

Context-Bounded Analysis of Multithreaded Programs with Dynamic Linked Structures

Ahmed Bouajjani¹, Séverine Fratani¹, and Shaz Qadeer²

¹ LIAFA, CNRS & University of Paris 7, {abou, fratani}@liafa.jussieu.fr

² Microsoft Research, Redmond, qadeer@microsoft.com

Abstract. Bounded context switch reachability analysis is a useful and efficient approach for detecting bugs in multithreaded programs. In this paper, we address the application of this approach to the analysis of multithreaded programs with procedure calls and dynamic linked structures. We define a program semantics based on concurrent pushdown systems with *visible heaps* as stack symbols. A visible heap is the part of the heap reachable from global and local variables. We use pushdown analysis techniques to define an algorithm that explores the entire configuration space reachable under given bounds on the number of context switches and the size of visible heaps.

1 Introduction

Automated analysis of software systems is a challenging problem. The behavior of these systems is usually complex and hard to predict due to aspects such as concurrency and memory management. Reasoning about these behaviors requires considering potentially infinite sets of configurations which makes verification problems undecidable in general. Therefore, approaches based on approximate analysis are needed. While over-approximations are useful for proving properties, under-approximations are useful for finding bugs. In this paper, we propose algorithmic (automata-based) techniques for under-approximate analysis of multithreaded programs that manipulate dynamic linked structures.

A simple way to get an under-approximate analysis is to bound the depth of the explored state space and use finite-state model checking techniques. This approach is interesting only if bugs appear after a small number of computation steps (e.g., [6]), which is unlikely to be the case for multithreaded programs. In such a program, processes interact through shared memory and are executed in alternation according to a schedule. Concurrency bugs appear only after a number of *context switches*, i.e. points in the execution where the active process is stopped and another one is resumed. Between context switches, a process may execute an unbounded number of computation steps. Therefore, a natural approach for under-approximate analysis of multithreaded program is to perform a (precise) analysis for a bounded number of context switches [9]. It has been demonstrated that this is indeed efficient for detecting bugs in multithread programs since they appear in many cases after a small number of context switches [7].

However, bounding the number of context-switches does not allow the use of finite-state model checkers since each process may be infinite-state, and unbounded computations are possible between context switches. For example, in concurrent programs with procedure calls, each process may have an unbounded stack. In [9, 2], automata-based

techniques for symbolic analysis of pushdown systems have been used in order to define bounded context switch analysis of programs with finite data domains. In this paper, we consider the more general case of concurrent programs (with procedure calls) which can create shared objects and manipulate references to them.

In the spirit of under-approximate analysis, we could bound the heap size and reduce this verification problem to the case with finite data domains. This simple approach is unlikely to be effective at getting good state-space coverage for real programs, in which many heap-allocated objects are used either as local variables or for passing parameters to a called procedure. For such programs, it is better to bound only the *visible heap*, i.e., the part of the heap that is reachable from the global variables and the local variables of the running procedure in the active process. Indeed, a program's global heap might be unbounded in spite of its visible heap being bounded due to unbounded nesting of procedure calls.

It is nontrivial to obtain an algorithm for reachability analysis of multithreaded programs based on bounding the number of context switches and the size of the visible heap. In [10, 11, 8], the idea of visible heaps was used for interprocedural analysis of sequential programs; these techniques are based on procedure summarization and cannot be used for multithreaded programs since the (infinite) sets of stack configurations reached by processes must be stored at each context switch. Moreover, [10, 11] consider abstract semantics, whereas we must consider an exact - sound and complete - semantics. Nevertheless, the idea of visible heaps can be used to define a program semantics where the heap is manipulated implicitly and not as a shared global structure: each procedure executed by some process manipulates a local heap structure which is a *copy* of its visible heap. Such an approach leads to a program model based on stack machines where locally visible heaps constitute the (potentially infinite) stack alphabet. To simulate correctly the "real" heap, our algorithm synchronizes the local views of procedures and processes at each procedure call or return and at each context switch. We prove that our new semantics is correct and use it to define an automata-based symbolic reachability analysis algorithm for bounded context switches and bounded visible heap. For lack of space, the proofs of our theorems are omitted here. They can be found in [3].

The contribution of our paper, although theoretical, has important practical applications. Since our algorithm constructs the set of reachable stack configurations, it allows the verification of reachability queries that require stack inspection. This expressiveness is important for specifying various resource-usage scenarios, e.g., user-space buffers are accessed by an operating system kernel only within the scope of an appropriate exception handler, or certain operations on security-critical objects are invoked only if a certain privileged procedure is present on the call stack. Our algorithm allows the verification of such expressive queries for multithreaded programs upto the context-switch and visible heap bounds and for single-threaded programs upto the visible heap bound. Of course, our algorithm can be used iteratively by systematically increasing the bounds in each iteration.

2 Multithreaded programs with dynamic memory

We consider programs with multiple threads and procedure calls, which can dynamically create objects and manipulate pointers on these objects. A sequential program is

given as a collection of control flow graphs, one graph for each of its procedures, defined by a set of control nodes N , and a set of transitions between these nodes labeled with actions and tests on the memory heap. We assume that the control flow graph of each procedure p has a unique initial control location n_{init}^p and a unique termination location n_{end}^p reachable after executing a return action. A multithread program consists of a parallel composition of a fixed number of sequential programs sharing the heap.

Let \mathbb{A} be a countable (infinite) domain of memory addresses (pointers), and assume that \mathbb{A} contains a special element \perp representing the null address and a special element \top representing an undefined address. Then, consider a class of objects where each object contains n successor fields s_1, \dots, s_n ranging over the pointer domain \mathbb{A} . Let $S = \{s_1, \dots, s_n\}$. (We omit aspects related to data manipulation. Data values over finite data domains can of course be encoded in the successor fields.)

We assume that a program has a set of *global* pointer variables G ranging over \mathbb{A} . These variables are shared by all parallel threads, and all procedures. Moreover, we consider that each procedure has a set of *local* pointer variables (also ranging over \mathbb{A}). We assume w.l.o.g. that all the procedures have the same set of local variables L . Given a (global/local) pointer variable v and a successor field $s \in S$, we denote by $v.s$ the pointer stored in the field s of the object pointed by v . This notation can be extended to sequences of successor fields $\sigma \in S^*$ in the obvious manner. When σ is the empty sequence, we consider that $v.\sigma$ is identical to v .

