Analysis of Recursively Parallel Programs

Ahmed Bouajjani, Michael Emmi

Analysis of Recursively Parallel Programs *

Ahmed Bouajjani
LIAFA, Université Paris Diderot, France
abou@liafa.jussieu.fr

Michael Emmi †
LIAFA, Université Paris Diderot, France
mje@liafa.jussieu.fr

Abstract

We propose a general formal model of isolated hierarchical parallel computations, and identify several fragments to match the concurrency constructs present in real-world programming languages such as Cilk and X10. By associating fundamental formal models (vector addition systems with recursive transitions) to each fragment, we provide a common platform for exposing the relative difficulties of algorithmic reasoning. For each case we measure the complexity of deciding state-reachability for finite-data recursive programs, and propose algorithms for the decidable cases. The complexities which include PTIME, NP, EXPSPACE, and 2EXPTIME contrast with undecidable state-reachability for recursive multi-threaded programs.

1. Introduction

Despite the ever-increasing importance of concurrent software (e.g., for designing reactive applications, or parallelizing computation across multiple processor cores), concurrent programming and concurrent program analysis remain challenging endeavors. The most widely available facility for designing concurrent applications is *multithreading*, where concurrently executing sequential threads nondeterministically interleave their accesses to shared memory. Such nondeterminism leads to rarely-occurring "Heisenbugs" which are notoriously difficult to reproduce and repair. To prevent such bugs programmers are faced with the difficult task of preventing undesirable interleavings, e.g., by employing lock-based synchronization, without preventing benign interleavings—otherwise the desired reactivity or parallelism is forfeited.

The complexity of multi-threaded program analysis seems to comply with the perceived difficulty of multi-threaded programming. The state-reachability problem for multi-threaded programs is PSPACE-complete [22] with a finite number of finite-state threads, and undecidable [31] with recursive threads. Current analysis approaches either explore an underapproximate concurrent semantics by considering relatively few interleavings [9, 23] or explore a coarse overapproximate semantics via abstraction [13, 19].

Explicitly-parallel programming languages have been advocated to avoid the intricate interleavings implicit in program syntax [25], and several such industrial-strength languages have been developed [2, 5, 6, 18, 26, 32, 34]. Such systems introduce various mechanisms for creating (e.g., **fork**, **spawn**, **post**) and consuming (e.g., **join**, **sync**) concurrent computations, and either encourage (through recommended programming practices) or ensure (through static analyses or runtime systems) that parallel computations execute in isolation without interference from others, through data-partitioning [6], data-replication [5], functional programming [18], message passing [29], or version-based memory access models [34], perhaps falling back on transactional mechanisms [24] when complete isolation is impractical. Although few of these systems behave deterministically, consuming one concurrent computation at a time, many are sensitive to the order in which multiple isolated computations are consumed. Furthermore, some allow computations creating an unbounded number of sub-computations, returning to their superiors an unbounded number of handles to unfinished computations. Even without multithreaded interleaving, nondeterminism in the order in which an unbounded number of computations are consumed has the potential to make program reasoning complex.

In this work we investigate key questions on the analysis of interleaving-free programming models. Specifically, we ask to what extent such models simplify program reasoning, how those models compare with each other, and how to design appropriate analysis algorithms. We attempt to answer these questions as follows:

- We introduce a general interleaving-free parallel programming model on which to express the features found in popular parallel programming languages (Section 2).
- We discover a surprisingly-complex feature of some existing languages: even simple classes of programs with the ability to pass unfinished computations both to and from subordinate computations have undecidable state-reachability problems (Section 2.4).
- We show that the concurrency features present in many real-world programming languages such as Cilk, X10, and Multilisp are captured precisely (modulo the possibility of interleaving) by various fragments of our model (Sections 4 and 6).
- For fragments corresponding to real-world language features, we measure the complexity of computing state-reachability for finite-data programs, and provide, in most cases, asymptotically optimal state-reachability algorithms (Sections 5 and 7).

Our focus on finite-data programs without interleaving is a means to measuring complexity for the sake of comparison, required since state-reachability for infinite-data or multi-threaded programs is generally undecidable. Applying our algorithms in practice may rely on data abstraction [16], and separately ensuring isolation [24], or approximating possible interleavings [9, 13, 19, 23]; still, our handling of computation-order non-determinism is precise.

The major distinguishing language features are whether a single or an arbitrary number of subordinate computations are waited for at once, and whether the scope of subordinate computations is confined. Generally speaking, reasoning for the "single-wait" case of Section 4 is less difficult than for the "multi-wait" case of Section 6, and we demonstrate a range of complexities[1] from PTIME, NP, EXPSPACE, and 2EXPTIME for various scoping restrictions in Sections 5 and 7. Despite these worst-case complexities, a promising line of work has

* Partially supported by the project ANR-09-SEGI-016 Veridyc.

† Supported by a post-doctoral fellowship from the Fondation Sciences Mathématiques de Paris.

[0] Proofs to technical results are contained in the appendices.

[1] In order to isolate concurrent complexity from the exponential factor in the number of program variables, we consider a fixed number of variables in each procedure frame; this allows us a PTIME point-of-reference for state-reachability in recursive sequential programs [33].

$$
\begin{array}{rcl}
P & ::= & (\ \textbf{proc}\ p\ (\textbf{var}\ \texttt{l}\colon T)\ s\ )^* \\
s & ::= & s;\ s \ \mid\ \texttt{l} := e \ \mid\ \textbf{skip} \ \mid\ \textbf{assume}\ e \\
 & \mid & \textbf{if}\ e\ \textbf{then}\ s\ \textbf{else}\ s \ \mid\ \textbf{while}\ e\ \textbf{do}\ s \\
 & \mid & \textbf{call}\ \texttt{l} := p\ e \ \mid\ \textbf{return}\ e \\
 & \mid & \textbf{post}\ r \leftarrow p\ e\ \vec{r}\ d \ \mid\ \textbf{ewait}\ r \ \mid\ \textbf{await}\ r
\end{array}
$$

Figure 1. The grammar of recursively parallel programs. Here T is an unspecified type, p ranges over procedure names, e over expressions, r over regions, and d over return-value handlers.

already demonstrated effective algorithms for practically-occurring EXPSPACE-complete state-reachability problem instances based on simultaneously computing iterative under- and over-approximations, and rapidly converging to a fixed point [15, 20].

We thus present a classification of concurrency constructs, connecting programming language features to fundamental formal models, which highlight the sources of concurrent complexity resulting from each feature, and provide a platform for comparing the difficulty of formal reasoning in each. We hope that these results may be used both to guide the design of impactful program analyses, as well as to guide the design and choice of languages appropriate for various programming problems.

2. Recursively Parallel Programs

We consider a simple concurrent programming model where computations are hierarchically divided into isolated parallely executing tasks. Each task executes sequentially while maintaining *regions* (i.e., containers) of *handles* to other tasks. The initial task begins without task handles. When a task t creates a subordinate (child) task u, t stores the handle to u in one of its regions, at which point t and u begin to execute in parallel. The task u may then recursively create additional parallel tasks, storing their handles in its own regions. At some later point when t requires the result computed by u, t must *await* the completion of u—i.e., blocking until u has finished—at which point t consumes its handle to u. When u does complete, the value it returns is combined with the current state of t via a programmer-supplied *return-value handler*. In addition to creating and consuming subordinate tasks, tasks can transfer ownership of their subordinate tasks to newly-created tasks—by initially passing to the child a subset of task handles—and to their superiors upon completion—by finally passing to the parent unconsumed tasks.

This model permits vastly concurrent executions. Each task along with all the tasks it has created execute completely in parallel. As tasks can create tasks recursively, the total number of concurrently executing tasks has no bound, even when the number of handles stored by each task is bounded.

2.1 Program Syntax

Let Procs be a set of procedure names, Vals a set of values, Exprs a set of expressions, Regs a set of region identifiers, and $\mathsf{Rets} \subseteq (\mathsf{Vals} \to \mathsf{Stmts})$ a set of return-value handlers. The grammar of Figure 1 describes our language of *recursively parallel programs*. We intentionally leave the syntax of expressions e unspecified, though we do insist Vals contains $\mathbf{true}$ and $\mathbf{false}$, and Exprs contains Vals and the *(nullary) choice operator* $\star$. We refer to the class of programs restricted to a finite set of values as *finite-value* programs, and to the class of programs restricted to at most $n \in \mathbb{N}$ (resp., 1) region identifiers as *n-region* (resp., *single-region*) or, more generally, *finite-region* programs. A *sequential program* is a program without $\mathbf{post}$, $\mathbf{ewait}$, and $\mathbf{await}$ statements.

Each program P declares a sequence of procedures named $p_0 \ldots p_i \in \mathsf{Procs}^*$, each p having single type-T parameter $\texttt{l}$ and a top-level statement denoted s_p; as statements are built inductively by composition with control-flow statements, s_p describes the entire body of p. The set of program statements s is denoted Stmts. Intuitively, a $\mathbf{post}\ r \leftarrow p\ e\ \vec{r}\ d$ statement stores the handle to a newly-created task executing procedure p in the region r; besides the procedure argument e, the newly-created task is passed a subset of the parent's task handles in regions $\vec{r}$, and a return-value handler d. The $\mathbf{ewait}\ r$ statement blocks execution until *some* task whose handle is stored in region r completes, at which point its return-value handler is executed. Similarly, the $\mathbf{await}\ r$ statement blocks execution until *all* tasks whose handles are stored in region r complete, at which point all of their return-value handlers are executed, in some order. We refer to the $\mathbf{call}$, $\mathbf{return}$, $\mathbf{post}$, $\mathbf{ewait}$ and $\mathbf{await}$ as *inter-procedural statements*, and the others as *intra-procedural statements*, and insist that return-value handlers are comprised only of intra-procedural statements. The $\mathbf{assume}\ e$ statement proceeds only when e evaluates to $\mathbf{true}$—we use this statement in subsequent sections to block undesired executions in our encodings of other parallel programming models.

Example 1. The Fibonacci function can be implemented as a single-region recursively parallel program as follows.

```
proc fib (var n: ℕ)
    var sum: ℕ
    if n < 2 then
        return 1
    else
        post r ← fib (n-1) ε (λv. sum := sum + v);
        post r ← fib (n-2) ε (λv. sum := sum + v);
        await r;
        return sum
```

Alternate implementations are possible, e.g., by replacing the $\mathbf{await}$ statement by two $\mathbf{ewait}$ statements, or storing the handles to the recursive calls in separate regions. Note that in this implementation task-handles are not passed to child tasks (ε specifies the empty region sequence) nor to parent tasks (all handles are consumed by the $\mathbf{await}$ statement before returning).

The programming language we consider is simple yet expressive, since the syntax of types and expressions is left free, and we lose no generality by considering only a single variable per procedure.

2.2 Parallel Semantics with Task-Passing

Unlike recursive sequential programs, whose semantics is defined over *stacks* of procedure frames, the semantics of recursively parallel programs is defined over *trees* of procedure frames. Intuitively, the frame of each posted task becomes a child of the posting task's frame. Each step of execution proceeds either by making a single intra-procedural step of some frame in the tree, creating a new frame by posting a task, or removing a frame by consuming a completed task; unconsumed sub-task frames of a completed task are added as children to the completed task's parent.

A *task* $\langle \ell, s, d \rangle$ is a valuation $\ell \in \mathsf{Vals}$ to the procedure-local variable $\texttt{l}$, along with a statement s to be executed, and a return-value handler $d \in \mathsf{Rets}$. (Here s describes the entire body of a procedure p that remains to be executed, and is initially set to p's top-level statement s_p.) A *tree configuration* c is a finite unordered tree of task-labeled vertices and region-labeled edges, and the set of configurations is denoted $\mathsf{Configs}$. Let $\mathbb{M}[\mathsf{Configs}]$ denote the set of configuration multisets. We represent configurations inductively, writing $\langle t, m \rangle$ for the tree with t-labeled root whose child subtrees are given by a *region valuation* $m : \mathsf{Regs} \to \mathbb{M}[\mathsf{Configs}]$: for $r \in \mathsf{Regs}$, the multiset $m(r)$ specifies the collection of subtrees connected to the root of $\langle t, m \rangle$ by an r-edge. The *initial region valuation* $m_\emptyset$ is defined by $m_\emptyset(r) \stackrel{\text{def}}{=} \emptyset$ for all $r \in \mathsf{Regs}$. The singleton region valuation $(r \mapsto c)$ maps r to $\{c\}$, and $r' \in \mathsf{Regs} \setminus \{r\}$ to $\emptyset$, and the union $m_1 \cup m_2$ of region valuations is

defined by the multiset union of each valuation: $(m_1 \cup m_2)(r) \stackrel{\text{def}}{=} m_1(r) \cup m_2(r)$ for all $r \in \mathsf{Regs}$. The projection $m \mid_{\vec{r}}$ of a region valuation m to a region sequence $\vec{r}$ is defined by $m \mid_{\vec{r}}(r') = m(r')$ when r' occurs in $\vec{r}$, and $m \mid_{\vec{r}}(r') = \emptyset$ otherwise.

