
 \end{aligned}$$

Here, $(\cdot)^\sharp$ is a transformation from types of expressions to sorts, defined by: $\text{unit}^\sharp = \circ$ and $(\tau_1 \rightarrow \tau_2)^\sharp = \tau_1^\sharp \rightarrow \tau_2^\sharp$. The function \mathcal{E} transforms an expression to an applicative term, defined inductively by:

$$\begin{aligned}
 \mathcal{E}(\$) &= E_\$ & \mathcal{E}(x) &= x & \mathcal{E}(F) &= F & \mathcal{E}(\text{if}_{-} e_1 e_2) &= E_{\text{if}_-} \mathcal{E}(e_1) \mathcal{E}(e_2) \\
 \mathcal{E}(e_1 e_2) &= \mathcal{E}(e_1) \mathcal{E}(e_2) & \mathcal{E}(\text{join}; e) &= E_{\text{join}} \mathcal{E}(e) & \mathcal{E}(\text{acq}_i; e) &= E_{\text{acq}_i} \mathcal{E}(e) \\
 \mathcal{E}(\text{rel}_i; e) &= E_{\text{rel}_i} \mathcal{E}(e) & \mathcal{E}(\text{spawn}(e_c); e) &= E_{\text{spawn}} \mathcal{E}(e) \mathcal{E}(e_c) & \mathcal{E}(e^\ell) &= E_\ell \mathcal{E}(e)
 \end{aligned}$$

The idea of the transformation above is quite simple: just replace each synchronization primitive op with a non-terminal E_{op} , which will generate a tree node $\langle \text{op} \rangle$ indicating that the operation op has been performed. Additionally, in order to generate all the intermediate states of an execution, we allow each non-terminal to be reduced to \perp or $\ell \in \text{Label}$. Note that $\text{order}(\Gamma(F)) = \text{order}(\mathcal{N}(F))$ holds for every function symbol F of p . Thus, $\text{order}(\mathcal{G}_p) = \max(\text{order}(p), 1)$.

Example 4. The program p of Example 1 is transformed to the recursion scheme $\mathcal{G}_p = (\Sigma, \mathcal{N}_\Gamma \cup \mathcal{N}_0, \mathcal{R}_\Gamma \cup \mathcal{R}_0, S)$ where

$$\begin{aligned}
 \mathcal{N}_\Gamma &= \{ S \mapsto \circ, F \mapsto (\circ \rightarrow \circ) \rightarrow \circ \rightarrow \circ, G \mapsto \circ \rightarrow \circ, H \mapsto (\circ \rightarrow \circ) \rightarrow \circ \rightarrow \circ \} \\
 \mathcal{R}_\Gamma &= \left\{ \begin{array}{l} S \rightarrow F G E_\$ \\ F g t \rightarrow E_{\text{if}_-} (E_{\text{spawn}} (F g t) (H g E_\$)) (E_{\text{join}} t) \\ G t \rightarrow E_\ell t \\ H g t \rightarrow E_{\text{acq}_i} (g (E_{\text{rel}_i} t)) \end{array} \right\}.
 \end{aligned}$$

The action tree in Figure 1 is generated by the following reduction sequence:

$$\begin{aligned}
 S &\longrightarrow F G E_\$ \longrightarrow E_{\text{if}_-} (E_{\text{spawn}} (F G E_\$) (H G E_\$)) (E_{\text{join}} E_\$) \\
 &\longrightarrow E_{\text{spawn}} (F G E_\$) (H G E_\$) \longrightarrow \langle \text{sp} \rangle (F G E_\$) (H G E_\$) \\
 &\longrightarrow^* \langle \text{sp} \rangle (E_{\text{if}_-} (E_{\text{spawn}} (F G E_\$) (H G E_\$)) (E_{\text{join}} E_\$)) (E_{\text{acq}_i} (G (E_{\text{rel}_i} E_\$))) \\
 &\longrightarrow^* \langle \text{sp} \rangle (E_{\text{join}} E_\$) (\langle \text{Acq}_i \rangle (G (E_{\text{rel}_i} E_\$))) \longrightarrow^* \langle \text{sp} \rangle (\langle \text{jo} \rangle E_\$) (\langle \text{Acq}_i \rangle (E_{\text{rel}_i} E_\$)) \\
 &\longrightarrow^* \langle \text{sp} \rangle (\langle \text{jo} \rangle \langle \$ \rangle) (\langle \text{Acq}_i \rangle (\langle \text{Rel}_i \rangle E_\$)) \longrightarrow \langle \text{sp} \rangle (\langle \text{jo} \rangle \langle \$ \rangle) (\langle \text{Acq}_i \rangle (\langle \text{Rel}_i \rangle \langle \$ \rangle))
 \end{aligned}$$

Now we show that the grammar generates all the action trees that corresponds to the reachable configurations of a program. We first prepare some definitions. To clarify the relationship between an applicative term (consisting of terminals and non-terminals of a HORS) and a configuration, we first define the terms that can appear at “run-time” and thus have unique corresponding configurations.