Programs can perform the following operations on heaps: $v.\sigma := v'.\sigma'$ (*pointer assignment*), $v.\sigma := null$ (*pointer annihilation*), and $v.\sigma := new$ (*object creation*), where v, v' are pointer variables, and $\sigma \in S \cup \{\epsilon\}, \sigma' \in S^*$. They can also perform the following tests: $v.\sigma \# v'.\sigma'$ with $\# \in \{=, \neq\}$ (*equality/disequality test*). In addition, they can perform the following actions: $call(p, v_1.\sigma_1, \dots, v_m.\sigma_m)$ (*procedure call with parameters*), and $return$ (*termination of a procedure call*), where p is a procedure name, v_1, \dots, v_m are pointer variables, and $\sigma_1, \dots, \sigma_m \in S^*$. The effect of executing a statement $call(p, v_1.\sigma_1, \dots, v_m.\sigma_m)$ is to call the procedure p after initialization of its local variables ℓ_1, \dots, ℓ_m with the pointer values $v_1.\sigma_1, \dots, v_m.\sigma_m$ respectively where the variables v_1, \dots, v_m are either global variables or local variables of the calling procedure.

3 Program semantics

Heaps as labeled graphs: A *global heap* is a finite directed graph where vertices correspond to memory addresses, edges are labeled by elements of S , and each element of G appears as a label of some vertex. Formally, a global heap is a tuple $GH = (A, \Delta, \Gamma)$ where (1) A is a finite subset of \mathbb{A} containing \perp and \top , (2) $\Delta : S \rightarrow (A \rightarrow A)$ associates with each $s \in S$ a successor mapping, and (3) $\Gamma : G \rightarrow A$ associates with a global variable g an object address. We assume that $\forall s \in S. \Delta(s)(\perp) = \top$ and $\Delta(s)(\top) = \top$, and $\forall a \in A, \Delta(s)(a)$ is defined.

Given $\sigma = s_{i_1} \dots s_{i_m} \in S^*$, we denote by $\Delta(\sigma)$ the mapping $\Delta(s_{i_m}) \circ \dots \circ \Delta(s_{i_1})$. When σ is the empty sequence, we consider that $\Delta(\sigma)$ is the identity mapping over A . Then, we define the reachability mapping to be $Reach_\Delta = \bigcup_{\sigma \in S^*} \Delta(\sigma)$.

Then, we consider that a *heap* is a global heap augmented by a mapping associating with each local variable an address. Formally, a heap is a pair $H = (GH, \Lambda)$, where $GH = (A, \Delta, \Gamma)$ is a global heap and $\Lambda : L \rightarrow A$.

Programs as concurrent heap pushdown systems: We associate with multithreaded programs concurrent stack machines which act on a (global) shared heap. The behavior of such concurrent stack machines is as follows: at each moment, there is one process which is running and has access to the heap through the global variables (shared by all processes), and its own local variables. Notice that due to nested procedure calls, the heap may be accessible from the local variables of the currently running procedure, and also from the local environments (of the pending procedure executions) stored in the stack of the process. While this process is running, all the other parallel processes are in an idle mode. When a context switch occurs, the active process is interrupted and some other one is resumed, becoming the new active process.

Let us sketch here the construction of the models associated with programs (details can be found in [3]). Let Σ be the (infinite) set of all pairs $\langle n, \Lambda \rangle$, where $n \in N$ and $\Lambda \in [L \rightarrow \mathbb{A}]$. We associate with each sequential program (given by a control flow graph) a *heap pushdown system* (H-PDS) which is a stack machine whose control states are global heaps, and whose stack alphabet is Σ : The semantics of basic operations and tests is defined by a transition relation \xrightarrow{op} between heaps, and procedure calls and returns are modeled as usual by push and pop operations on the stack. Then, we associate with a multithreaded program with m parallel threads an m -dim *concurrent heap pushdown system* (CH-PDS). Control states of such a model are pairs (i, GH) , where $i \in [1, m]$ is the index of the active process, and GH is a global heap. A configuration of a CH-PDS is a tuple $((i, GH), [u_1, \dots, u_m])$ where $u_1, \dots, u_m \in \Sigma^*$ are the local configurations of each of the threads, i.e., the contents of their stacks. A transition relation \Rightarrow is defined between configurations: computations steps are either local to a process (the process i in the global state (i, GH)), or correspond to a context switch (i.e., substitution of i by some $j \neq i$ in the global state (i, GH)).

Bounded heap depth programs: Let $k \geq 1$. Let $c = ((i, GH), [u_1, \dots, u_m])$ be a configuration, where $GH = (A, \Delta, \Gamma)$. Then, c is k -bounded iff for every $\langle n, \Lambda \rangle$ appearing in any u_i , for any $i \in \{1, \dots, m\}$, we have $|\text{Reach}_\Delta(\Gamma(G) \cup \Lambda(L))| \leq k$. A program is k -bounded for a set of initial configurations C iff every \Rightarrow -reachable configuration from C is k -bounded. A program has *bounded heap depth* if it is k -bounded for some $k \geq 1$.

Notice that a heap with a bounded depth may have an unbounded size. Indeed, due to unbounded nesting of procedure calls, it is possible to have an unbounded stack of environments each of them pointing to a different but bounded part of the heap.

4 Program semantics based on locally visible heaps

We define in this section a program semantics in terms of concurrent pushdown systems. The difference with the CH-PDS based semantics of the previous section is that the heap is not manipulated explicitly as a global shared structure, but rather implicitly: each procedure executed by some process manipulates a local heap structure which is a *copy* of the reachable part of the heap from its local and global variables. Therefore, we can associate with programs stack machines where locally visible heaps constitute the stack alphabet. However, to simulate correctly operations on the “real” heap, procedures and processes must exchange (pass/update) their local informations about the heap at each procedure call/return and context switch. These manipulations of the local heaps are quite delicate. Let us give an informal description of the main ideas behind them.

Consider first the case of a sequential program. At each procedure call, the caller passes to the callee a copy of a part of its local heap. At the return of the call, the callee gives back its new local heap, the last one before termination, to the caller which updates accordingly its own local heap (which corresponds to its view of the heap just before the call). To perform correctly these caller/callee communications, a relation between vertices in different copies of a same piece of the “real” heap must be maintained, relating vertices which are copies of each other, representing the same address in the heap. Moreover, since pointer manipulation may disconnect some vertices from the visible part of a procedure during its execution, the deletion of these vertices from the local structure can be problematic since these vertices may still be relevant for further executions of procedures in the call stack. To handle this problem, we introduce a notion of *cut points*: when a heap is passed from the caller to the callee, cut points represent the first vertices in the visible part by the callee which are reachable from the caller or from other procedures in the call stack. Then, during the execution of the procedure, a vertex can be removed from the locally visible heap only if it becomes unreachable from the local variables, the global variables, and the cut points.