For expressions without program variables, we assume the existence of an evaluation function $[\![\cdot]\!]_{\mathrm{e}} : \mathsf{Exprs} \to \wp(\mathsf{Vals})$ such that $[\![\star]\!]_{\mathrm{e}} = \mathsf{Vals}$. For convenience, we define

$$e(\langle \ell, s, d \rangle) \stackrel{\text{def}}{=} e(\ell) \stackrel{\text{def}}{=} [\![e[\ell/\mathtt{1}]]\!]_{\mathrm{e}}$$

—as $\mathtt{1}$ is the only variable, the expression $e[\ell/\mathtt{1}]$ has no free variables.

To reduce clutter and focus on the relevant parts of transition rules in the program semantics, we introduce a notion of contexts. A *configuration context* C is a tree with a single $\diamond$-labeled leaf, task-labeled vertices and leaves otherwise, and region-labeled edges. We write $C[c]$ for the configuration obtained by substituting a configuration c for the unique $\diamond$-labeled leaf of C. We use configuration contexts to isolate individual task transitions, writing, for instance $C[\langle t, m \rangle] \to C[\langle t', m \rangle]$ to indicate an intra-procedural transition of the task t. Similarly a *statement context* $S = \diamond;\ s_1;\ \ldots;\ s_i$ is a $\diamond$-led sequence of statements, and we write $S[s_0]$ for the statement obtained by substituting a statement s_0 for the unique occurrence of $\diamond$ as the first symbol of S, indicating that s_0 is the next-to-be-executed statement. A *task-statement context* $T = \langle \ell, S, d \rangle$ is a task with a statement context S in place of a statement, and we write $T[s]$ to indicate that s is the next statement to be executed in the task $\langle \ell, S[s], d \rangle$. Finally, we write $C[\langle T[s_1], m \rangle] \to C[\langle T[s_2], m' \rangle]$ to denote a transition of a task executing a statement s_1 and replacing s_1 by s_2—normally s_2 is the **skip** statement. Since the current statement s of a task $T[s]$ does not effect expression evaluation, we liberally write $e(T)$ to denote the evaluation $e(T[s])$.

We say a task $t = \langle \ell, S[s], d \rangle$ is *completed* when its next-to-be-executed statement s is $\mathbf{return}\ e$, in which case we define $\mathsf{rvh}(t) \stackrel{\text{def}}{=} \{d(v) : v \in e(\ell)\}$ as the set of possible return-value handler statements for t; $\mathsf{rvh}(t)$ is undefined when t is not completed.

Figure 2 and Figure 3 define the transition relation $\to^{\mathrm{rpp/p}}$ of recursively parallel programs as a set of operational steps on configurations. The intra-procedural transitions $\to^{\mathrm{seq}}$ of individual tasks in Figure 2 are standard. More interesting are the inter-procedural transitions of Figure 3, which implicitly include a transition $C[\langle t_1, m \rangle] \to_P^{\mathrm{rpp/p}} C[\langle t_2, m \rangle]$ whenever $t_1 \to_P^{\mathrm{seq}} t_2$. The POST-T rule creates a procedure frame to execute in parallel, and links it to the current frame by the given region, passing ownership of tasks in the specified region sequence to the newly-created frame. The ∃WAIT-T rule consumes the result of a single child frame in the given region, and applies the return-value handler to update the parent frame's local valuation. Similarly, the ∀WAIT-NEXT-T and ∀WAIT-DONE-T rules consume the results of every child frame in the given region, applying their return handlers in the order they are consumed. The semantics of **call** statements reduces to that of **post** and **ewait**: supposing an unused region identifier $\mathtt{r}_{call}$, we translate each statement $\mathbf{call}\ \mathtt{1}\ \mathtt{:=}\ p\ e$ into the sequence

$$\mathbf{post}\ \mathtt{r}_{call} \leftarrow p\ e\ \varepsilon\ \mathtt{d}_{call};$$
$$\mathbf{ewait}\ \mathtt{r}_{call},$$

where $\mathtt{d}_{call}(v) \stackrel{\text{def}}{=} \mathtt{1}\ \mathtt{:=}\ v$ is the return-value handler which simply writes the entire return value v into the local variable $\mathtt{1}$, and ε denotes an empty sequence of region identifiers.

A *parallel execution of a program P (from c_0 to c_j)* is a configuration sequence $c_0 c_1 \ldots c_j$ where $c_i \to_P^{\mathrm{rpp/p}} c_{i+1}$ for $0 \le i < j$. An initial condition $\iota = \langle p_0, \ell_0 \rangle$ is a procedure $p_0 \in \mathsf{Procs}$ along with a value $\ell_0 \in \mathsf{Vals}$. A configuration $\langle \langle \ell_0, s, d \rangle, m_\emptyset \rangle$ is called $\langle p_0, \ell_0 \rangle$-*initial* when s is the top-level statement of p_0. A configuration c_f is called ℓ_f-*final* when there exists a context C such that $c_f = C[\langle t, m \rangle]$ and $\mathtt{1}(t) = \ell_f$. We say a valuation ℓ is *reachable in P* from ι when there exists an execution of P from some c_0 to c_f, where c_0 is ι-initial and c_f is ℓ-final.

$$\textsc{Post-T}\quad \frac{v \in e(T) \qquad m' = m \setminus m \mid_{\vec{r}} \cup \left(r \mapsto \left\langle \langle v, s_p, d \rangle, m \mid_{\vec{r}} \right\rangle\right)}{C[\langle T[\mathbf{post}\ r \leftarrow p\ e\ \vec{r}\ d], m \rangle] \xrightarrow[P]{\mathrm{rpp/p}} C\left[\left\langle T[\mathbf{skip}], m' \right\rangle\right]}$$

$$\exists\textsc{wait-T}\quad \frac{m_1 = (r \mapsto \langle t_2, m_2 \rangle) \cup m_1' \qquad s \in \mathsf{rvh}(t_2)}{C[\langle T_1[\mathbf{ewait}\ r], m_1 \rangle] \xrightarrow[P]{\mathrm{rpp/p}} C\left[\left\langle T_1[s], m_1' \cup m_2 \right\rangle\right]}$$

$$\forall\textsc{wait-Next-T}\quad \frac{m_1 = (r \mapsto \langle t_2, m_2 \rangle) \cup m_1' \qquad s \in \mathsf{rvh}(t_2)}{C[\langle T_1[\mathbf{await}\ r], m_1 \rangle] \xrightarrow[P]{\mathrm{rpp/p}} C\left[\left\langle T_1[s;\ \mathbf{await}\ r], m_1' \cup m_2 \right\rangle\right]}$$

$$\forall\textsc{wait-Done-T}\quad \frac{m(r) = \emptyset}{C[\langle T[\mathbf{await}\ r], m \rangle] \xrightarrow[P]{\mathrm{rpp/p}} C[\langle T[\mathbf{skip}], m \rangle]}$$

Figure 3. The tree-based transition relation for parallely-executing recursively parallel programs with task-passing.

Problem 1 (State-Reachability)**.** *The* state-reachability problem *is to determine, given an initial condition ι of a program P and a valuation ℓ, whether ℓ is reachable in P from ι.*

2.3 Sequential Semantics with Task-Passing

Parallelly-executing tasks do not have the ability to influence the values each other compute. Thus, when it comes to state-reachability, the manner in which executions of parallely-executing tasks interleave is unimportant. In this section we leverage this fact and focus on a particular execution order in which at any moment only a single task is enabled. When the currently enabled task encounters and **ewait/await** statement, suspending execution to wait for a subordinate task t, t becomes the currently-enabled task; when t completes, control returns to its waiting parent. At any moment only the tasks along one path ρ in the configuration tree have ever been enabled, and all but the last task in ρ are waiting for their child in ρ to complete. This execution order can be encoded as an equivalent stack-based operational semantics, which immediately reduces analysis on recursively parallel programs to analysis on sequential programs with a particular (generally unbounded) auxiliary storage device, needed to store an arbitrary number of pending tasks. We then interpret the **ewait** and **await** statements as (synchronous) procedure calls to compute the values returned by pending tasks.

We define a *frame* to be a configuration in the sense of the tree-based semantics of Section 2.2, i.e., a finite unordered tree of task-labeled vertices and region-labeled edges. (Here all non-root nodes in the tree are posted tasks that have yet to take a single step of execution.) In our stack-based semantics, a *stack configuration* c is a sequence of frames, representing a procedure activation stack.

Figures 2 and 4 define the sequential transition relation $\to^{\mathrm{rpp/s}}$ of recursively parallel programs as a set of operational steps on configurations. The inter-procedural transitions of Figure 4 implicitly include a transition $\langle t_1, m \rangle\, c \to_P^{\mathrm{rpp/s}} \langle t_2, m \rangle\, c$ whenever $t_1 \to_P^{\mathrm{seq}} t_2$. Interesting here are the rules for **ewait** and **await**. The ∃WAIT-S rule blocks the currently executing frame to obtain the result for a single, nondeterministically chosen, frame c_0 in the given region, by pushing c_0 onto the activation stack. Similarly, the ∀WAIT-NEXT-S and ∀WAIT-DONE-S rules block the currently executing frame to obtain the results for every task in the given region, in a

$$\textsc{Skip}\ \frac{}{T[\mathbf{skip};\ s] \xrightarrow[P]{\text{seq}} T[s]} \qquad \textsc{Assume}\ \frac{\mathbf{true} \in e(T)}{T[\mathbf{assume}\ e] \xrightarrow[P]{\text{seq}} T[\mathbf{skip}]} \qquad \textsc{If-Then}\ \frac{\mathbf{true} \in e(T)}{T[\mathbf{if}\ e\ \mathbf{then}\ s_1\ \mathbf{else}\ s_2] \xrightarrow[P]{\text{seq}} T[s_1]} \qquad \textsc{If-Else}\ \frac{\mathbf{false} \in e(T)}{T[\mathbf{if}\ e\ \mathbf{then}\ s_1\ \mathbf{else}\ s_2] \xrightarrow[P]{\text{seq}} T[s_2]}$$

$$\textsc{Assign}\ \frac{\ell' \in e(\ell)}{\langle \ell, S[\texttt{l := e}], d\rangle \xrightarrow[P]{\text{seq}} \langle \ell', S[\mathbf{skip}], d\rangle} \qquad \textsc{Loop-Do}\ \frac{\mathbf{true} \in e(T)}{T[\mathbf{while}\ e\ \mathbf{do}\ s] \xrightarrow[P]{\text{seq}} T[s; \mathbf{while}\ e\ \mathbf{do}\ s]} \qquad \textsc{Loop-End}\ \frac{\mathbf{false} \in e(T)}{T[\mathbf{while}\ e\ \mathbf{do}\ s] \xrightarrow[P]{\text{seq}} T[\mathbf{skip}]}$$

Figure 2. The intra-procedural transition relation for recursively parallel programs.

$$\textsc{Post-S}\ \frac{v \in e(T) \qquad m' = m \setminus m\mid_{\vec{r}} \cup \left(r \mapsto \left\langle \langle v, s_p, d\rangle, m\mid_{\vec{r}}\right\rangle\right)}{\langle T[\mathbf{post}\ r \leftarrow p\ e\ \vec{r}\ d], m\rangle\ c \xrightarrow[P]{\text{rpp/s}} \langle T[\mathbf{skip}], m'\rangle\ c}$$

$$\exists\textsc{wait-S}\ \frac{m = (r \mapsto c_0) \cup m'}{\langle T[\mathbf{ewait}\ r], m\rangle\ c \xrightarrow[P]{\text{rpp/s}} c_0\ \langle T[\mathbf{skip}], m'\rangle\ c}$$

$$\forall\textsc{wait-Next-S}\ \frac{m = (r \mapsto c_0) \cup m'}{\langle T[\mathbf{await}\ r], m\rangle\ c \xrightarrow[P]{\text{rpp/s}} c_0\ \langle T[\mathbf{skip};\ \mathbf{await}\ r], m'\rangle\ c}$$

$$\forall\textsc{wait-Done-S}\ \frac{m(r) = \emptyset}{\langle T[\mathbf{await}\ r], m\rangle\ c \xrightarrow[P]{\text{rpp/s}} \langle T[\mathbf{skip}], m\rangle\ c}$$

$$\textsc{Return-S}\ \frac{s \in \mathsf{rvh}(t_1)}{\langle t_1, m_1\rangle\ \langle T_2[\mathbf{skip}], m_2\rangle\ c \xrightarrow[P]{\text{rpp/s}} \langle T_2[s], m_1 \cup m_2\rangle\ c}$$

Figure 4. The stack-based transition relation for sequentially-executing recursively parallel programs with task-passing.

nondeterministically-chosen order. Finally, the RETURN-S applies a completed task's return-value handler to update the parent frame's local valuation. The definitions of *sequential execution*, *initial*, and *reachable* are nearly identical to their parallel counterparts.