Definition 11. *An applicative term t is called a **run-time term** if (1) t has sort \mathfrak{o} , (2) t contains no labels nor \perp , and (3) no terminals occur in the arguments of non-terminals in t .*

Definition 12. *Let t be a run-time term of sort \mathfrak{o} . The action tree t^\perp is defined by:*

$$(E_\ell t)^\perp = \ell \quad (A t_1 \dots t_n)^\perp = \perp \quad (\text{if } A \notin \{E_\ell \mid \ell \in \text{Label}\}) \\ (a t_1 \dots t_n)^\perp = a t_1^\perp \dots t_n^\perp$$

We extend the map $\mathcal{X}_{\pi,L,S}(\cdot)$ to that on run-time trees, by:

$$\mathcal{X}_{\pi,L,S}(\langle \text{sp} \rangle t_1 t_2) = \mathcal{X}_{\pi,L,S+1}(t_1) \cup \mathcal{X}_{\pi \cdot s, \epsilon, 0}(t_2) \quad \mathcal{X}_{\pi,L,S}(\langle \text{jo} \rangle t) = \mathcal{X}_{\pi,L,S}(t) \\ \mathcal{X}_{\pi,L,S}(\langle \$ \rangle) = \emptyset \quad \mathcal{X}_{\pi,L,S}(\langle \text{Acq}_i \rangle t) = \mathcal{X}_{\pi,L,i,S}(t) \quad \mathcal{X}_{\pi,L,i,S}(\langle \text{Rel}_i \rangle t) = \mathcal{X}_{\pi,L,S}(t) \\ \mathcal{X}_{\pi,L,S}(t) = \{ \pi \mapsto (\mathcal{E}^{-1}(t), L, s) \} \quad (\text{for the other cases})$$

Here, \mathcal{E}^{-1} is the inverse of \mathcal{E} . We write $\mathcal{X}(t)$ for $\mathcal{X}_{\epsilon,\epsilon,0}(t)$.

The following lemmas establish the correspondence between p and \mathcal{G}_p .

Lemma 2. *Suppose $S \rightarrow_{\mathcal{G}_p}^* t$ where t is a run-time term and t^\perp is join-lock sensitive. Then, $\{ \epsilon \mapsto (S, \epsilon, 0) \} \rightarrow_p^* \mathcal{X}(t)$ holds.*

Lemma 3. *If $\{ \epsilon \mapsto (S, \epsilon, 0) \} \rightarrow_p^* c$, then there exists a run-time term t such that $S \rightarrow_{\mathcal{G}_p}^* t$, $\mathcal{X}(t) = c$, and t^\perp is join-lock-sensitive.*

By the above lemmas, the pairwise reachability problem on p is reduced to higher-order model checking on \mathcal{G}_p .

Theorem 2. *Let p be a program and (ℓ_1, ℓ_2) be a pair of labels. Let L_{ℓ_1, ℓ_2} be the set $\{ \gamma \mid \exists u_1, u_2. \gamma(u_1) = \ell_1 \wedge \gamma(u_2) = \ell_2 \wedge u_1 \neq u_2 \}$ of action trees. Then, $p \models \ell_1 \parallel \ell_2$ if and only if $\mathcal{L}(\mathcal{G}_p) \not\subseteq \overline{L_{\text{sensitive}}} \cup \overline{L_{\ell_1, \ell_2}}$ holds.*

Proof. Suppose $p \models \ell_1 \parallel \ell_2$. Then there exists c such that $\{ (S, \epsilon, 0) \} \rightarrow_p^* c$ with $c(\pi_1) = (e_1^{\ell_1}, L_1, s_1)$, $c(\pi_2) = (e_2^{\ell_2}, L_2, s_2)$, and $\pi_1 \neq \pi_2$. By Lemma 3, there exists a run-time term t such that $S \rightarrow_{\mathcal{G}_p}^* t$, $\mathcal{X}(t) = c$, and t^\perp is join-lock-sensitive. By the conditions $\mathcal{X}(t) = c$ and $t^\perp, t^\perp \in L_{\ell_1, \ell_2}$. Since $S \rightarrow_{\mathcal{G}_p}^* t \rightarrow_{\mathcal{G}_p}^* t^\perp$, we have $t^\perp \in \mathcal{L}(\mathcal{G}_p) \cap L_{\text{sensitive}} \cap L_{\ell_1, \ell_2}$, i.e., $\mathcal{L}(\mathcal{G}_p) \not\subseteq \overline{L_{\text{sensitive}}} \cup \overline{L_{\ell_1, \ell_2}}$.