In the case of parallel programs, an additional but similar mechanism has to be introduced for passing/updating locally visible heaps at context switches. Intuitively, parallel processes synchronize their views of the heap at each context switch. The active process passes its local heap to all of the other processes which can update accordingly their local heaps (corresponding to the configuration at the previous context switch), taking into account the modifications on the heap performed by the active process while they were idle. For that, the active process maintains in its control state, which is a global state for all its procedures, a relation between its current local heap and its initial local heap corresponding to the configuration at the last context switch, allowing to determine for each vertex in the (old) local heaps whether it has a copy in the (new) local heap returned by the active process. Also, to deal with vertex deletion, it is necessary to extend the use of cut points by considering the reachable vertices from the stacks of each process. In fact, each process needs to distinguish his own cut points from the ones of the other processes.

We prove that the new semantics is correct (sound and complete) w.r.t. the original semantics (in section 3) in the sense that they define bisimilar transition systems. An important property allowing to prove this fact is that isomorphism between locally visible heaps preserves bisimilarity (i.e., substitution of isomorphic local heaps does not modify behaviors). This holds because the performed operations and tests on the heaps do not refer to the precise values of the addresses. (Notice that there is an infinite number of isomorphic heaps of a same size since the address domain is infinite). Then, we consider a normalization operation for visible heaps which associates to each of them an equivalent heap modulo isomorphism, and we prove that the semantics based on normalized visible heaps is also correct w.r.t. the original semantics. Our normalization operation matches all the isomorphic visible heaps to a finite number of representatives. (We do not have a unique representative due to the presence of cut points.) Therefore, for sequential programs, our pushdown model construction terminates for bounded heap depth programs. For concurrent programs, the construction terminates for a finite number of a context switches, given a bound on the heap depth.

4.1 Locally visible heaps

A *cut heap* is a tuple $CH = (A, \Delta, \Gamma, \Lambda, \vec{C})$ where $(A, \Delta, \Gamma, \Lambda)$ is a heap and $\vec{C} \in (2^A)^m$ is a vector of sets of cut points. For each $i \in \{1, \dots, m\}$, $\vec{C}(i)$ is the set of cut points reachable from the local environments (of the pending procedure executions) stored in the stack of the process i . Then, a *visible heap* is a cut heap such that $A = \text{Reach}_\Delta(\Gamma(G) \cup \Lambda(L) \cup \bigcup_{i=1}^m \vec{C}(i) \cup \{\perp\})$, i.e., it is a cut heap without garbage. We define an operation of *garbage elimination* allowing to obtain a visible heap from a cut heap: given a cut heap $CH = (A, \Delta, \Gamma, \Lambda, \vec{C})$, the visible heap $\text{Clean}(CH)$ is given by $(A', \Delta', \Gamma, \Lambda, \vec{C})$ where $A' = A \cap \text{Reach}_\Delta(\Gamma(G) \cup \Lambda(L) \cup \bigcup_{i=1}^m \vec{C}(i) \cup \{\perp\})$ and $\forall s \in S. \Delta'(s) = \Delta(s) \cap [A' \rightarrow A']$, i.e., the restriction of $\Delta(s)$ to A' .

We define an equivalence relation \simeq between visible heaps which is essentially graph isomorphism modulo renaming of vertices. Let $VH_1 = (A_1, \Delta_1, \Gamma_1, \Lambda_1, \vec{C}_1)$ and $VH_2 = (A_2, \Delta_2, \Gamma_2, \Lambda_2, \vec{C}_2)$ be two visible heaps and let $\beta : A_1 \rightarrow A_2$ be a bijection s.t. $\beta(\top) = \top$ and $\beta(\perp) = \perp$. Then, $VH_1 \simeq_\beta VH_2$ iff (i) $\forall s \in S. \forall a \in A_1. \beta(\Delta_1(s)(a)) = \Delta_2(s)(\beta(a))$, (ii) $\forall v \in G. \Gamma_2(v) = \beta(\Gamma_1(v))$, and $\forall v \in L. \Lambda_2(v) = \beta(\Lambda_1(v))$, and (iii) $\forall i \in [1, m], \vec{C}_2(i) = \{\beta(c) : c \in \vec{C}_1(i)\}$. Then, $VH_1 \simeq VH_2$ if there is a β s.t. $VH_1 \simeq_\beta VH_2$.

4.2 Sequential programs as pushdown systems

Let us define an *environment* to be a tuple $e = \langle n, VH, \pi \rangle$ where $n \in N$, VH is a visible heap, and $\pi \subseteq \mathbb{A} \times \mathbb{A}$ is an injective function (π relates vertices in the heap of the caller procedure with vertices in the current visible heap VH). Let \mathbb{V} be the set of all environments. Then, we associate with each sequential process a *visible heap pushdown system* (VH-PDS) which is a stack machine whose stack alphabet is \mathbb{V} and whose control states are injective functions from \mathbb{A} to \mathbb{A} . These functions are used to maintain a link between the current visible heap of the running sequential process and its initial visible heap since its last activation (at the initial configuration of the system or at the last context switch). This information is needed at the next context switch for updating the visible heaps of the idle processes. (Information in control states can be omitted in the case of a purely sequential program.) Now, let i be the index of a sequential process. We define the (infinite) set of transition rules R_i of the stack machine associated with process i by means of the three inference rules given hereafter.

Basic operations and test: We extend the relation \xrightarrow{op} defined on heaps (see section 3) to a relation on visible heaps (which also informs about the correspondence between vertices in the original and final heaps). Let $VH_1 = (H_1, \vec{C}_1)$, and $VH_2 = (H_2, \vec{C}_2)$ be two visible heaps. For every function $\pi : \mathbb{A} \times \mathbb{A}$, we have $VH_1 \xrightarrow{op, \pi} VH_2$ iff there exists a heap H'_2 s.t. $H_1 \xrightarrow{op} H'_2$ and $\text{Clean}(H'_2, \vec{C}_1) \simeq_\pi VH_2$. Then:

$$\frac{n_1 \xrightarrow{op} n_2 \quad VH_1 \xrightarrow{op, \pi} VH_2}{(\Pi, \langle n_1, VH_1, \pi_1 \rangle) \hookrightarrow (\Pi \circ \pi, \langle n_2, VH_2, \pi_1 \circ \pi \rangle) \in R_i} \text{HeapOp}$$

Procedure calls: Let $i \in [1, m]$ be the index of the active sequential process, let $VH_1 = (A_1, \Delta_1, \Gamma_1, \Lambda_1, \vec{C}_1)$ be the visible heap of the caller procedure, and let $v_1.\sigma_1, \dots, v_k.\sigma_k$ be the effective parameters which must be assigned respectively to the local variables ℓ_1, \dots, ℓ_k of the called procedure (considered as its formal parameters). The visible heap passed to the callee procedure (of the process i) is obtained as follows: (1) construct first the cut heap $CutPass_i(VH_1)$ which is a copy of VH_1 , where local variables of the callee procedure are assigned with their new values, and local variables of the caller procedure are memorized as cut points, (2) then determine the new vector of sets of cut points: cut points of processes $j \neq i$ stay unchanged, cut points of the process i are the first addresses reachable from cut points of $CutPass_i(VH_1)$ in the sub-heap corresponding to reachable addresses from the new local variables, the global variables, and the cut points of all processes, (3) finally remove the garbage of $CutPass_i(VH_1)$ according to the new set of cut points.