Lemma 1. *The parallel semantics and the sequential semantics are indistinguishable w.r.t. state reachability, i.e., for all initial conditions ι of a program P, the valuation ℓ is reachable in P from ι by a parallel execution if and only if ℓ is reachable in P from ι by a sequential execution.*

2.4 Undecidability of State-Reachability with Task-Passing

Recursively parallel programs allow pending tasks to be passed *bidirectionally*: both from completed tasks and to newly-created tasks. This capability makes the state-reachability problem undecidable—even for the very simple cases recursive programs with at least one region, and for non-recursive programs with at least two regions. Essentially, when pending tasks can be passed to newly-created tasks, it becomes possible to construct and manipulate unbounded task-chains by keeping a handle to most-recently created task, after having passed the handle of the previously-most-recently created task to the most-recently created task. We can then show that such unbounded chains of pending tasks can be used to simulate an arbitrary unbounded and ordered storage device.

Definition 1 (Task passing)**.** A program which contains a statement $\mathbf{post}\ r \leftarrow p\ e\ \vec{r}\ d$, such that $|\vec{r}| > 0$ is called *task-passing*.

The *task-depth* of a program P is the maximum length of a sequence $p_1 \ldots p_i$ of procedures in P such that each p_j contains a statement $\mathbf{post}\ r \leftarrow p_{i+j}\ e\ \vec{r}\ d$, for $0 < j < i$, and some $r \in \mathsf{Regs}$, $e \in \mathsf{Exprs}$, $\vec{r} \in \mathsf{Regs}^*$, and $d \in \mathsf{Rets}$. Programs with unbounded task-depth are *recursive*, and are otherwise *non-recrusive*.

Theorem 1. *The state-reachability problem for n-region finite-value task-passing parallel programs is undecidable for*

(a) non-recursive programs with $n > 1$, and
(b) recursive programs with $n > 0$.

The proof of Theorem 1 is given by two separate reductions from the emptiness problem for Turing machines to "single-wait" programs, i.e., those using **ewait** statements but not **await** statements. In essence, as each task-handle can point to an unbounded chain of task-handles, we can construct an unbounded Turing machine tape by using one task-chain to store the contents of cells to the left of the tape head, and another chain to store the contents of cells to the right of the tape head. If only one region is granted but recursion is allowed (i.e., as in (b)), we can still construct the tape using the task-chain for the cells right of the tape head, while using the (unbounded) procedure-stack to store the cells left of the head. When only one region is granted and recursion is not allowed, neither of these reductions work. Without recursion we can bound the procedure stack, and then we can show that single-stack machine suffices to encode the single unbounded chain of tasks.

3. Programs without Task Passing

Due to the undecidability result of Theorem 1 and our desire to compare the analysis complexities of parallel programming models, we consider, henceforth, unless otherwise specified, only non-task-passing programs, simplifying program syntax by writing $\mathbf{post}\ r \leftarrow p\ e\ d$. When task-passing is not allowed, region valuations need not store an entire configuration for each newly-posted task, since the posted task's initial region valuation is empty. As this represents a significant simplification on which our subsequent analysis results rely, we redefine here a few key notions.

3.1 Sequential Semantics without Task-Passing

A *region valuation* is a (non-nested) mapping $m : \mathsf{Regs} \to \mathbb{M}[\mathsf{Tasks}]$ from regions to multisets of tasks, a *frame* $\langle t, m\rangle$ is a task $t \in \mathsf{Tasks}$ paired with a region valuation m, and a *configuration* c is a sequence of frames representing a procedure activation stack. The transition relation $\rightarrow^{\text{rpp}}$ of Figures 2 and 5 implicitly include a transition $\langle t_1, m\rangle\ c \rightarrow^{\text{rpp}}_P \langle t_2, m\rangle\ c$ whenever $t_1 \rightarrow^{\text{seq}}_P t_2$. The definitions of *sequential execution*, *initial*, and *reachable* are nearly identical to their task-passing parallel and sequential counterparts. Since pending tasks need not store initial region-valuations in non-task-passing programs, this simpler semantics is equivalent to the previous stack-based semantics.

Lemma 2. *For all initial conditions ι non-task-passing programs P, the valuation ℓ is reachable in P from ι by a sequential execution*

POST
$$\frac{v \in e(T) \qquad m' = m \cup (r \mapsto \langle v, s_p, d\rangle)}{\langle T[\mathbf{post}\ r \leftarrow p\ e\ d], m\rangle\ c \xrightarrow[P]{\text{rpp}} \langle T[\mathbf{skip}], m'\rangle\ c}$$

∃WAIT
$$\frac{m = (r \mapsto t_2) \cup m'}{\langle T_1[\mathbf{ewait}\ r], m\rangle\ c \xrightarrow[P]{\text{rpp}} \langle t_2, \emptyset\rangle\ \langle T_1[\mathbf{skip}], m'\rangle\ c}$$

∀WAIT-NEXT
$$\frac{m = (r \mapsto t_2) \cup m'}{\langle T_1[\mathbf{await}\ r], m\rangle\ c \xrightarrow[P]{\text{rpp}} \langle t_2, \emptyset\rangle\ \langle T_1[\mathbf{skip};\ \ \mathbf{await}\ r], m'\rangle\ c}$$

∀WAIT-DONE
$$\frac{m(r) = \emptyset}{\langle T[\mathbf{await}\ r], m\rangle\ c \xrightarrow[P]{\text{rpp}} \langle T[\mathbf{skip}], m\rangle\ c}$$

RETURN
$$\frac{s \in \mathsf{rvh}(t_1)}{\langle t_1, m_1\rangle\ \langle T_2[\mathbf{skip}], m_2\rangle\ c \xrightarrow[P]{\text{rpp}} \langle T_2[s], m_1 \cup m_2\rangle\ c}$$

Figure 5. The stack-based transition relation for sequentially-executing recursively parallel programs without task-passing.

with task-passing if and only if ℓ is reachable in P from ι by a sequential execution without task-passing.

Even with this simplification, we do not presently know whether the state-reachability problem for (finite-value) recursively parallel programs is decidable in general. In the following sections, we identify several decidable, and in some cases tractable, restrictions to the program model which correspond to the concurrency mechanisms found in real-world parallel programming languages.

3.2 Recursive Vector Addition Systems with Zero-Test Edges

Fix $k \in \mathbb{N}$. A *recursive vector addition system (RVASS)* $\mathcal{A} = \langle Q, \delta\rangle$ *of dimension* k is a finite set Q of states, along with a finite set $\delta = \delta_1 \uplus \delta_2 \uplus \delta_3$ of transitions partitioned into *additive* transitions $\delta_1 \subset Q \times \mathbb{N}^k \times \mathbb{N}^k \times Q$, *recursive* transitions $\delta_2 \subseteq Q \times Q \times Q \times Q$, and *zero-test* transitions $\delta_3 \subseteq Q \times Q$. We write

$$\begin{aligned} q \xrightarrow{\vec{n}_1 \vec{n}_2} q' \quad &\text{when } \langle q, \vec{n}_1, \vec{n}_2, q'\rangle \in \delta_1, \text{ and} \\ q \xrightarrow{q_1 q_2} q' \quad &\text{when } \langle q, q_1, q_2, q'\rangle \in \delta_2. \\ q \hookrightarrow q' \quad &\text{when } \langle q, q'\rangle \in \delta_3. \end{aligned}$$

A *(non-recursive) vector addition system (with states) (VASS)* is a recursive vector addition system $\langle Q, \delta\rangle$ such that δ contains only additive transitions.

An *(RVASS) frame* $\langle q, \vec{n}\rangle$ is a state $q \in Q$ along with a vector $\vec{n} \in \mathbb{N}^k$, and an *(RVASS) configuration* $c \in (Q \times \mathbb{N}^k)^+$ is a non-empty sequence of frames representing a stack of non-recursive sub-computations. The transition relation $\rightarrow^{\text{rvas}}$ for recursive vector addition systems is defined in Figure 9. The ADDITIVE rule updates the top frame $\langle q, \vec{n}\rangle$ by subtracting the vector $\vec{n}_1$ from $\vec{n}$, adding the vector $\vec{n}_2$ to the result, and updating the control state to q'. The CALL rule pushes on the frame-stack a new frame $\langle q_1, \mathbf{0}\rangle$ from which the RETURN rule will eventually pop at some point when the control state is q_2; when this happens, the vector $\vec{n}_1$ of the popped frame is added to the vector $\vec{n}_2$ of the frame below. We describe an application of the CALL (resp., RETURN) rule as a *call (resp., return) transition*. Finally, the ZERO rule proceeds only when the top-most frame's vector equals $\mathbf{0}$.

An *execution of a RVASS $\mathcal{A}$ (from c_0 to c_j)* is a configuration sequence $c_0 c_1 \ldots c_j$ where $c_i \rightarrow^{\text{rvas}} c_{i+1}$ for $0 \leq i < j$. A configuration $\langle q, \vec{n}\rangle$ is called q_0*-initial* when $q = q_0$ and $\vec{n} = \mathbf{0}$, and a configuration c_f is called q_f*-final* when $c_f = \langle q_f, \vec{n}\rangle\ c$ for some configuration c and $\vec{n} \in \mathbb{N}^k$. We say a state q_f is *reachable in* $\mathcal{A}$ from q_0 when there exists an execution of $\mathcal{A}$ from some q_0-initial configuration c_0 to some q_f-final configuration c_f. The *state-reachability problem* for recursive vector addition systems is to determine whether a given state q is reachable from some q_0.

Recently Demri et al. [8] have proved that state-reachability in branching vector addition systems (BVAS)—a very similar formal model to which RVASS reduces—is in 2EXPTIME. This immediately gives us an upper-bound on computing state-reachability in RVASS without zero-test edges. Though state-reachability in non-recursive systems is EXPSPACE-complete [27, 30], for the moment, we do not know matching upper and lower bounds for RVASS.

Lemma 3. *The state-reachability problem for recursive (resp., non-recursive) vector addition systems without zero-test edges is EXPSPACE-hard, and in 2EXPTIME (resp., EXPSPACE).*

3.3 Encoding Finite-Data Programs without Task-Passing as Recursive Vector Addition Systems with Zero-Test Edges

When the value set Vals of a given program P is taken to be finite, the set Tasks also becomes finite since there are finitely many statements and return-value handlers occurring in P. As finite-domain multisets are equivalently encoded with a finite number of counters (i.e., one counter per element), we can encode each region valuation $m \in \mathsf{Regs} \rightarrow \mathbb{M}[\mathsf{Tasks}]$ by a vector $\vec{n} \in \mathbb{N}^k$ of counters, where $k = |\mathsf{Regs} \times \mathsf{Tasks}|$. To clarify the correspondence, we fix an enumeration $\mathsf{cn} : \mathsf{Regs} \times \mathsf{Tasks} \rightarrow \{1, \ldots, k\}$, and associate each region valuation m with a vector $\vec{n}$ such that for all $r \in \mathsf{Regs}$ and $t \in \mathsf{Tasks}$, $m(r)(t) = \vec{n}(\mathsf{cn}(r, t))$. Let $\vec{n}_i$ denote the unit vector of dimension i, i.e., $\vec{n}_i(i) = 1$ and $\vec{n}_i(j) = 0$ for $j \neq i$.