Conversely, suppose $\mathcal{L}(\mathcal{G}_p) \not\subseteq \overline{L_{\text{sensitive}}} \cup \overline{L_{\ell_1, \ell_2}}$, i.e., $\gamma \in \mathcal{L}(\mathcal{G}_p) \cap L_{\text{sensitive}} \cap L_{\ell_1, \ell_2}$ for some action tree γ . Then there exists a run-time term t such that $t^\perp = \gamma$ and $S \rightarrow_{\mathcal{G}_p}^* t$. By Lemma 2, we have $\{ \epsilon \mapsto (S, \epsilon, 0) \} \rightarrow_p^* \mathcal{X}(t)$. By the conditions $t^\perp = \gamma$ and $\gamma \in L_{\ell_1, \ell_2}$, t has two distinct sub-terms (at redex-positions) of the form $E_{\ell_1} t_1$ and $E_{\ell_2} t_2$. Thus, there exist π_1, π_2 such that $\mathcal{X}(t)(\pi_1) = (e_1^{\ell_1}, L_1, s_1)$ and $\mathcal{X}(t)(\pi_2) = (e_2^{\ell_2}, L_2, s_2)$ with $\pi_1 \neq \pi_2$. Therefore, we have $p \models \ell_1 \parallel \ell_2$ as required. \square

Because $\overline{L_{\text{sensitive}}} \cup \overline{L_{\ell_1, \ell_2}}$ is a regular tree language, by Theorems 1 and 2, the pairwise reachability $p \models \ell_1 \parallel \ell_2$ is decidable.

Complexity. Recall that the order of \mathcal{G}_p is $\max(\text{order}(p), 1)$ and the model checking of order- k HORS is k -EXPTIME [15]. Therefore, the pairwise reachability analysis for order- k programs is k -EXPTIME for $k \geq 1$. As for the lower-bound, the problem of checking whether $\ell \in \mathcal{L}(\mathcal{G})$ (where ℓ is a singleton tree consisting of the leaf ℓ) is already $(k - 1)$ -EXPTIME-hard [12]. Since it can be easily reduced to a pairwise reachability analysis problem, the pairwise reachability is $(k - 1)$ -EXPTIME-hard. It should be noted, however, that if a regular tree language L is fixed, and also if both the largest arity and order of symbols are fixed, then $\mathcal{L}(\mathcal{G}) \subseteq L$ can be decided in time linear in the size of \mathcal{G} [11]. In our method, the regular tree language $\overline{L}_{\text{sensitive}} \cup \overline{L}_{\ell_1, \ell_2}$ is determined by the sets `Label` and `Lock`. We can fix `Label` as $\{\ell_1, \ell_2\}$ by omitting the other labels from the input program. Therefore, if we fix (i) the largest arity and order of functions in a program and (ii) the number of locks used in the program (i.e., $|\text{Lock}|$), then the pairwise reachability can also be decided in time linear in the size of the program.

4 Preliminary Experiments

We carried out preliminary experiments to check the feasibility of our verification method. As the underlying model checker, we used HorSat [1]. At the time of writing this paper, we have not yet fully automated the translation from programs to HORS, but doing so is not difficult.