Let us define formally the operation $CutPass_i$. Let $VH = (A, \Delta, \Gamma, \Lambda, \vec{C})$ be a visible heap. Then, $CutPass_i(VH, [\ell_1, v_1.\sigma_1], \dots, [\ell_k, v_k.\sigma_k])$ is the cut heap $(A, \Delta, \Gamma, \Lambda', \vec{C}')$ where (1) $\forall j \in [1, k], \Lambda'(\ell_j) = \Delta(\sigma_j)((\Lambda \cup \Gamma)(v_j))$ and $\Lambda'(\ell) = \top$ for all other variables $\ell \in L$, (2) $\forall j \in [1, m], j \neq i$, we have $\vec{C}'(j) = \vec{C}(j)$, and (3) $\vec{C}'(i) = \vec{C}(i) \cup \Lambda(L)$.

We give now the formal definition of the visible heap passed by the caller to the callee. Consider $CH = (A, \Delta, \Gamma, \Lambda, \vec{C})$ to be any cut heap. Let $A' = Reach_\Delta(\Gamma(G) \cup \Lambda(L) \cup \bigcup_{j \neq i} \vec{C}(j))$ (the set A' is the visible part of the heap from local and global variables and from cut points in $\vec{C}(j)$, for $j \neq i$), and let Δ' be the restriction of Δ to vertices in A' . Then, let $Visible_i(CH) = Clean(A, \Delta, \Gamma, \Lambda, \vec{C}')$ where, for every $j \in [1, m]$ s.t. $j \neq i$, $\vec{C}'(j) = \vec{C}(j)$, and $\vec{C}'(i) = (A' \cap Reach_\Delta(\vec{C}(i))) \setminus \Delta'(S)(Reach_\Delta(\vec{C}(i)))$. The set $\vec{C}'(i)$ contains the first vertices in A' which are Δ -reachable from $\vec{C}(i)$. The cleaning operation removes all vertices which are Δ -reachable only from $\vec{C}(i)$, i.e., the set of vertices in $Visible_i(CH)$ is precisely A' .

Finally, given a visible heap VH , we define $Pass_i(VH, p, v_1.\sigma_1, \dots, v_k.\sigma_k)$ to be the set of pairs (VH_π, π) such that $Visible_i(CutPass_i(VH, [\ell_1, v_1.\sigma_1], \dots, [\ell_k, v_k.\sigma_k])) \simeq_\pi VH_\pi$. The injective relation π allows to relate addresses in the visible heap of the caller with addresses in the new visible heap. It is used to update the heap of the caller after termination of the callee procedure. Then, we consider the inference rule:

$$\frac{n_1 \xrightarrow{call(p, v_1.\sigma_1, \dots, v_k.\sigma_k)} n_2 \quad (VH_2, \pi_2) \in Pass_i(VH_1, p, v_1.\sigma_1, \dots, v_k.\sigma_k)}{(\Pi, \langle n_1, VH_1, \pi_1 \rangle) \hookrightarrow (\Pi \circ \pi_2, \langle n_{init}^p, VH_2, \pi_2 \rangle \langle n_2, VH_1, \pi_1 \rangle) \in R_i} \quad \text{Call}$$

Procedure returns: Given the current visible heap VH_2 of the terminating procedure and the visible heap VH_1 of the caller procedure stored in the stack, we define an operation updating VH_1 according to the effect of the procedure call on the heap.

Let $\pi_2 \subseteq A_1 \times A_2$ be an injective relation giving the correspondence between the vertices of VH_1 and VH_2 (i.e., if $(a_1, a_2) \in \pi$ and $(a'_1, a_2) \in \pi$, then $a_1 = a'_1$). We suppose that $A_1 \cap A_2 = \{\top, \perp\}$ (otherwise, instead of VH_2 and π_2 , consider a new visible heap VH'_2 and $\pi_2 \circ \pi$ for some π s.t. $VH_2 \simeq_\pi VH'_2$).

Then, let $B_1 = \{a \in A_1 : \nexists a' \in A_2. (a, a') \in \pi_2\}$, let $B_2 = \{a \in A_1 : \exists a' \in A_2. (a, a') \in \pi_2\}$, and let $B_3 = \{a' \in A_2 : \exists a \in A_1. (a, a') \in \pi_2\}$.

Intuitively, B_1 is the set of vertices in A_1 such that either they were not reachable in the initial visible heap passed to the called procedure, and therefore they should be restored in the heap of the caller procedure after the call return, or they became invisible during the call execution due to garbage deletion and therefore they should not appear in the heap after the call return since they were not reachable from cut points. (We explain below how to get rid of the vertices in this second category). Vertices in B_2 are those which were present before the call, and which are still present after termination of the call. Finally, B_3 is the set of the created vertices during the call.

It can be easily seen that these three sets are disjoint. Moreover, we have $B_3 \cup \pi_2(B_2) = A_2$. Let us consider the bijection $\beta : A_2 \rightarrow B_2 \cup B_3$ defined in the obvious way (for every $a \in A_2$, if $a \in B_3$ then $\beta(a) = a$, otherwise $\beta(a) = \pi_2^{-1}(a)$).

Let $(A'_1, \Delta'_1, \Gamma'_1, \Lambda'_1, \vec{C}'_1)$ be the cut heap such that (1) $A'_1 = B_1 \cup B_2 \cup B_3$, (2) $\forall g \in G. \Gamma'_1(g) = \beta(\Gamma_2(g))$, (3) $\forall \ell \in L. \Lambda'_1(\ell) = \Lambda_1(\ell)$, (4) $\vec{C}'_1(i) = \vec{C}_1(i)$ and $\forall j \neq i. \vec{C}'_1(j) = \beta(\vec{C}_2(j))$, and (5) $\forall s \in S$, (i) $\forall a \in B_1. \Delta'_1(s)(a) = \Delta_1(s)(a)$, (ii) $\forall a \in B_2. \Delta'_1(s)(a) = \beta(\Delta_2(s)(\pi_2(a)))$, and (iii) $\forall a \in B_3. \Delta'_1(s)(a) = \beta(\Delta_2(s)(a))$.