Given a finite-data recursively parallel program P without task-passing, we associate a corresponding recursive vector addition system $\mathcal{A}_P = \langle Q, \delta\rangle$. We define $Q \stackrel{\text{def}}{=} \mathsf{Tasks} \cup \mathsf{Tasks}^3$, and define δ formally in Figure 7. Intra-procedural transitions translate directly to additive transitions. The $\mathbf{call}$ statements are handled by recursive transitions between entry and exit points t_0 and t_f of the called procedure. The $\mathbf{post}$ statements are handled by additive transitions that increment the counter corresponding to a region-task pair. The $\mathbf{ewait}$ statements are handled in two steps: first an additive transition decrements the counter corresponding to region-task pair $\langle r, t_0\rangle$, then a recursive transition between entry and exit points t_0 and t_f of the corresponding procedure is made, applying the return-value handler of t_f upon the return. (Here we use an intermediate state $\langle T[\mathbf{skip}], t_0, t_f\rangle \in Q$ to connect the two transitions, in order to differentiate the intermediate steps of other $\mathbf{ewait}$ transitions.) The $\mathbf{await}$ statements are handled similarly, except the $\mathbf{await}$ statement must be repeated again upon the return. Finally, a zero-test transition allows $\mathcal{A}_P$ to eventually step past each $\mathbf{await}$ statement.

Notice that ignoring intermediate states $\langle t_1, t_2, t_3\rangle \in Q$, the frames $\langle t, \vec{n}\rangle$ of $\mathcal{A}_P$ correspond directly to frames $\langle t, m\rangle$ of the given program P, given the correspondence between vectors and region valuations. This correspondence between frames indeed extends to configurations, and ultimately to the state-reachability problems between $\mathcal{A}_P$ and P.

Lemma 4. *For all programs P without task-passing, procedures $p_0 \in \mathsf{Procs}$, and values $\ell_0, \ell \in \mathsf{Vals}$, ℓ is reachable from $\langle \ell_0, p_0\rangle$ in P if and only if there exist $s \in \mathsf{Stmts}$ and $d_0, d \in \mathsf{Rets}$ such that $\langle \ell, s, d\rangle$ is reachable from $\langle \ell_0, s_{p_0}, d_0\rangle$ in $\mathcal{A}_P$.*

Our analysis algorithms in the following sections use Lemma 4 to compute state-reachability of a program P without task-passing by computing state-reachability on the corresponding RVASS $\mathcal{A}_P$.

$$\textsc{Additive}\ \frac{q \xhookrightarrow{\vec{n}_1 \vec{n}_2} q' \qquad \vec{n} \geq \vec{n}_1}{\langle q, \vec{n}\rangle\, c \xrightarrow{\text{rvas}} \langle q', \vec{n} \ominus \vec{n}_1 \oplus \vec{n}_2\rangle\, c} \qquad \textsc{Call}\ \frac{q \xhookrightarrow{q_1 q_2} q'}{\langle q, \vec{n}\rangle\, c \xrightarrow{\text{rvas}} \langle q_1, \mathbf{0}\rangle\, \langle q, \vec{n}\rangle\, c}$$

$$\textsc{Return}\ \frac{q \xhookrightarrow{q_1 q_2} q'}{\langle q_2, \vec{n}_1\rangle\, \langle q, \vec{n}_2\rangle\, c \xrightarrow{\text{rvas}} \langle q', \vec{n}_1 \oplus \vec{n}_2\rangle\, c} \qquad \textsc{Zero}\ \frac{q \hookrightarrow q'}{\langle q, \mathbf{0}\rangle\, c \xrightarrow{\text{rvas}} \langle q', \mathbf{0}\rangle\, c}$$

Figure 6. The transition relation for recursive vector addition systems. To simplify presentation, we assume that there is at most one recursive transition originating from each state, i.e., for all $q \in Q$, $\left|\delta_2 \cap (\{q\} \times Q^3)\right| \leq 1$. We denote by $\mathbf{0}$ the vector $\langle 0, 0, \ldots, 0\rangle$, and by $\oplus$ and $\ominus$ the usual vector addition and subtraction operators.

$$\frac{v_0 \in e(T) \qquad i = \mathsf{cn}(r, \langle v_0, s_p, d\rangle)}{T[\mathbf{post}\ r \leftarrow p\ e\ d] \xrightarrow{\mathbf{0}\vec{n}_i} T[\mathbf{skip}]} \qquad \frac{}{T[\mathbf{await}\ r] \longhookrightarrow T[\mathbf{skip}]}$$

$$\frac{v_0 \in e(T) \qquad t_0 = \langle v_0, s_p, \mathtt{d}_{call}\rangle \qquad (\mathtt{l}\ \mathtt{:=}\ v_f) \in \mathsf{rvh}(t_f)}{T[\mathbf{call}\ \mathtt{l}\ \mathtt{:=}\ p\ e] \xrightarrow{t_0 t_f} T\big[\mathtt{l}\ \mathtt{:=}\ v_f\big]}$$

$$\frac{t_1 \xrightarrow[P]{\text{seq}} t_2}{t_1 \xrightarrow{\mathbf{00}} t_2} \qquad \frac{i = \mathsf{cn}(r, t_0) \qquad s \in \mathsf{rvh}(t_f)}{T[\mathbf{ewait}\ r] \xrightarrow{\vec{n}_i \mathbf{0}} \big\langle T[\mathbf{skip}], t_0, t_f\big\rangle \xrightarrow{t_0 t_f} T[s]}$$

$$\frac{i = \mathsf{cn}(r, t_0) \qquad s \in \mathsf{rvh}(t_f)}{T[\mathbf{await}\ r] \xrightarrow{\vec{n}_i \mathbf{0}} \big\langle T[\mathbf{skip}], t_0, t_f\big\rangle \xrightarrow{t_0 t_f} T[s;\ \mathbf{await}\ r]}$$

Figure 7. The transitions of the RVASS $\mathcal{A}_P$ encoding the behavior of a finite-data recursively parallel program P.

In general, our algorithms compute sets of region valuation vectors

$$\mathsf{sms}(t_0, t_f, P) \stackrel{\text{def}}{=} \{\vec{n} : \langle t_0, \mathbf{0}\rangle \xrightarrow[\mathcal{A}_P]{\text{rvas}} * \langle t_f, \vec{n}\rangle\},$$

summarizing the execution of a procedure between an entry point t_0 and exit point t_f, where we write $\rightarrow^{\text{rvas}}_{\mathcal{A}_P} *$ to denote zero or more applications of $\rightarrow^{\text{rvas}}_{\mathcal{A}_P}$. Given an effective way to compute such a function, we could systematically replace inter-procedural program steps (i.e., of the **call**, **ewait**, and **await** statements) with intra-procedural edges performing their net effect. Note however that even if the set of tasks is finite, the set $\mathsf{sms}(t_0, t_f, \mathcal{A}_P)$ of summaries between t_0 and t_f need not be finite; the ability to compute this set is thus the key to our summarization-based algorithms in the following sections.

4. Single-Wait Programs

Definition 2 (Single wait)**.** A *single-wait program* is a program which does not contain the **await** statement.

Single-wait programs can wait only for a single pending task at any program point. Many parallel programming constructs can be modeled as single-wait programs.

4.1 Parallel Programs with Futures

The **future** annotation of Multilisp [18] has become a widely adopted parallel programming construct, included, for example, in X10 [6] and in Leijen et al. [26]'s task parallel library. Flanagan and Felleisen [12] provide a principled description of its semantics. The future construct leverages the procedural program structure for parallelism, essentially adding a "lazy" procedure call which immediately returns control to the caller with a placeholder for a value that may not yet have been computed, along with an operation for ensuring that a given placeholder has been filled in with a computed value. Syntactically, futures add two statements,

$$\mathbf{future}\ x\ \mathtt{:=}\ p\ e \qquad \mathbf{touch}\ x,$$

where x ranges over program variables, $p \in \mathsf{Procs}$, and $e \in \mathsf{Exprs}$. Though it is not necessarily present in the syntax of a source language with futures, we assume every use of a variable assigned by a **future** statement is explicitly preceded by a **touch** statement. Semantically, the **future** statement creates a new process in which to execute the given procedure, which proceeds to execute in parallel with the caller—and all other processes created in this way. The **touch** statement on a variable x blocks execution of the current procedure until the future procedure call which assigned to x completes, returning a value with which is copied into x. Even though each procedure can only spawn a bounded number of parallel processes—i.e., one per program variable—there is in general no bound on the total number of parallelly-executing processes, since procedure calls—even parallel ones—are recursive.

Example 2. The Fibonacci function can be implemented as a parallel algorithm using futures as follows.

```
proc fib (var n: ℕ)
    var x, y: ℕ
    if n < 2 then
        return 1
    else
        future x := fib (n-1);
        future y := fib (n-2);
        touch x;
        touch y;
        return x + y
```

As opposed to the usual (naïve) sequential implementation operating in time $\mathcal{O}(n^2)$, this parallel implementation runs in time $\mathcal{O}(n)$.

The semantics of futures is readily expressed with handle-passing programs using the **post** and **ewait** statements. Assuming a region identifier $\mathtt{r}_x$ and return handler $\mathtt{d}_x$ for each program variable x, we encode

$\mathbf{future}\ x\ \mathtt{:=}\ p\ e$	as	$\mathbf{post}\ \mathtt{r}_x \leftarrow p\ e\ \vec{\mathtt{r}}\ \mathtt{d}_x$
$\mathbf{touch}\ x$	as	$\mathbf{ewait}\ \mathtt{r}_x$

where $\mathtt{d}_x(v) \stackrel{\text{def}}{=} x\ \mathtt{:=}\ v$ simply assigns the return value v to the variable x, and the vector $\vec{\mathtt{r}}$ contains each $\mathtt{r}_y$ such that the variable y appears in e.

4.2 Parallel Programs with Revisions

Burckhardt et al. [5]'s revisions model of concurrent programming proposes a mechanism analogous to (software) version control systems such as CVS and subversion, which promises to naturally and easily parallize sequential code in order to take advantage of multiple computing cores. There, each sequentially executing process is referred to as a *revision*. A revision can branch into two revisions, each continuing to execute in parallel on their own separate copies of data, or merge a previously-created revision, provided a programmer-defined *merge function* to mitigate the updates to data which each have performed. Syntactically, revisions add two statements,

$$x\ \mathtt{:=}\ \mathbf{rfork}\ s \qquad \mathbf{join}\ x,$$

where x ranges over program variables, and $s \in \mathsf{Stmts}$. Semantically, the **rfork** statement creates a new process to execute the given statement, which proceeds to execute in parallel with the invoker—and all other processes created in this way. The assignment stores a *handle* to the newly-created revision in a *revision variable* x. The **join** statement on a revision variable x blocks execution of the current revision until the revision whose handle is stored in x completes; at that point the current revision's data is updated according to a programmer-supplied merge function $m : (\mathsf{Vals} \times \mathsf{Vals} \times \mathsf{Vals}) \to \mathsf{Vals}$: when v_0, v_1 are, resp., the initial and final data values of the merged revision, and v_2 is the current data value of the current revision, the current revisions data value is updated to $m(v_0, v_1, v_2)$.

The semantics of revisions is readily expressed with handle-passing programs using the **post** and **ewait** statements. Assuming a region identifier $\mathtt{r}_x$ for each program variable x, and a programmer-supplied merge function m, we encode

x `:=` **rfork** s	as	**post** $\mathtt{r}_x \leftarrow p_s$ `l` $\vec{\mathtt{r}}$ `d`
join x	as	**ewait** $\mathtt{r}_x$

where p_s is a procedure declared as

proc p_s (**var** `l`: T)
  var $\mathtt{l}_0$ `:= l`
  s`;`
  return $(\mathtt{l}_0$`,l)`

and $\mathtt{d}(\langle v_0, v_1\rangle) \stackrel{\text{def}}{=}$ `l := m(`v_0`,l,`v_1`)` updates the current local valuation based on the joined revision's initial and final valuations $v_0, v_1 \in \mathsf{Vals}$, and the joining revision's current local valuation stored in `l`. The vector $\vec{\mathtt{r}}$ contains each $\mathtt{r}_y$ for which the revision variable y is accessed in s.[2]

4.3 Parallel Programs with Asynchronous Procedures

The model of so-called *asynchronous programs* [14, 20, 35] is of particular interest for modeling reactive systems, such as device drivers, web servers, and graphical user interfaces, with low-latency requirements. Essentially, a program is made up of a collection of short-lived tasks running one-by-one and accessing a global store, which post other tasks to be run at some later time. Tasks are initially posted by an initial procedure, and may also be generated by external system events. An *event loop* repeatedly chooses a pending task from its collection to execute to completion, adding the tasks it posts back to the task collection. Syntactically, asynchronous programs add two statements,

async $p\ e$ **eventloop**

such that **eventloop** is invoked only once as the last statement of the initial procedure. Semantically, the **async** statement initializes a procedure call and returns control immediately, without waiting for the call to return. The **eventloop** statement repeatedly dispatches pending—i.e., called but not yet returned—procedures, and executing them to completion; each procedure can make both synchronous calls, as well as additional asynchronous procedure calls. It is assumed that each procedure executes atomically, i.e., one-by-one, without interference from other pending procedures. Each procedure can call an unbounded number of asynchronous procedures, and the order in which procedure calls are dispatched is chosen non-deterministically.