Table 1 shows the experimental results. The column “Order” and “# of functions” indicate the order and the number of function definitions of each program. The rightmost column shows the times spent for higher-order model checking (excluding those for translations, which can be performed instantly once automated). Note that the Reachability column shows the answers for pairwise reachability, not for the original verification problems. The benchmark program `example` has been taken from Example 1. The program `example_wrong` is a variation of `example`, obtained by omitting `join` operation of F . The other benchmark programs have been obtained by encoding exceptions, Java-style “synchronized” constructs (but with non-reentrant locks), and lists. The encoding uses higher-order functions, so it demonstrates an advantage of being able to deal with higher-order programs. For example, `exception` models the following OCaml-like program:

```
let rec read_and_update x =
  let n = read_int(x) (* may raise an Eof exception *) in
  lock g; c := c+n; (* may raise an Overflow exception *)
  unlock g; read_and_update x;;
let rec f file =
  let x = open_in file in
  try read_and_update(x) with
    Eof -> close x | Overflow -> unlock g;
spawn(f("foo"));spawn(f("bar"));join();print c
```

Table 1. The experimental results

Program	Order	# of functions	Checked pair	Reachability	Elapsed time [sec.]
example	2	4	(ℓ, ℓ)	no	0.11
example_wrong	2	4	(ℓ, ℓ)	yes	0.08
exception	3	7	(ℓ, ℓ)	no	0.25
exception_wrong	3	7	(ℓ, ℓ)	yes	0.12
synchronized	3	7	(ℓ_1, ℓ_1)	no	0.28
			(ℓ_1, ℓ_2)	yes	1.01
list	4	8	(ℓ_1, ℓ_1)	no	0.81
			(ℓ_1, ℓ_2)	no	1.75

Here, the goal of verification is to check that no race occurs on the shared variable c . The above program is encoded into the following order-3 program of the language in Section 2;

$$\begin{aligned}
 R\ h\ k &= \mathbf{if_}(h\ True) (\mathbf{acq}_g; (\mathbf{if_}(\mathbf{rel}_g; R\ h\ k) (h\ False))^\ell) \\
 F\ k &= R\ H\ k \quad H\ b = b\ \$ (\mathbf{rel}_g; \$) \quad P\ k = k^\ell \\
 S &= \mathbf{spawn}(F); (\mathbf{spawn}(F); (\mathbf{join}; (P\ \$))) \quad True\ x\ y = x \quad False\ x\ y = y
 \end{aligned}$$

Here, the function R corresponds to `read_and_update`; we have abstracted away x and instead added an exception handler h and a continuation parameter k in order to precisely model exception primitives. The program `exception_wrong` is a variation of `exception` obtained by omitting `acqg` and `relg` operations.

All the benchmark programs have been verified within a few seconds. Although the programs are very small, this is encouraging, considering the worst-case complexity of higher-order model checking. As discussed at the end of the previous section, the pairwise reachability can be decided in time linear *with respect to the size of HORS* under a certain assumption; thus, the results indicate that our method may scale for larger programs.

5 Related Work

As already mentioned in Section 1, the present work is based on the series of work on verification of pushdown systems with nested locking [7,14,4], and extends it to deal with higher-order programs. The target language of our verification is more expressive than dynamic pushdown networks and corresponds to an extension of collapsible pushdown systems [6] (which are higher-order pushdown systems extended with collapse operations) with concurrency primitives. Gawlitza et al. [4] used a clever encoding of a configuration of a dynamic pushdown network, so that the (forward) reachable set of configurations can be represented as a regular tree language and the reachability problem can be reduced to the inclusion between regular tree languages. One of our insights was that thanks to the decidability of higher-order model checking, actually we need not represent the reachable set as a regular language; a higher-order tree language (generated

by a HORS) suffices. This has enabled not only the higher-order extension, but also a conceptual simplification of the verification method in our opinion.

Higher-order model checking has recently been applied to program verification, but most of them have been for sequential programs [9,13,16]. Kobayashi and Igarashi [10] have shown that the reachability problem for higher-order concurrent programs with a bounded number of context switches can be reduced to higher-order model checking. (They have also shown a reduction from verification of higher-order concurrent programs to an extension of higher-order model checking, but the latter is undecidable.) Hague [5] has shown the decidability of reachability of ordered, phase-bounded and scope-bounded concurrent collapsible pushdown systems. These methods underapproximate the reachable set of ordinary higher-order concurrent programs (without the “bound” conditions). To our knowledge, there is no realistic implementation of those methods. There are also overapproximation approaches to static analysis or verification of higher-order concurrent programs [2,3,8].

6 Conclusion

We have shown the decidability of pairwise reachability of higher-order concurrent programs with recursion, dynamic process creation, joins, and nested locking. To our knowledge, this is the first realistic application of higher-order model checking to verification of concurrent programs. Despite the extremely high worst-case complexity of higher-order model checking, preliminary experiments show that our approach is feasible at least for small programs.

Acknowledgement. We would like to thank Markus Müller-Olm for the discussion on the subject and information about dynamic pushdown networks, and anonymous referees for useful comments. We would also like to thank Taku Terao for providing his higher-order model checker for the experiments. This work was supported by JSPS Kakenhi 23220001.