Then, we define $\text{Update-seq}_i(VH_1, VH_2, \pi_2)$ to be the set of all pairs (VH_π, π) such that $\text{Clean}(A'_1, \Delta'_1, \Gamma'_1, \Lambda'_1, \vec{C}'_1) \simeq_\pi VH_\pi$.

Notice that (1) the cleaning operation removes the vertices of A_1 which were garbage collected during the procedure call, and (2) for every VH'_2 , and every β s.t. $VH_2 \simeq_\beta VH'_2$, $\text{Update-seq}_i(VH_1, VH_2, \pi_2) = \text{Update-seq}_i(VH_1, VH'_2, \pi_2 \circ \beta)$.

Then, we consider the inference rule:

$$\frac{n \xrightarrow{\text{return}} n_{\text{end}} \quad (VH'_1, \pi') \in \text{Update-seq}_i(VH_1, VH_2, \pi_2)}{(\Pi, \langle n, VH_2, \pi_2 \rangle \langle n', VH_1, \pi_1 \rangle) \leftrightarrow (\Pi \circ \pi', \langle n', VH'_1, \pi_1 \circ \pi' \rangle) \in R_i} \quad \text{Return}$$

4.3 Multithreaded programs as concurrent pushdown systems

We associate with a multithreaded program with m parallel threads an m -dim *concurrent visible heap pushdown system* (CVH-PDS). The stack alphabet is \mathbb{V} , and the (infinite) set of control states is $\mathbb{S} = \{(i, \vec{\Pi}) : i \in [1, m], \vec{\Pi} = (\vec{\Pi}(1), \dots, \vec{\Pi}(m)), \vec{\Pi}(j) : \mathbb{A} \rightarrow \mathbb{A}\}$. We define hereafter the set of transition rules \mathcal{R} of the model.

Local transitions: Transitions of each sequential process are obviously transitions of the whole system. For every $i \in [1, m]$, let R_i be the set of transition rules associated with the process of index i (defined in the previous subsection). Then, we have:

$$\frac{(\Pi_i, w_i) \leftrightarrow (\Pi'_i, w'_i) \in R_i \quad \vec{\Pi}(i) = \Pi_i \quad \vec{\Pi}'(i) = \Pi'_i \quad \forall j \neq i. w_j \in \mathbb{V} \text{ and } \vec{\Pi}(j) = \vec{\Pi}'(j)}{((i, \vec{\Pi}), [w_1, \dots, w_i, \dots, w_m]) \leftrightarrow ((i, \vec{\Pi}'), [w_1, \dots, w'_i, \dots, w_m]) \in \mathcal{R}} \quad \text{Local}$$

Context switches: Finally, consider the case of a context switch. Assume that process i was the last active process and let $VH_i = (A_i, \Delta_i, \Gamma_i, \Lambda_i, \vec{C}_i)$ be its last visible heap. At the context switch, process i communicates to all the other ones (which were idle)

informations about the new heap configuration. For that, the part which is *shared* by all processes is extracted from VH_i and passed to them. Then, each process $j \neq i$ updates accordingly its old visible heap (which corresponds to the heap configuration at the previous context switch) before the next active process starts its computation.

Formally, let us define $Shared_i(VH_i) = Visible_i(CutPass_i(VH_i, \forall \ell \in L. [\ell, \top]))$, and assume that $Shared_i(VH_i) = (A'_i, \Delta'_i, \Gamma'_i, \Lambda'_i, \vec{C}'_i)$. Notice that we remove from A_i only vertices that are reachable from local variables in L or from cut points in $\vec{C}'_i(i)$, but which are not reachable from global variables, nor from cut points of other threads. These vertices are not visible from other threads and therefore they do not belong to the shared part of the heap. The cut points in $\vec{C}'_i(i)$ allow to know which vertices in the shared part are reachable from local variables in the stack of the current thread. The removed vertices will be added to the heap when the current process will resume later. This is done by an updating operation described below.

Let $j \in [1, m]$, let $VH_j = (A_j, \Delta_j, \Gamma_j, \Lambda_j, \vec{C}_j)$ be the visible heap of the process j , and suppose $\Pi \subseteq A_j \times A'_i$ is an injective relation connecting vertices in VH_j with vertices in $Shared_i(VH_i)$. We suppose that $A_i \cap A_j = \{\top, \perp\}$ (otherwise, instead of VH_j and Π , consider VH'_j and $\Pi \circ \pi$ for some π s.t. $VH_j \simeq_\pi VH'_j$).

Then, let $B_1 = \{a \in A_j : \exists a' \in A'_i. (a, a') \in \Pi\}$, let $B_2 = \{a \in A_j : \exists a' \in A'_i. (a, a') \in \Pi\}$, and let $B_3 = \{a' \in A'_i : \exists a \in A_j. (a, a') \in \Pi\}$. These sets are disjoint and we have $B_3 \cup \Pi(B_2) = A'_i$. Then, consider the bijection $\beta_j : A'_i \rightarrow B_2 \cup B_3$ defined in the obvious way (for every $a \in A'_i$, if $a \in B_3$ then $\beta_j(a) = a$, otherwise $\beta_j(a) = \Pi^{-1}(a)$).

We define the operation $Update\text{-}par_j$, for $j \in [1, m]$, which updates the local heap of process j using the shared heap passed by process i . For $j \neq i$, let $(A'_j, \Delta'_j, \Gamma'_j, \Lambda'_j, \vec{C}'_j)$ be the cut heap s.t. (1) $A'_j = B_1 \cup B_2 \cup B_3$, (2) $\forall g \in G. \Gamma'_j(g) = \beta_j(\Gamma'_i(g))$, (3) $\forall \ell \in L. \Lambda'_j(\ell) = \Lambda_j(\ell)$, (4) $\vec{C}'_j(j) = \vec{C}_j(j)$, and $\forall k \neq j. \vec{C}'_j(k) = \beta_j(\vec{C}'_i(k))$, and (5) $\forall s \in S$, (i) $\forall a \in B_1. \Delta'_j(s)(a) = \Delta_j(s)(a)$, (ii) $\forall a \in B_2. \Delta'_j(s)(a) = \beta_j(\Delta'_i(s)(\Pi(a)))$, and (iii) $\forall a \in B_3. \Delta'_j(s)(a) = \beta_j(\Delta'_i(s)(a))$. Then, $Update\text{-}par_j(VH_j, VH_i, \Pi)$ is defined to be the set of all (VH'_j, π, Π_j) such that (1) $Clean(A'_j, \Delta'_j, \Gamma'_j, \Lambda'_j, \vec{C}'_j) \simeq_\pi VH'_j$, and (2) $\Pi_j = \pi^{-1} \circ \beta_j^{-1}$. Moreover, for every mapping $\Pi_i : \mathbb{A} \rightarrow \mathbb{A}$, we define $Update\text{-}par_i(VH_i, VH_i, \Pi_i)$ to be the set of all (VH'_i, π, π^{-1}) such that $VH_i \simeq_\pi VH'_i$. Then, we consider the inference rule:

$$\frac{\begin{array}{l} i, k \in \{1, \dots, m\} \quad i \neq k \quad \forall j \in \{1, \dots, m\}. w_j = \langle n_j, VH_j, \pi_j \rangle \\ \forall j \in \{1, \dots, m\}. (VH'_j, \pi'_j, \Pi_j) \in Update\text{-}par_j(VH_j, VH_i, \vec{\Pi}(j) \circ \vec{\Pi}(i)) \\ \forall j \in \{1, \dots, m\}. w'_j = \langle n_j, VH'_j, \pi_j \circ \pi'_j \rangle \end{array}}{((i, \vec{\Pi}), [w_1, \dots, w_m]) \hookrightarrow ((k, (\Pi_1, \dots, \Pi_m) \circ \Pi_k^{-1}), [w'_1, \dots, w'_m]) \in \mathcal{R}} \quad \text{Switch}$$

with the notational convention $(\Pi_1, \dots, \Pi_m) \circ \Pi = (\Pi_1 \circ \Pi, \dots, \Pi_m \circ \Pi)$.

4.4 Associated transition system and correctness of the semantics

A *configuration* of a CVH-PDS is a tuple $((i, \vec{\Pi}), [v_1, \dots, v_m]) \in \mathbb{S} \times [\mathbb{V}^*]^m$ where $i \in [1, m]$ is the index of the active process, and $v_1, \dots, v_m \in \mathbb{V}^*$ are the local configurations of all the processes. Let \mathbb{C} be the set of all configurations.

Given $r = ((i_1, \vec{\Pi}_1), [u_1, \dots, u_m]) \hookrightarrow ((i_2, \vec{\Pi}_2), [u'_1, \dots, u'_m]) \in \mathcal{R}$, let $\mapsto_r \subseteq \mathbb{C} \times \mathbb{C}$ be the relation s.t. $b \mapsto_r b'$ iff $b = ((i_1, \vec{\Pi}_1), [v_1, \dots, v_m])$, $b' = ((i_2, \vec{\Pi}_2), [v'_1, \dots, v'_m])$, and $\forall k \in [1, m], \exists w_k \in \mathbb{V}^*$ s.t. $v_k = u_k w_k$ and $v'_k = u'_k w_k$.

Let \mapsto_{loc} (resp. \mapsto_{sw}) be the union of all relations \mapsto_r where r is a Local rule (resp. Switch rule), and let \mapsto be the union of all the relations \mapsto_r , for $r \in \mathcal{R}$. Then, we consider the relation $\sim = \mapsto_{\text{loc}}^* \circ \mapsto_{\text{sw}}$. For each $K \geq 1$, the relation \sim^K (K th power of \sim) corresponds to \mapsto -reachability with $K - 1$ context switches (or K consecutive contexts).

We extend the equivalence relation \simeq defined on visible heaps to environments in \mathbb{V} : we consider that $\langle n_1, VH_1, \pi_1 \rangle \simeq_\beta \langle n_2, VH_2, \pi_2 \rangle$ iff (1) $n_1 = n_2$, and (2) β is a bijection from A_1 to A_2 such that $VH_1 \simeq_\beta VH_2$ and $\pi_2 = \pi_1 \circ \beta$. Given two environments e_1, e_2 , we write $e_1 \simeq e_2$ iff there exists β such that $e_1 \simeq_\beta e_2$. We extend this equivalence relation to sequences of environments in the obvious manner ($e_1 \dots e_j \simeq e'_1 \dots e'_k$ iff $j = k$ and for every $i \in [1, j]$, $e_i \simeq e'_i$).

Finally, we extend \simeq to configurations: let $b = ((i, \vec{\Pi}), [e_1 \alpha_1, \dots, e_m \alpha_m])$ and $b' = ((j, \vec{\Pi}'), [e'_1 \alpha'_1, \dots, e'_m \alpha'_m])$ be two configurations. Then, $b \simeq b'$ iff (1) $i = j$, (2) $\forall k \in [1, m], \alpha_k \simeq \alpha'_k$, and (3) $\exists \pi_k : \mathbb{A} \rightarrow \mathbb{A}$ s.t. $e'_k \simeq_{\pi_k} e_k$ and $\vec{\Pi}'(k) \circ \pi_k = \vec{\Pi}(k)$.

Proposition 1. *For every configurations b_0, b, b' , if $b_0 \mapsto^* b$ and $b' \simeq b$, then $b_0 \mapsto^* b'$.*

We prove that given a multithreaded program, its associated CH-PDS and CVH-PDS are bisimilar. For that, we exhibit a bisimulation relation between configurations of the two systems. Intuitively, this relation maps a configuration of the CVH-PDS to a configuration of the CH-PDS by applying the updating operations through the configuration. (Indeed, visible heaps stored in the stacks of the CVH-PDS model are not up to date since they correspond to views of the heap at the moment of their memorization.)

Theorem 1. *The relations \Rightarrow and \mapsto define bisimilar transition systems.*

4.5 Program semantics based on normalized visible heaps

Normalized visible heaps: Visible heaps in *normal form* are obtained by numbering nodes according to a depth-first traversal of the heap, for a given ordering of global and local variables, and the fixed order on the successor fields $\{s_1, \dots, s_n\}$.

Let us consider a bijection $\eta : L \cup G \rightarrow [1, |L \cup G|]$. Then, given an environment $\langle n, VH, \pi \rangle$ with $VH = (A, \Delta, \Gamma, \Lambda)$, we define the class $\llbracket \langle n, VH, \pi \rangle \rrbracket_\eta$ to be the set of all environments $\langle n, VH', \pi \circ \beta \rangle$ with $VH' = (A', \Delta', \Gamma', \Lambda')$ s.t. (1) $VH \simeq_\beta VH'$, (2) $A' = [1, |A|]$, and (3) in the graph (V, Δ'_V) where $V = \text{Reach}_{\Delta'}(\Lambda'(L) \cup \Gamma'(G))$ and $\Delta'_V = \Delta' \cap [S \rightarrow (V \rightarrow V)]$, vertices correspond to the depth-first-traversal induced by the order η on root nodes (and the fixed order on successor fields labeling the edges of the graph).