The semantics of asynchronous programs is readily expressed with (non-deterministic) recursively parallel programs using the **post** and **ewait** statements. Assuming a single region identifier $\mathtt{r}_0$, we encode

async $p\ e$	as	**post** $\mathtt{r}_0 \leftarrow p'\ e$ `d`
eventloop	as	**while** `true` **do ewait** $\mathtt{r}_0$

where given that p has top-level statement s which accesses a shared global variable `g`, in addition to the procedure parameter `l`, p' is declared as

proc p' (**var** `l`: T)
  var $\mathtt{g}_0$ `:=` $\star$
  var `g :=` $\mathtt{g}_0$
  s`;`
  return $(\mathtt{g}_0$`,g)`

and $\mathtt{d}(\langle v_0, v_1\rangle) \stackrel{\text{def}}{=}$ **assume** `l =` v_0`;` `l :=` v_1 models the atomic update p performs from an initial (guessed) shared global valuation v_0. This guessing allows us to simulate the communication of a shared global state `g`, which is later ensured to have begun at the value v_0 which the previously-executed asynchronous task had written.

5. Single-Wait Analysis

The absence of **await** edges in a program P implies the absence of zero-test transitions in the corresponding recursive vector addition system $\mathcal{A}_P$. To compute state-reachability in P via procedure summarization, we must summarize the recursive transitions of $\mathcal{A}_P$ by additive transitions (in a non-recursive system) accounting for the left-over pending tasks returned by reach procedure. This is not trivial in general, since the space of possibly returned region valuations is infinite. In increasing difficulty, we isolate three special cases of single-wait programs, whose analysis problems are simpler than the general case. In the simplest "non-aliasing" case where the number of tasks stored in each region of a procedure frame is limited to one, the execution of **ewait** statements are deterministic. When the number of tasks stored in each region is not limited to one, non-determinism arises from the choice of which completed task to pick at each **ewait** statement (see the ∃WAIT rule of Figure 5). This added power makes the state-reachability problem at least as hard as state-reachability in vector addition systems—i.e., EXPSPACE-hard, though the precise complexity depends on the scope of pending tasks. After examining the PTIME-complete non-aliasing case, we examine two EXPSPACE-complete cases by restricting the scope of task handles, before moving to the general case.

5.1 Single-Wait Analysis without Aliasing

Many parallel programming languages consume only the computations of precisely-addressed tasks. In futures, for example, the **touch** x statement applies to the return value of a particular procedure—the last one whose future result was assigned to x. Similarly, in revisions, the **join** x statement applies to the last revision whose handle was stored in x. Indeed in the single-wait program semantics of each case, we are guaranteed that the corresponding region, $\mathtt{r}_x$, contains at most one task handle. Thus the non-determinism arising (from choosing between tasks in a given region) in the ∃WAIT rule of Figure 3 disappears.

Definition 3 (Non aliasing)**.** We say a region $r \in \mathsf{Regs}$ is *aliased* in a region valuation $m : \mathsf{Regs} \to \mathbb{M}[\mathsf{Tasks}]$ when $|m(r)| > 1$. We say r is *aliasing* in a program P if there exists a reachable configuration $C[\langle t, m\rangle]$ of P in which r is aliased in m. A *non-aliasing program* is a program in which no region is aliasing.

Note that the set of non-aliasing region valuations is finite when the number of program values is. The non-aliasing restriction thus allow us immediately to reduce the state-reachability problem for single-wait programs to reachability in a recursive finite-data sequential program. To compute state-reachability we consider a sequence

[2] Actually $\vec{\mathtt{r}}$ must in general be chosen non-deterministically, as each revision handle may be joined either by the parent revision or its branch.

$\mathcal{A}_0\mathcal{A}_1\ldots$ of finite-state systems iteratively under-approximating the recursive system $\mathcal{A}_P$ given from a single-wait program P. Initially, $\mathcal{A}_0$ has only the transitions of $\mathcal{A}_P$ corresponding to intra-procedural and **post** transitions of P. At each step $i > 0$, we add to $\mathcal{A}_i$ an additive edge summarizing an **ewait** transition

$$T[\mathbf{ewait}\ r] \xhookrightarrow{\vec{n}_j\vec{n}} T[s]\,,$$

for some $t_0, t_f \in \mathsf{Tasks}$ such that $j = \mathsf{cn}(r, t_0)$, $s \in \mathsf{rvh}(t_f)$, and $\vec{n}$ is reachable at t_f from t_0 in $\mathcal{A}_{i-1}$, i.e., $\vec{n} \in \mathsf{sms}(t_0, t_f, \mathcal{A}_{i-1})$. This $\mathcal{A}_0\mathcal{A}_1\ldots$ sequence is guaranteed to reach a fixed-point $\mathcal{A}_k$, since the set of non-aliasing region valuation vectors, and thus the number of possibly added edges, is finite. Furthermore, as each $\mathcal{A}_i$ is finite-state, only finite-state reachability queries are needed to determine the reachable states of $\mathcal{A}_k$, which are precisely the same reachable states of $\mathcal{A}_P$. Note that the number of region valuations grows exponentially in the number of regions.

Theorem 2. *The state-reachability problem for non-aliasing single-wait finite-value finite-region programs is PTIME-complete for a fixed number of regions, and EXPTIME-complete in the number of regions.*

Note that these results apply to non-aliasing single-wait programs *without* task-passing (see Section **??**). Although futures (Section 4.1) and revisions (Section 4.2) are both captured by non-aliasing single-wait programs, they are not necessarily free of task passing.

5.2 Local-Scope Single-Wait Analysis with Aliasing

Definition 4 (Local scope)**.** A *local-scope program* is a program in which tasks only return with empty region valuations; i.e., for all reachable configurations $C[\langle t[\mathbf{return}\ e]\,, m\rangle]$ we have $m = m_\emptyset$.

To solve state-reachability in local-scope single-wait programs, we compute a sequence $\mathcal{A}_0\mathcal{A}_1\ldots$ of non-recursive vector addition systems iteratively under-approximating the recursive system $\mathcal{A}_P$ arising from a program P. The initial system $\mathcal{A}_0$ has only the transitions of $\mathcal{A}_P$ corresponding to intra-procedural and **post** transitions of P, and at each step $i > 0$, we add to $\mathcal{A}_i$ an additive edge summarizing an **ewait** transition

$$T[\mathbf{ewait}\ r] \xhookrightarrow{\vec{n}_j\mathbf{0}} T[s]$$

for some $t_0, t_f \in \mathsf{Tasks}$ such that $j = \mathsf{cn}(r, t_0)$, $s \in \mathsf{rvh}(t_f)$, and $\vec{n} \in \mathsf{sms}(t_0, t_f, \mathcal{A}_{i-1})$. As P is a local-scope program, every such $\vec{n}$ must be equal to $\mathbf{0}$. Since the number of possibly added edges is polynomial in P, the $\mathcal{A}_0\mathcal{A}_1$ sequence is guaranteed to reach in a polynomial number of steps a fixed-point $\mathcal{A}_k$ whose reachable states are exactly those of $\mathcal{A}_P$. The entire procedure is EXPSPACE-complete, since each procedure-summarization reachability query is equivalent to computing state-reachability in vector addition systems.

Theorem 3. *The state-reachability problem for local-scope single-wait finite-value finite-region programs is EXPSPACE-complete.*

5.3 Global-Scope Single-Wait Analysis with Aliasing

Another relatively simple case of interest is when pending tasks are allowed to leave the scope in which they are posted, but can only be consumed by a particular, statically declared, task in an enclosing scope. This is the case, for example, in asynchronous programs [35], though here we allow for slightly more generality, since tasks can be posted to multiple regions, and arbitrary control in the initial procedure frame is allowed.

Definition 5 (Global scope)**.** A *global-scope programs* is a program in which the **ewait** (and **await**) statements are used only in the initial procedure frame.

The global-scope restriction also makes state-reachability in finite-value single-wait programs equivalent to the state-reachability problem for non-recursive vector addition systems, albeit through a completely different reduction.

Since each non-initial procedure p cannot consume tasks, the set of posted tasks along any execution of p from t_0 to t_f, including tasks posted by procedures p has called recursively, is a semi-linear set, described by the Parikh image of a context-free language. Following the approach of Ganty and Majumdar [14], for each $t_0, t_f \in \mathsf{Tasks}$ we construct a polynomial-sized vector addition system $\mathcal{A}(t_0, t_f)$ characterizing this semi-linear set of tasks (recursively) posted between t_0 and t_f. Then, we use each $\mathcal{A}(t_0, t_f)$ as a component of a non-recursive vector addition system $\mathcal{A}'_P$ representing execution of the initial frame. In particular, $\mathcal{A}'_P$ contains transitions to and from the component $\mathcal{A}(t_0, t_f)$ for each $t_0, t_f \in \mathsf{Tasks}$,

$$T[\mathbf{ewait}\ r] \xhookrightarrow{\vec{n}_j\mathbf{0}} \langle q_0, T[\mathbf{skip}]\rangle \quad \langle q_f, T[\mathbf{skip}]\rangle \xhookrightarrow{\mathbf{00}} T[s]\,,$$

for all $r \in \mathsf{Regs}$ such that $j = \mathsf{cn}(r, t_0)$, $s \in \mathsf{rvh}(t_f)$, and q_0 and q_f are the initial and final states of $\mathcal{A}(t_0, t_f)$. We assume each $\mathcal{A}(t_0, t_f)$ has unique initial and final states, distinct from the states of other components $\mathcal{A}(t'_0, t'_f)$. In order to transition to the correct state $T[s]$ upon completion, $\mathcal{A}(t_0, t_f)$ carries an auxiliary state-component $T[\mathbf{skip}]$. In this way, for each task t' posted to region r' in an execution between t_0 and t_f, the component $\mathcal{A}(t_0, t_f)$ does the incrementing of the $\mathsf{cn}(r', t')$-component of the region-valuation vector. As each of the polynomially-many components $\mathcal{A}(t_0, t_f)$ are constructed in polynomial time [14], this method constructs $\mathcal{A}'_P$ in polynomial time. Thus state-reachability in P is computed by state-reachability in the non-recursive vector addition system $\mathcal{A}'_P$, in exponential space. The EXPSPACE complexity of this algorithm is asymptotically optimal since global-scope single-wait programs are powerful enough to capture state-reachability in vector addition systems.

Theorem 4. *The state-reachability problem for global-scope single-wait finite-value finite-region programs is EXPSPACE-complete.*

The proof is a generalization of Ganty and Majumdar [14]'s proof of EXPSPACE-completeness for asynchronous programs.

5.4 Single-Wait Analysis with Aliasing

In general, the state-reachability problem for finite-value single-wait programs is as hard as state-reachability in recursive vector addition systems without zero-test edges.

Theorem 5. *The state-reachability problem for single-wait finite-value finite-region programs is EXPSPACE-hard, and in* 2*EXPTIME.*

Demri et al. [8]'s proof of membership in 2EXPTIME relies on a non-deterministically chosen reachability witness without materializing a practical algorithm for the search of said witness. Here we give a summarization-based algorithm for computing state-reachability.

To compute the set of reachable states, we consider again a sequence $\mathcal{A}_0\mathcal{A}_1\ldots$ of non-recursive vector addition systems successively under-approximating the recursive system $\mathcal{A}_P$ arising from a single-wait program P. Initially, $\mathcal{A}_0$ has only the transitions of $\mathcal{A}_P$ corresponding to intra-procedural and **post** transitions of P. At each step $i > 0$, we add to $\mathcal{A}_i$ an additive edge summarizing an **ewait** transition

$$T[\mathbf{ewait}\ r] \xhookrightarrow{\vec{n}_j\vec{n}} T[s]\,,$$

for some $t_0, t_f \in \mathsf{Tasks}$ such that $j = \mathsf{cn}(r, t_0)$, $s \in \mathsf{rvh}(t_f)$, and $\vec{n} \in \mathsf{sms}(t_0, t_f, \mathcal{A}_{i-1})$. Even though the set of possible added additive edges summarizing recursive transitions is infinite, with careful analysis we can show that this very simple algorithm

terminates, provided we can bound the edge-labels $\vec{n}$ needed to compute state-reachability in $\mathcal{A}_P$. It turns out we can bound these edge labels, by realizing that the minimal vectors required to reach a target state from any given program location are bounded.