References

1. Broadbent, C.H., Kobayashi, N.: Saturation-based model checking of higher-order recursion schemes. In: Rocca, S.R.D. (ed.) CSL. LIPIcs, vol. 23, pp. 129–148. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2013)
2. D’Osualdo, E., Kochems, J., Ong, C.-H.L.: Automatic verification of erlang-style concurrency. In: Logozzo, F., Fähndrich, M. (eds.) Static Analysis. LNCS, vol. 7935, pp. 454–476. Springer, Heidelberg (2013)
3. Feret, J.: Abstract interpretation of mobile systems. *Journal of Logic and Algebraic Programming* 63(1) (2005)
4. Gawlitza, T.M., Lammich, P., Müller-Olm, M., Seidl, H., Wenner, A.: Join-lock-sensitive forward reachability analysis for concurrent programs with dynamic process creation. In: Jhala, R., Schmidt, D. (eds.) VMCAI 2011. LNCS, vol. 6538, pp. 199–213. Springer, Heidelberg (2011)

5. Hague, M.: Saturation of concurrent collapsible pushdown systems. In: Proceedings of FSTTCS 2013. LIPIcs, vol. 24, pp. 313–325 (2013)
6. Hague, M., Murawski, A., Ong, C.-H.L., Serre, O.: Collapsible pushdown automata and recursion schemes. In: Proceedings of 23rd Annual IEEE Symposium on Logic in Computer Science, pp. 452–461. IEEE Computer Society (2008)
7. Kahlon, V., Ivančić, F., Gupta, A.: Reasoning about threads communicating via locks. In: Etesami, K., Rajamani, S.K. (eds.) CAV 2005. LNCS, vol. 3576, pp. 505–518. Springer, Heidelberg (2005)
8. Kobayashi, N.: Type systems for concurrent programs. In: Aichernig, B.K. (ed.) Formal Methods at the Crossroads. From Panacea to Foundational Support. LNCS, vol. 2757, pp. 439–453. Springer, Heidelberg (2003)
9. Kobayashi, N.: Model checking higher-order programs. *J. ACM* 60(3), 20 (2013)
10. Kobayashi, N., Igarashi, A.: Model-checking higher-order programs with recursive types. In: Felleisen, M., Gardner, P. (eds.) Programming Languages and Systems. LNCS, vol. 7792, pp. 431–450. Springer, Heidelberg (2013)
11. Kobayashi, N., Ong, C.-H.L.: A type system equivalent to the modal mu-calculus model checking of higher-order recursion schemes. In: Proceedings of LICS 2009, pp. 179–188 (2009)
12. Kobayashi, N., Ong, C.-H.L.: Complexity of model checking recursion schemes for fragments of the modal mu-calculus. *Logical Methods in Computer Science* 7(4) (2011)
13. Kobayashi, N., Sato, R., Unno, H.: Predicate abstraction and CEGAR for higher-order model checking. In: Proceedings of PLDI 2011, pp. 222–233 (2011)
14. Lammich, P., Müller-Olm, M., Wenner, A.: Predecessor sets of dynamic pushdown networks with tree-regular constraints. In: Bouajjani, A., Maler, O. (eds.) CAV 2009. LNCS, vol. 5643, pp. 525–539. Springer, Heidelberg (2009)
15. Ong, C.-H.L.: On model-checking trees generated by higher-order recursion schemes. In: LICS, pp. 81–90 (2006)
16. Ong, C.-H.L., Ramsay, S.: Verifying higher-order programs with pattern-matching algebraic data types. In: Proceedings of POPL 2011, pp. 587–598 (2011)
17. Plotkin, G.D.: Call-by-name, call-by-value and the lambda-calculus. *Theor. Comput. Sci.* 1(2), 125–159 (1975)
18. Qadeer, S., Rehof, J.: Context-bounded model checking of concurrent software. In: Halbwachs, N., Zuck, L.D. (eds.) TACAS 2005. LNCS, vol. 3440, pp. 93–107. Springer, Heidelberg (2005)
19. Ramalingam, G.: Context-sensitive synchronization-sensitive analysis is undecidable. *ACM Trans. Program. Lang. Syst.* 22(2), 416–430 (2000)
20. Sato, R., Unno, H., Kobayashi, N.: Towards a scalable software model checker for higher-order programs. In: Proceedings of PEPM 2013, pp. 53–62 (2013)