Notice that, if all vertices in VH are reachable from $L \cup G$, then $\llbracket \langle n, VH, \pi \rangle \rrbracket_\eta$ contains a *single* element. Otherwise, there must be cut points (and may be other vertices reachable from them) which are not reachable from $L \cup G$. In that case, $\llbracket \langle n, VH, \pi \rangle \rrbracket_\eta$ contains environments (with identical reachable parts from $L \cup G$) corresponding to different permutations of the vertices reachable from cut points but not from $L \cup G$.

Proposition 2. *For every η , the set $\llbracket \langle n, VH, \pi \rangle \rrbracket_\eta$ is finite (if VH is finite).*

Concurrent pushdown systems on normalized visible heaps: Let us fix η . We define a program semantics where visible heaps are considered modulo the η -equivalence. Let \mathbb{V}_η be the set of all $\llbracket e \rrbracket_\eta$, for $e \in \mathbb{V}$. We consider the set of transition rules \mathcal{R}_η corresponding to the restriction of the set \mathcal{R} to the alphabet \mathbb{V}_η :

$$\mathcal{R}_\eta = \{((i, \vec{\Pi}), [\alpha_1, \dots, \alpha_m]) \hookrightarrow ((j, \vec{\Pi}'), [\alpha'_1, \dots, \alpha'_m]) \in \mathcal{R} \mid \forall j, \alpha_j, \alpha'_j \in \mathbb{V}_\eta^*\}.$$

Let \mathbb{C}_η be the set of all configurations $((i, \vec{\Pi}), [\alpha_1, \dots, \alpha_m])$ such that $\forall j, \alpha_j \in \mathbb{V}_\eta^*$. Then, let $\mapsto_{r, \eta}$ be the restriction of the transition relation \mapsto_r to configurations in \mathbb{C}_η . Let \mapsto_η be the union of the relations $\mapsto_{r, \eta}$ for all transition rules r . The relations $\mapsto_{\text{loc}, \eta}$, $\mapsto_{\text{sw}, \eta}$, and \rightsquigarrow_η are also defined as previously in terms of the restricted relations $\mapsto_{r, \eta}$.

Theorem 2. *The relations \mapsto and \mapsto_η define bisimilar transition systems.*

5 Reachability analysis

Bounded visible heap depth programs: Let $k \in \mathbb{N}$. A visible heap $VH = (A, \Delta, \Gamma, \Lambda, C)$ is k -bounded if (1) $|\text{Reach}_\Delta(\Lambda(L) \cup \Gamma(G))| \leq k$, and (2) $\forall i \in [1, m]. |\text{Reach}_\Delta(C(i))| \leq k$. Notice that k -boundedness does not imply that $|A| \leq k$ since there may exist vertices which are reachable from cut points but not from local/global variables. A sequence $\langle n_1, VH_1, \pi_1 \rangle \dots \langle n_j, VH_j, \pi_j \rangle \in \mathbb{V}^*$ is k -bounded iff $\forall i \in [1, j], VH_i$ is k -bounded. A configuration $((i, \vec{\Pi}), [\alpha_1, \dots, \alpha_m])$ is k -bounded iff $\forall j \in [1, m], \alpha_j$ is k -bounded.

Theorem 3. *Given $k \geq 1$, for every k -bounded \Rightarrow -computation in the CH-PDS model of a program there is a bisimilar k -bounded \mapsto -computation in its CVH-PDS model.*

Bounded heap depth reachability analysis of sequential programs: We extend to VH-PDS the automata-based construction of *post*/pre** images for pushdown systems (see, e.g., [1, 5, 4]). We assume in the sequel that visible heaps are in normal form according to some fixed ordering η on local and global variables. Sets of configurations are recognized by finite-state automata over the alphabet of normalized environments called *Conf-automata*. More precisely, initial states in these automata correspond to mappings Π , and edges are labeled by elements of \mathbb{V}_η . Then, a configuration (Π, α) is accepted by the automaton if starting from the initial state Π there is an accepting run for the sequence $\alpha \in \mathbb{V}_\eta^*$.

Given $k \geq 1$, and a regular set of k -bounded (local) configurations recognized by a Conf-automaton \mathcal{A} , we apply a saturation based algorithm $\text{Closure}_\eta(\mathcal{A}, k)$ which constructs a sequence of Conf-automata with increasing languages, each of them being obtained from the previous one using one of the transition rules of the pushdown system. The difference here with the existing constructions for pushdown systems is that the stack alphabet of the built automata may increase at each step. Therefore, we restrict the construction to the k -bounded environments. Then, by Proposition 2, the algorithm terminates and produces a finite Conf-automaton representing all (forward/backward) reachable configuration by k -bounded \mapsto_η -computations, which is a subset of the set of all reachable configurations from $L(\mathcal{A})$ (by Theorems 1 and 2), and which contains the set of all reachable configurations by k -bounded \Rightarrow -computations (by Theorem 3).

Reachability analysis of concurrent programs: Let us first introduce some definitions and notations: A *special* Conf-automaton is a Conf-automaton for which there exists a pair (Π, e) such that, for every local configuration $(\Pi', \alpha) \in L(\mathcal{A})$, $\Pi' = \Pi$ and there exists $\alpha' \in \mathbb{V}_\eta^*$ such that $\alpha = e\alpha'$, i.e., all words accepted by \mathcal{A} start with (Π, e) . The pair (Π, e) is then denoted $\widehat{\mathcal{A}}$. An *aggregate* is a tuple $(\mathcal{A}_1, \dots, \mathcal{A}_m)$ of special Conf-automata. Such an aggregate defines the set of global configurations $L(\mathcal{A}_1, \dots, \mathcal{A}_m) = \{((i, \vec{\Pi}), [\alpha_1, \dots, \alpha_m]) : i \in [1, m], \forall j \in [1, m]. (\vec{\Pi}(j), \alpha_j) \in L(\mathcal{A}_j)\}$. A finite set of aggregates defines the union of the languages defined by all its elements.

Given a Conf-automaton \mathcal{A} and (Π, e) with $e \in \mathbb{V}_\eta$, let $\mathcal{A}^{(\Pi, e)}$ be an automaton recognizing the language $L(\mathcal{A}) \cap (\Pi \times e\mathbb{V}_\eta^*)$. Clearly, $\mathcal{A}^{(\Pi, e)}$ is a special automaton. Given (Π, e) and a special automaton \mathcal{A} , we denote by $(\Pi, e) \triangleright \mathcal{A}$ the special automaton recognizing the language $\{(\Pi, e\alpha) \mid (\Pi', e'\alpha) \in L(\mathcal{A})\}$, i.e., the language of \mathcal{A} where the first symbol of every word is replaced by e and the initial state is replaced by Π .