First, in order to reason backward about executions to the target state, consider the non-recursive system $\mathcal{A}'_i$ obtained by adding "return" transitions $t_f \overset{\mathbf{oo}}{\hookrightarrow} T[s]$ from every procedure exit point $t_f = T_f[\mathbf{return}\ e]$ and procedure return point $T[\mathbf{ewait}\ e]$ occurring in P such that $s \in \mathsf{rvh}(t_f)$. These extra transitions in $\mathcal{A}'_i$ simulate a return from t_f to t, transferring all of the pending tasks from a frame at t_f to a frame at $T[s]$, without any contribution from the $T[s]$'s intra-procedural predecessor $T[\mathbf{ewait}\ e]$.

Then define a sequence of functions $\eta_0, \eta_1, \ldots : \mathsf{Tasks} \to \wp(\mathbb{N}^k)$, each η_i mapping each $t \in \mathsf{Tasks}$ to the (possibly empty, upward-closed) set of vectors $\eta_i(t)$ such that for any $\vec{n} \in \eta_i(t)$, a configuration $\langle t, \vec{n}\rangle$ is guaranteed to reach the target reachable state in $\mathcal{A}'_i$—and thus $\langle t, \vec{n}\rangle\, c$ is guaranteed to reach the target reachable state in $\mathcal{A}_P$; each η_i can be computed in exponential space, by backward reachability in the non-recursive vector addition system $\mathcal{A}'_i$ [30]. Since each $\mathcal{A}_i$ contains at least the transitions of $\mathcal{A}_{i-1}$, the η_i-sequence non-decreasing; i.e., more and more configurations can reach the target state; i.e., for all $t \in \mathsf{Tasks}$ we have $\eta_{i-1}(t) \subseteq \eta_i(t)$. Since there can be no ever-increasing sequence of upward-closed sets of vectors over natural numbers, the η_i sequence must stabilize after a finite number of steps.

Furthermore, since any $\vec{n} \in \eta_i(t)$ is guaranteed to reach the target state, it suffices to consider only vectors $\vec{n}'$ bounded by the minimals of the upward-closed set $\eta_i(t)$. To see why, notice that if some $\vec{n} \in \eta_i(t)$ labels an edge between t_0 and t, then *every* configuration at t_0 is guaranteed to reach the target state, since this edge adds the vector guaranteed to reach the target from t. Additionally, any vector grater than a minimal of $\eta_i(t)$ is already guaranteed to be present in $\eta_i(t)$, since $\eta_i(t)$ is upward closed. Thus we need only consider edge-labels bounded by the decreasing $\eta_0\eta_1 \ldots$ sequence.

Thus since the $\mathcal{A}_0\mathcal{A}_1 \ldots$ sequence stabilizes after a finite number of steps, and each step can be computed by backward-reachability in a (non-recursive) vector addition system, our summarization algorithm is guaranteed to terminate.

6. Multi-Wait Programs

Though single-wait programs capture many parallel programming constructs, they can not express waiting for each and every of an unbounded number of tasks to complete. Some programming languages require this dual notion, expressed here with **await**.

Definition 6 (Multi wait)**.** A *multi-wait program* is a program which does not contain the **ewait** statement.

Thus, multi-wait programs can wait only on every pending task (in a given region) at any program point. Many parallel programming constructs can be modeled as multi-wait programs.

6.1 Parallel Programming in Cilk

The Cilk parallel programming language [32] is an industrial-strength language with an accompanying runtime system which is used in a spectrum of environments, from modest multi-core computations to massively parallel computations with supercomputers. Similarly to futures (see Section 4.1), Cilk adds a form of procedure call which immediately returns control to the caller. Instead of an operation to synchronize with a *particular* previously-called procedure, Cilk only provides an operation to synchronize with *every* previously-called procedure. At such a point, the previously-called procedures communicate their results back to the caller one-by-one with atomically-executing procedure in-lined in scope of the caller.

Syntactically, Cilk adds two statements

$$\mathbf{spawn}\ p\ e\ p' \qquad \mathbf{sync},$$

where p ranges over procedures, e over expressions, and p' over procedures declared by

$$\mathbf{inlet}\ p'\ \texttt{(var rv:}\ T\texttt{)}\ s.$$

Here s ranges over intra-procedural program statements containing two variables: `rv`, corresponding to the value returned from a spawned procedure, and `l`, corresponding to the local variable of the spawning procedure. Semantically, the **spawn** statement creates a new process in which to execute the given procedure, which proceeds to execute in parallel with the caller—and all other processes created in this way. The **sync** statement blocks execution of the current procedure until each spawned procedure completes, and executes its associated inlet. The inlets of each procedure execute atomically. Each procedure can spawn an unbounded number of parallel processes, and the order in which the inlets of procedures execute is chosen non-deterministically.

Example 3. The Fibonacci function can be implemented as a parallel algorithm using Cilk as follows.

```
proc fib (var n: ℕ)
    var sum: ℕ
    if n < 2 then
        return 1
    else
        spawn fib (n-1) summer;
        spawn fib (n-2) summer;
        sync;
        return sum

inlet summer (var i: ℕ)
    sum := sum + i
```

As opposed to the usual (naïve) sequential implementation operating in time $\mathcal{O}(n^2)$, this parallel implementation runs in time $\mathcal{O}(n)$.

The semantics of Cilk is ready expressed with recursively parallel programs using the **post** and **await** statements. Assuming a region identifier $\mathtt{r}_0$, we encode

$\mathbf{spawn}\ p\ e\ p'$	as	$\mathbf{post}\ \mathtt{r}_0 \leftarrow p\ e\ \mathtt{d}_{p'}$
$\mathbf{sync}$	as	$\mathbf{await}\ \mathtt{r}_0$

where $\mathtt{d}_{p'}(v) \overset{\text{def}}{=} s_{p'}[v/\mathtt{rv}]$ executes the top-level statement of the inlet p' with input parameter v.

6.2 Parallel Programming with Asynchronous Statements

The async/finish pair of constructs in X10 [6] introduces parallelism through asynchronously executing statements and synchronization blocks. Essentially, an asynchronous statement immediately passes control to a following statement, executing itself in parallel. A synchronization block executes as any other program block, but does not pass control to the following statements/block until every asynchronous statement within has completed. Syntactically, this mechanism is expressed with two statements,

$$\mathbf{async}\ s \qquad \mathbf{finish}\ s$$

where s ranges over program statements. Semantically, the **async** statement creates a new process to execute the given statement, which proceeds to execute in parallel with the invoker—and all other processes created in this way. The **finish** statement executes the given statement s, then blocks execution until every process created within s has completed.

Example 4. The Fibonacci function can be implemented as a parallel algorithm using asynchronous statements as follows.

```
proc fib (var n: ℕ)
    var x, y: ℕ
    if n < 2 then
        return 1
    else
        finish
            async call x := fib (n-1);
            async call y := fib (n-2);
        return x + y
```

As opposed to the usual (naïve) sequential implementation operating in time $\mathcal{O}(n^2)$, this parallel implementation runs in time $\mathcal{O}(n)$.

Asynchronous statements are readily expressed with (non-deterministic) recursively parallel programs using the **post** and **await** statements. Let N be the maximum depth of nested **finish** statements. Assuming region identifiers $\mathtt{r}_1, \ldots, \mathtt{r}_N$, we encode

$$\textbf{async}\ s \quad \text{as} \quad \textbf{post}\ \mathtt{r}_i \leftarrow p_s\ \star\ \mathtt{d}$$

$$\textbf{finish}\ s \quad \text{as} \quad \textbf{await}\ \mathtt{r}_i$$

where $i-1$ is number of enclosing **finish** statements, and p_s is a procedure declared as

```
proc p_s (var l: T)
    var l_0 := l
    s;
    return (l_0,l)
```

and $\mathtt{d}(\langle v_0, v_1\rangle) \stackrel{\text{def}}{=} \textbf{assume}\ \mathtt{l} = v_0;\ \mathtt{l} := v_1$ models the update p performs from an initial (guessed) local valuation v_0. Using the same trick we have used to model asynchronous programs in Section 4.3, we model the sequencing of asynchronous tasks by initially guessing the value v_0 which the previously-executed asynchronous tasks had written, and validating that value when the return-value handler of a given task is finally run. Note that although X10 allows, in general, asynchronous tasks to interleave their memory accesses, our model captures only non-interfering tasks, by assuming either data-parallelism (i.e., disjoint accesses to data), or by assuming tasks are properly synchronized to ensure atomicity.

6.3 Structured Parallel Programming

So-called structured parallel constructs are becoming a standard parallel programming feature, adopted, for instance, in X10 [6] and in Leijen et al. [26]'s task parallel library. These constructs leverage normally sequential control structures to express parallelism. A typical syntactic instance of this is the parallel for-each loop:

$$\textbf{foreach}\ x\ \textbf{in}\ e\ \textbf{do}\ s$$

where x ranges over program variables, e over expressions, and s over statements. Semantically, the **foreach** statement creates a collection of new processes in which to execute the given statement—one for each valuation of the loop variable. After creating these processes, the **foreach** statement then block execution, waiting for each to complete.

The semantics of the for-each loop is readily expressed with recursively parallel programs using the **post** and **await** statements. With a region identifier $\mathtt{r}_0$, we encode $\textbf{foreach}\ x\ \textbf{in}\ e\ \textbf{do}\ s$ as

```
for x in e do post r_0 ← p_s (x,⋆) d;
await r_0
```

and given that both x and $\mathtt{l}$ are free variables in s, p_s is a procedure declared as

```
proc p_s (var x: T, l: T)
    var l_0 := l
    s;
    return (l_0,l)
```

and $\mathtt{d}(\langle v_0, v_1\rangle) \stackrel{\text{def}}{=} \textbf{assume}\ \mathtt{l} = v_0;\ \mathtt{l} := v_1$ models the update p performs from an initial (guessed) local valuation v_0.

7. Multi-Wait Analysis

The presence of **await** edges implies the presence of zero-test transitions in the recursive vector addition system $\mathcal{A}_P$ corresponding to a multi-wait program P. As we have done for single-wait programs, we first examine the easier sub-case of local-scope programs, which in the multi-wait setting corresponds concurrency in the Cilk [32] language (modulo task interleaving), as well as structured parallel programming constructs such as the **foreach** parallel loop in X10 [6] and in Leijen et al. [26]'s task parallel library (see Section 6.3). The concurrent behavior of the asynchronous statements (Section 6.2) in X10 [6] does not satisfy the local-scope restriction, since **async** statements can include recursive procedure calls which are nested without interpolating **finish** statements. There computing state-reachability is equivalent to determining whether a particular vector is reachable in a non-recursive vector addition system—a decidable problem which is known to be EXPSPACE-hard, but for which the only known algorithms are non-primitive recursive. encountered use only a single-region, we restrict our attention at present to single-region multi-wait programs.

7.1 Local-Scope Single-Region Multi-Wait Analysis

With the local-scoping restriction, executions of each procedure $p \in \mathsf{Procs}$ between entry point $t_0 \in \mathsf{Tasks}$ and exit point $t_f \in \mathsf{Tasks}$ are completely summarized by a Boolean indicating whether or not t_f is reachable from t_0. However, as executions of p may encounter **await** statements, modeled by zero-test edges in the recursive vector addition system $\mathcal{A}_P$, computing this Boolean requires determining the reachable program valuations between each pair of consecutive "synchronization points" (i.e., occurrences of the **await** statement), which in principle requires deciding whether the vector $\mathbf{0}$ is reachable in a vector addition system describing execution from the program point just after the first **await** statement to the point just after the second; i.e., when $T_1[\textbf{await}\ r]$ and $T_2[\textbf{await}\ r]$ are consecutively-occurring synchronization points, we must determine whether $\langle T_1[\textbf{skip}], \mathbf{0}\rangle$ can reach $\langle T_2[\textbf{skip}], \mathbf{0}\rangle$.

A careful analysis of our reachability problem reveals it does not have the EXPSPACE-hard complexity of determining vector-reachability in general, due to the special structure of our reachability query. We notice that between two synchronization points t_1 and t_2 of p, execution proceeds in two phases. In the first, **post** statements made by p only increment the vector valuations. In the second phase, starting when the second **await** statement is encountered, the **await** statement repeatedly consumes tasks, only decrementing the vector valuations—the vector valuations can not be re-incremented again because of the local-scope restriction: each consumed task is forbidden from returning addition tasks. Due to this special structure, deciding reachability between t_1 and t_2 reduces to deciding if a particular integer linear program $I(t_1, t_2)$ has a solution.