We consider w.l.o.g. that the reachability analysis starts from a single initial configuration where the heap is empty and all pointer variables are equal to *null*: For each $i \in [1, m]$, let \mathcal{A}_0^i be the special automaton recognizing the (singleton) language $\{Id_{\{\top, \perp\}}, \langle n_{init}^i, VH_0, Id_{\{\top, \perp\}} \rangle\}$, where n_{init}^i is the entry node of the main procedure of process i , and $VH_0 = (\{\top, \perp\}, \Delta_\top, \Gamma_\perp, \Lambda_\perp, \vec{\emptyset})$ with Δ_\top being the function mapping \top to each address for all successor fields, Γ_\perp (resp. Λ_\perp) being the functions mapping \perp to each global (resp. local) variable, and $\vec{\emptyset}$ being the vector of functions mapping an empty set of cut points to each process.

Finally, given h_1, \dots, h_m such that for every $\ell \in [1, m]$, $h_\ell = (\vec{\Pi}(\ell), e_\ell)$, and given $i, j \in [1, m]$, $i \neq j$, we denote by $\text{Update}_{i,j}(h_1, \dots, h_m)$ the set of tuples (h'_1, \dots, h'_m) such that: (1) for every $\ell \in [1, m]$, $h'_\ell = (\vec{\Pi}'(\ell), e'_\ell)$, and (2) $((i, \vec{\Pi}), [e_1, \dots, e_m]) \hookrightarrow_{\text{sw}} ((j, \vec{\Pi}'), [e'_1, \dots, e'_m])$, where $\hookrightarrow_{\text{sw}}$ refers to the Switch inference rule (see section 4.3).

We are now ready to present our bounded reachability analysis algorithm. The algorithm is given in Figure 1. The input is a CVH-PDS, the bound on context switches K , and the bound on the size of visible heaps k . The algorithm computes (in *Reach*) the set of all reachable configurations by k -bounded \rightsquigarrow_η^K -computations: A set of *tasks* is maintained in a structure *todo*. A task is a triple (ℓ, i, \mathcal{A}) where ℓ is the number of context switches done so far, i is the index of the process chosen to be active, and \mathcal{A} is an aggregate. Initially, *todo* contains the initial configuration of the program with all possible starting active process (lines 2-3). Then, the treatment of a task is as follows: if ℓ has already reached K then \mathcal{A} is added to *Reach* (lines 6-7). Otherwise, the set of reachable k -bounded configurations from \mathcal{A} by process i is computed (line 9), and then new tasks are produced corresponding to all possible context switches from i to some other process j (lines 10-13). For that, a case splitting is performed according to all possible visible heaps reached by process i (line 10), corresponding to head environments of its possible stacks. Then, the local heaps of all processes are updated, and tasks (where ℓ is incremented) are defined for all possible next active processes $j \neq i$ (lines 11-12) and added to *todo* (line 13).

Input: An m -dim CVH-PDS, and two integers $K, k \geq 1$
Output: The set of reachable configurations by k -bounded $\rightsquigarrow_{\eta}^K$ -computations.

```

1  Reach  $\leftarrow \emptyset$ ;
2  for all  $i \in \{1, \dots, m\}$ 
3    todo  $\leftarrow \{(0, i, \mathcal{A}_0^1, \dots, \mathcal{A}_0^m)\}$ ;
4  while todo  $\neq \emptyset$  do
5    pop (level,  $i, \mathcal{A}_1, \dots, \mathcal{A}_m$ ) from todo with minimal level;
6    if level =  $K$  then
7      Reach  $\leftarrow$  Reach  $\cup \{(\mathcal{A}_1, \dots, \mathcal{A}_m)\}$ ;
8    else
9      let  $\mathcal{B}_i = \text{Closure}_{\eta}(\mathcal{A}_i, k)$  in
10     for all  $(\Pi, e'_i)$  such that  $L(\mathcal{B}_i^{(\Pi, e'_i)}) \neq \emptyset$ 
11       for all  $j \in \{1, \dots, m\}$  such that  $j \neq i$ 
12         for all  $(h_1, \dots, h_m) \in \text{Update}_{i,j}(\widehat{\mathcal{A}}_1, \dots, (\Pi, e'_i), \dots, \widehat{\mathcal{A}}_m)$ 
13           todo  $\leftarrow$  (level + 1,  $j, h_1 \triangleright \mathcal{A}_1, \dots, h_i \triangleright \mathcal{B}_i^{(\Pi, e'_i)}, \dots, h_m \triangleright \mathcal{A}_m$ );
14  return Reach

```

Fig. 1. Algorithm for bounded context-switch/heap depth reachability on CVH-PDS.

References

1. A. Bouajjani, J. Esparza, and O. Maler. Reachability analysis of pushdown automata: Application to model-checking. In *CONCUR'97*. LNCS 1243, 1997.
2. A. Bouajjani, J. Esparza, S. Schwoon, and J. Strejcek. Reachability analysis of multithreaded software with asynchronous communication. In *FSTTCS'05*. LNCS 3821, 2005.
3. A. Bouajjani, S. Fratani, and S. Qadeer. Bounded context switch analysis of multithreaded programs with dynamic linked structures. Technical Report 2007-02, LIAFA lab, January 2007. Available at <http://www.liafa.jussieu.fr/~abou/publis.html>.
4. J. Esparza, D. Hansel, P. Rossmanith, and S. Schwoon. Efficient algorithms for model checking pushdown systems. In *CAV'00*. LNCS 1855, 2000.
5. A. Finkel, B. Willems, and P. Wolper. A direct symbolic approach to model checking pushdown systems. *Electr. Notes Theor. Comput. Sci.*, 9, 1997.
6. P. Godefroid. Model checking for programming languages using Verisoft. In *POPL'97*. ACM, 1997.
7. M. Musuvathi and S. Qadeer. Iterative context bounding for systematic testing of multithreaded programs. In *PLDI'07*. ACM, 2007.
8. S. Qadeer and S. Rajamani. Deciding assertions in programs with references. Technical Report MSR-TR-2005-08, Microsoft Research, September 2005.
9. S. Qadeer and J. Rehof. Context-bounded model checking of concurrent software. In *TACAS'05*. LNCS 3440, 2005.
10. N. Rinetzky, J. Bauer, T. Reps, M. Sagiv, and R. Wilhelm. A semantics for procedure local heaps and its abstractions. In *POPL'05*. ACM, 2005.
11. N. Rinetzky, M. Sagiv, and E. Yahav. Interprocedural shape analysis for cutpoint-free programs. In *SAS'05*. LNCS 3672, 2005.