Since consuming tasks in the **await**-loop requires using the summaries computed for other procedures, we consider a sequence $\mathcal{A}_0\mathcal{A}_1\ldots$ of non-recursive vector addition systems iteratively under-approximating the recursive system $\mathcal{A}_P$. Initially $\mathcal{A}_0$ has only the transitions of $\mathcal{A}_P$ corresponding to intra-procedural and **post** transitions of P. At each step $i > 0$, we add to $\mathcal{A}_i$ one of two edges types. One type is an additive procedure-summary edge, used to describe a single task-consumption step of an **await** transition,

$$T[\textbf{await}\ r] \xrightarrow{\vec{n}_j \mathbf{0}} T[s;\ \textbf{await}\ r],$$

for some $t_0, t_f \in \mathsf{Tasks}$ such that $j = \mathsf{cn}(r, t_0)$, $s \in \mathsf{rvh}(t_f)$, and $\mathsf{sms}(t_0, t_f, \mathcal{A}_{i-1}) \neq \emptyset$. The second possibility is an additive synchronization-point summary edge, summarizing an entire of se-

quence of program transitions between two synchronization points,

$$T_1[\mathbf{skip}] \xhookrightarrow{\mathbf{oo}} T_2[\mathbf{skip}]\,,$$

where $T_1[\mathbf{await}\ r]\,, T_2[\mathbf{await}\ r] \in \mathsf{Tasks}$ are consecutive synchronization points occurring P, and $\mathbf{0} \in \mathsf{sms}(T_1[\mathbf{skip}]\,, T_2[\mathbf{skip}]\,, \mathcal{A}_P)$. The procedure-summary edges are computed using only finite-state reachability between program states, using the synchronization-point summary edges, while the synchronization-point summary edges are computed by reduction to integer linear programming. As the number of possible edges is bounded polynomially in the program size, the $\mathcal{A}_0\mathcal{A}_1$ sequence is guaranteed to reach a fixed-point $\mathcal{A}_k$ in a polynomial number of steps, though each step may take nondeterministic-polynomial time, in the worst case, due to computing solutions to integer linear programs. Furthermore, the reachable states of $\mathcal{A}_k$ are precisely the same reachable states of $\mathcal{A}_P$.

Theorem 6. *The state-reachability problem for local-scope multi-wait single-region finite-value programs is NP-complete.*

7.2 Single-Region Multi-Wait Analysis

Without the local-scoping restriction, each execution of each procedure $p \in \mathsf{Procs}$ between entry point $t_0 \in \mathsf{Tasks}$ and exit point $t_f \in \mathsf{Tasks}$ is summarized by the tasks posted between the last-encountered **await** statement, at a "synchronization point" $t_s \in \mathsf{Tasks}$ (note that $t_s = t_0$ if no **await** statements are encountered), and a **return** statement, at the exit point t_f. Since p can make recursive procedure calls between t_s and t_f, and each called procedure can again return pending tasks, the possible sets of pending tasks upon p's return at t_f is described by the Parikh image of a context-free language $L(t_0, t_f)$. It turns out we can describe $\Pi(L(t_0, t_f))$ as the set of vectors computed by a polynomially-sized vector addition system $\mathcal{A}^L(t_0, t_f)$ without recursion and zero-test edges [14]. We use thus computations of $\mathcal{A}^L(t_0, t_f)$ to summarize the set of possible region-valuations reached in an execution from t_0 to t_f. However, computing $\mathcal{A}^L(t_0, t_f)$ is not immediate, since between t_0 and the last-encountered synchronization point t_s, execution of the given procedure p may encounter **await** statements (necessarily so when $t_0 \neq t_s$). Since we use zero-test edges to express **await** statements, we also need to summarize execution between synchronization points (i.e., between the procedure entry point and among **await** statements) using only additive edges. To further complicate matters, each such summarization requires, in turn, the summaries $\mathcal{A}^L(t_0', t_f')$ computed for other procedures!

We break the circular dependence between procedure summaries and synchronization-point summaries by iteratively computing both. In particular, we compute a sequence $\mathcal{A}_0^L\mathcal{A}_1^L\ldots$ of procedure summary vector addition systems along with a sequence $\mathcal{A}_0\mathcal{A}_1\ldots$ of vector addition systems such that each $\mathcal{A}_i^L$, for $i > 0$, is computed using the transitions of $\mathcal{A}_{i-1}$, and $\mathcal{A}_i$, for $i \geq 0$ is computed using the procedure summaries of $\mathcal{A}_i^L$. Initially $\mathcal{A}_0^L$ contains only the pending-task sets reachable without taking **await** transitions, and $\mathcal{A}_0$ contains only the transitions of $\mathcal{A}_P$ corresponding to intra-procedural and **post** transitions of P, along with transitions to components $\mathcal{A}_0^L$. For $i \geq 0$, $\mathcal{A}_i$ contains transitions to and from the components $\mathcal{A}_i^L(t_0, t_f)$

$$T[\mathbf{await}\ r] \xrightarrow{\vec{n}_j\mathbf{0}} \langle q_0, T[\mathbf{skip}]\rangle \quad \langle q_f, T[\mathbf{skip}]\rangle \xrightarrow{\mathbf{oo}} T[s;\ \mathbf{await}\ r]$$

for each $t_0, t_f \in \mathsf{Tasks}$ such that $j = \mathsf{cn}(r, t_0)$, $s \in \mathsf{rvh}(t_f)$, and q_0 and q_f are the unique initial and final states of $\mathcal{A}_i^L(t_0, t_f)$. (We assume each component $\mathcal{A}_i^L(t_0, t_f)$ has unique initial and final states, distinct from the states of other components. Additionally, we equip each $\mathcal{A}^L(t_0, t_f)$ with auxiliary state to carry the identity $T[\mathbf{skip}]$ of the invoking task to ensure the proper return of control when $\mathcal{A}^L(t_0, t_f)$ completes.)

At each step $i > 0$, we add to $\mathcal{A}_i$ an additive edge summarizing the execution between two synchronization points $T_1[\mathbf{await}\ r]$ and $T_2[\mathbf{await}\ r]$ occurring in P:

$$T_1[\mathbf{skip}] \xhookrightarrow{\mathbf{oo}} T_2[\mathbf{skip}]$$

such that $T_2[\mathbf{skip}]$ is reachable in $\mathcal{A}_{i-1}$ from $T_1[\mathbf{skip}]$, i.e., $\mathbf{0} \in \mathsf{sms}(T_1[\mathbf{skip}]\,, T_2[\mathbf{skip}]\,, \mathcal{A}_{i-1})$. Note that when $T[\mathbf{await}\ r]$ is a synchronization point occurring in P, $T[\mathbf{skip}]$ refers to the program point immediately after the **await** statement. Since there are only polynomially-many such edges that can possibly be added, we are guaranteed to reach a fixed-point $\mathcal{A}_k$ of $\mathcal{A}_0\mathcal{A}_1\ldots$ in a polynomial number of steps. Furthermore, the reachable states of $\mathcal{A}_k$ are precisely the same reachable states of $\mathcal{A}_P$. However, computing $\mathbf{0} \in \mathsf{sms}(t_1, t_2, \mathcal{A}_{i-1})$ at each step is difficult due to the zero-test edge in the **await** statement immediately preceding t_2; this is computationally equivalent to computing reachability of a particular vector in non-recursive vector addition systems.

Theorem 7. *The state-reachability problem for multi-wait single-region finite-value programs is decidable.*

Since practical algorithms to compute vector-reachability is a difficult open problem, we remark that it is possible to obtain algorithms to approximate our state-reachability problem. Consider, for instance, the over-approximate semantics given by transforming each **await** r statement into **while** $\star$ **do ewait** r. Though many more behaviors are present in the resulting program, since not every task is necessarily consumed during the **while** loop, practical algorithmic solutions are more probable (see Section 5.4).

8. Related Work

Formal modeling and verification of multi-threaded programs has been heavily studied, including but not limited to identifying decidable sub-classes [21], and effective over-approximate [13, 19] and under-approximate [9, 23] analyses. To our knowledge little has been done for formal modeling and verification of programs written in explicitly-parallel languages free of thread interleaving. Sen and Viswanathan [35]'s asynchronous programs, which falls out as a special case of our single-wait programs, is perhaps most similar to our work. Practical verification algorithms [20] and complexity analysis [14] of asynchronous programs have been studied. Though decidability results of parallel models have been reported [4, 10] (Bouajjani and Esparza [3] survey of this line of work), these works target abstract computation models, and do not identify precise complexities and optimal algorithms for real-world parallel programming languages, nor do they handle the case where procedures can return unbounded sets of unfinished computations to their callers.

9. Conclusion

We have proposed a general model of recursively parallel programs which captures the concurrency constructs in a variety of popular programming languages. By isolating the fragments corresponding to various language features, we are able to associate corresponding formal models, measure the complexity of state-reachability, and provide precise analysis algorithms. We hope our complexity measurements may be used to guide the design and choice of concurrent programming languages and program analyses. Figure 8 summarizes our results.

Acknowledgments

We greatly appreciate formative discussions with Arnaud Sangnier and Peter Habermehl, and the feedback of Pierre Ganty, Giorgio

State-Reachability in Recursively Parallel Programs			
	result	complexity	language/feature
Handle-Passing			
general	Thm. 1	undecidable	futures, revisions
Single-Wait			
non-aliasing	Thm. 2	PTIME	futures†, revisions†
local scope	Thm. 3	EXPSPACE	—
global scope	Thm. 4	EXPSPACE	asynchronous programs
general	Thm. 5	2EXPTIME	—
† For programs without handle-passing.			
Multi-Wait (single region)			
local scope	Thm. 6	NP	Cilk
general	Thm. 7	decidable	async (X10)

Figure 8. Summary of results for computing state-reachability for finite-value recursively parallel programs.

Delzanno, Rupak Majumdar, Tom Ball, Sebastian Burckhardt, and the anonymous POPL reviewers.

References

[1] P. A. Abdulla, K. Cerans, B. Jonsson, and Y.-K. Tsay. General decidability theorems for infinite-state systems. In *LICS '96: Proc. 11th IEEE Symposium on Logic in Computer Science*, pages 313–321. IEEE Computer Society, 1996.

[2] E. Allen, D. Chase, V. Luchangco, J.-W. Maessen, S. Ryu, G. L. S. Jr., and S. Tobin-Hochstadt. The Fortress language specification. Technical report, Sun Microsystems, Inc., 2006.

[3] A. Bouajjani and J. Esparza. Rewriting models of boolean programs. In *RTA '06: Proc. 17th International Conference on Term Rewriting and Applications*, volume 4098 of *LNCS*, pages 136–150. Springer, 2006.

[4] A. Bouajjani, M. Müller-Olm, and T. Touili. Regular symbolic analysis of dynamic networks of pushdown systems. In *CONCUR '05: Proc. 16th International Conference on Concurrency Theory*, volume 3653 of *LNCS*, pages 473–487. Springer, 2005.

[5] S. Burckhardt, A. Baldassin, and D. Leijen. Concurrent programming with revisions and isolation types. In *OOPSLA '10: Proc. 25thACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications*, pages 691–707. ACM, 2010.

[6] P. Charles, C. Grothoff, V. A. Saraswat, C. Donawa, A. Kielstra, K. Ebcioglu, C. von Praun, and V. Sarkar. X10: an object-oriented approach to non-uniform cluster computing. In *OOPSLA '05: Proc. 20th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications*, pages 519–538. ACM, 2005.

[7] S. Chaudhuri. Subcubic algorithms for recursive state machines. In *POPL '08: Proc. 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 159–169. ACM, 2008.

[8] S. Demri, M. Jurdzinski, O. Lachish, and R. Lazic. The covering and boundedness problems for branching vector addition systems. In *FSTTCS '09: IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, pages 181–192, 2009.

[9] J. Esparza and P. Ganty. Complexity of pattern-based verification for multithreaded programs. In *POPL '11: Proc. 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 499–510. ACM, 2011.

[10] J. Esparza and A. Podelski. Efficient algorithms for pre* and post* on interprocedural parallel flow graphs. In *POPL*, pages 1–11, 2000.

[11] A. Finkel and P. Schnoebelen. Well-structured transition systems everywhere! *Theor. Comput. Sci.*, 256(1-2):63–92, 2001.

[12] C. Flanagan and M. Felleisen. The semantics of future and an application. *J. Funct. Program.*, 9(1):1–31, 1999.

[13] C. Flanagan and S. Qadeer. Thread-modular model checking. In *SPIN '03: Proc. 10th International Workshop on Model Checking Software*, volume 2648 of *LNCS*, pages 213–224. Springer, 2003.

[14] P. Ganty and R. Majumdar. Algorithmic verification of asynchronous programs. *CoRR*, abs/1011.0551, 2010. `http://arxiv.org/abs/1011.0551`.

[15] G. Geeraerts, J.-F. Raskin, and L. V. Begin. Expand, enlarge and check: New algorithms for the coverability problem of wsts. *J. Comput. Syst. Sci.*, 72(1):180–203, 2006.

[16] S. Graf and H. Saïdi. Construction of abstract state graphs with pvs. In *CAV '97: Proc. 9th International Conference on Computer Aided Verification*, volume 1254 of *LNCS*, pages 72–83. Springer, 1997.

[17] S. Haddad and D. Poitrenaud. Recursive petri nets. *Acta Inf.*, 44(7-8): 463–508, 2007.

[18] R. H. Halstead Jr. Multilisp: A language for concurrent symbolic computation. *ACM Trans. Program. Lang. Syst.*, 7(4):501–538, 1985.

[19] T. A. Henzinger, R. Jhala, R. Majumdar, and S. Qadeer. Thread-modular abstraction refinement. In *CAV '03: Proc. 15th International Conference on Computer Aided Verification*, volume 2725 of *LNCS*, pages 262–274. Springer, 2003.

[20] R. Jhala and R. Majumdar. Interprocedural analysis of asynchronous programs. In *POPL '07: Proc. 34th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 339–350. ACM, 2007.

[21] V. Kahlon. Boundedness vs. unboundedness of lock chains: Characterizing decidability of pairwise cfl-reachability for threads communicating via locks. In *LICS '09: Proc. 24th Annual IEEE Symposium on Logic in Computer Science*, pages 27–36. IEEE Computer Society, 2009.

[22] D. Kozen. Lower bounds for natural proof systems. In *FOCS '77: Proc. 18th Annual Symposium on Foundations of Computer Science*, pages 254–266. IEEE Computer Society, 1977.

[23] A. Lal and T. W. Reps. Reducing concurrent analysis under a context bound to sequential analysis. *Formal Methods in System Design*, 35(1): 73–97, 2009.

[24] J. R. Larus and R. Rajwar. *Transactional Memory*. Morgan & Claypool, 2006. `http://www.morganclaypool.com/doi/abs/10.2200/S00070ED1V01Y200611CAC002`.

[25] E. A. Lee. The problem with threads. *IEEE Computer*, 39(5):33–42, 2006.

[26] D. Leijen, W. Schulte, and S. Burckhardt. The design of a task parallel library. In *OOPSLA '09: Proc. 24thACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications*, pages 227–242. ACM, 2009.

[27] R. J. Lipton. The reachability problem requires exponential space. Technical Report 62, Yale University, 1976.

[28] C. H. Papadimitriou. *Computational Complexity*. Addison Wesley, 1993.

[29] P. Pratikakis, H. Vandierendonck, S. Lyberis, and D. S. Nikolopoulos. A programming model for deterministic task parallelism. In *MSPC '11: Proc. 2011 ACM SIGPLAN Workshop on Memory Systems Performance and Correctness*, pages 7–12. ACM, 2011.

[30] C. Rackoff. The covering and boundedness problems for vector addition systems. *Theor. Comput. Sci.*, 6:223–231, 1978.

[31] G. Ramalingam. Context-sensitive synchronization-sensitive analysis is undecidable. *ACM Trans. Program. Lang. Syst.*, 22(2):416–430, 2000.

[32] K. H. Randall. *Cilk: Efficient Multithreaded Computing*. PhD thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, May 1998.

[33] T. W. Reps, S. Horwitz, and S. Sagiv. Precise interprocedural dataflow analysis via graph reachability. In *POPL '95: Proc. 22th ACM*

SIGPLAN-SIGACT Symposium on Principles of Programming Languages, pages 49–61. ACM, 1995.

[34] C. Segulja and T. S. Abdelrahman. Synchronization-free and deterministic coarse-grain parallelism: Architectural support and programming model. In *FASPP '11: Proc. First International Workshop on Future Architectural Support for Parallel Programming*, 2011.

[35] K. Sen and M. Viswanathan. Model checking multithreaded programs with asynchronous atomic methods. In *CAV '06: Proc. 18th International Conference on Computer Aided Verification*, volume 4144 of *LNCS*, pages 300–314. Springer, 2006.

A. Syntactic Sugar

The following syntactic extensions are reducible to the original program syntax of Section 2. Here we freely assume the existence of various type- and expression-constructors. This does not present a problem since our program semantics does not restrict the language of types nor expressions.

Synchronous Procedure calls. Synchronous procedure calls **call** `l := p e` can be encoded by immediately waiting for p to return, i.e., by the sequence

post $\mathtt{r}_{call}$ `:=` p e $\mathtt{d}_{call}$`;`
await $\mathtt{r}_{call}$`,`

where $\mathtt{d}_{call}(v) \stackrel{\text{def}}{=} v$ and $\mathtt{r}_{call}$ is a region identifier.

Multiple types. Multiple type labels $T_1, \ldots, T_j$ can be encoded by systematically replacing each T_i with the sum-type $T = \sum_{i=1}^{j} T_i$. This allows local and global variables with distinct types.

Multiple variables. Additional variables $\mathtt{x}_1\mathtt{:}\ T_1\mathtt{,}\ \ldots\mathtt{,}\ \mathtt{x}_j\mathtt{:}\ T_j$ can be encoded with a single record-typed variable `x:` T, where T is the record type

`{` $\mathtt{f}_1\mathtt{:}\ T_1\mathtt{,}\ \ldots\mathtt{,}\ \mathtt{f}_j\mathtt{:}\ T_j$ `}`

and all occurrences of $\mathtt{x}_i$ are replaced by $\mathtt{x.f}_i$. When combined with the extension allowing multiple types, this allows each procedure to declare any number and type of local variable parameters, distinct from the number and type of global variables.

Local variable declarations. Additional (non-parameter) local variable declarations **var** `l':` T to a procedure p can be encoded by adding `l'` to the list of parameters, and systematically adding an initialization expression (e.g., the choice expression $\star$, or `false`) to the corresponding position in the list of arguments at each call site of p to ensure that `l'` begins correctly (un)initialized.

Unused values. Call assignments **call** `x :=` p e, where `x` is not subsequently used, can be written as **call** `_ :=` p e, where `_:` T is an additional unread local variable, or simpler yet as **call** p e.

Let bindings. Let bindings of the form **let** `x:` T `=` e **in** can be encoded by declaring `x` as a local variable **var** `x:` T immediately followed by an assignment `x :=` e. This construct is used to explicate that the value of `x` remains constant once initialized. The binding **let** `x:` T **in** is encoded by the binding **let** `x:` T `=` $\star$ **in** where $\star$ is the choice expression.

Tuples. Assignments $(\mathtt{x}_1\mathtt{,}\ \ldots\mathtt{,}\ \mathtt{x}_j)$ `:=` e to a tuple of variables $\mathtt{x}_1 \ldots \mathtt{x}_j$ are encoded by the sequence

let `r:` `{` $\mathtt{f}_1\mathtt{:}\ T_1\mathtt{,}\ \ldots\mathtt{,}\ \mathtt{f}_j\mathtt{:}\ T_j$ `}` `=` e **in**
$\mathtt{x}_1$ `:= r.`$\mathtt{f}_1$`;` `...;` $\mathtt{x}_j$ `:= r.`$\mathtt{f}_j$

where `r` is a fresh variable. A tuple expression $(\mathtt{x}_1\mathtt{,}\ \ldots\mathtt{,}\ \mathtt{x}_j)$ occurring in a statement s is encoded as

let `r:` `{` $\mathtt{f}_1\mathtt{:}\ T_1\mathtt{,}\ \ldots\mathtt{,}\ \mathtt{f}_j\mathtt{:}\ T_j$ `}` `=` `{` $\mathtt{f}_1$ `=` $\mathtt{x}_1\mathtt{,}\ \ldots\mathtt{,}$
$\mathtt{f}_j$ `=` $\mathtt{x}_j$ `}` **in**
$s[\mathtt{r}/(\mathtt{x}_1\mathtt{,}\ \ldots\mathtt{,}\ \mathtt{x}_j)]$

where `r` is a fresh variable, and $s[e_1/e_2]$ replaces all occurrences of e_2 in s with e_1. When a tuple-element $\mathtt{x}_i$ on the left-hand side of an assignment is unneeded (e.g., from the return value of a **call**), we may replace the occurrence of $\mathtt{x}_i$ with the `_` variable—see the "unused values" desugaring.

Arrays. Finite arrays with j elements of type T can be encoded as records of type `{` $\mathtt{f}_1\mathtt{:}\ T\mathtt{,}\ \ldots\mathtt{,}\ \mathtt{f}_j\mathtt{:}\ \mathtt{T}$ `}`, where $\mathtt{f}_1 \ldots \mathtt{f}_j$ are fresh names. Occurrences of terms `a[i]` are replaced by $\mathtt{a.f}_i$, and array-expressions `[` $e_1\mathtt{,}\ \ldots\mathtt{,}\ e_j$ `]` are replaced by record-expressions `{` $\mathtt{f}_1$ `=` $e_1\mathtt{,}\ \ldots\mathtt{,}\ \mathtt{f}_j$ `=` e_j `}`.

ToDo: unbounded arrays?

B. Proofs to Selected Theorems

To begin with we introduce notation and simplifying assumptions in order to simplify the proof arguments in the following subsections.

B.1 Notation and Simplifying Assumptions

Explicit Control-Location Representation In the definition of recursively parallel programs, the body of each procedure is defined by a top-level statement, which is built inductively from smaller sub-statements. Control locations (e.g., program-counter values) are represented only implicitly, by considering a procedure name and statement pair $\langle p, s\rangle \in \mathsf{Procs} \times \mathsf{Stmts}$. In order to simplify many of the constructions in the following sections, we will often assume instead an explicit representation of control locations. As the analysis complexity results only apply to finite-data programs (all complexities are undecidable otherwise!), we will further assume that this explicit control-location representation includes a local-variable valuation.

In particular, given a program P over finite sets Procs of procedures of Vals of values, we assume P can be represented by a set Locs of *control locations*, along with *intra-procedural transitions* $\hookrightarrow_P \subseteq \mathsf{Locs} \times \mathsf{Stmts} \times \mathsf{Locs}$ between control locations of the same procedure, labelled by one of the following statements s:

`l :=` e, **skip**, **assume** e, **post** $r \leftarrow p\ e\ d$,
return e, **ewait** r, **await** r.

Furthermore, each $m \in \mathsf{Locs}$ contains a component $\mathsf{val}(m)$ giving a local-variable valuation, in addition to a program-counter location $\mathsf{ctrl}(m)$. We require that each transition $\langle m_1, s, m_2\rangle \in \hookrightarrow_P$ agrees with the rules of Figure 2 and Figure 3; i.e.,

$$
\begin{array}{rlr}
m_1 \overset{\mathtt{l := } e}{\hookrightarrow_P} m_2 & \text{implies} & \mathsf{val}(m_2) = e(\mathsf{val}(m_1)), \\
m_1 \overset{\mathbf{skip}}{\hookrightarrow_P} m_2 & \text{implies} & \mathsf{val}(m_1) = \mathsf{val}(m_2), \\
m_1 \overset{\mathbf{assume}\ e}{\hookrightarrow_P} m_2 & \text{implies} & e(\mathsf{val}(m_1)) = \mathbf{true}, \mathsf{val}(m_1) = \mathsf{val}(m_2) \\
m_1 \overset{\mathbf{post}\ r \leftarrow p\ e\ d}{\hookrightarrow_P} m_2, & \text{implies} & \mathsf{val}(m_1) = \mathsf{val}(m_2).
\end{array}
$$

Words & Languages A Σ-*word* is a finite sequence $w \in \Sigma^*$ of symbols from an *alphabet* Σ; ε denotes the empty word. A *language* $L \subseteq \Sigma^*$ is a set of words. The *Parikh image* $\Pi(w)$ of a word $w \in \Sigma^*$ is the multiset $m \in \mathbb{M}[\Sigma]$ (equivalently, the vector $\vec{n} \in \mathbb{N}^{|\Sigma|}$) such that for each $a \in \Sigma$, $m(a)$ (resp., , $\vec{n}(a)$) is the number of occurrences of a in w; the *Parikh image* of a language $L \subseteq \Sigma^*$ is the set of Parikh images of each constituent word $\Pi(L) = \{\Pi(w) : w \in L\}$. Two languages L_1 and L_2 are *Parikh equivalent* when $\Pi(L_1) = \Pi(L_2)$.

Finite-State Automata A *finite-state automaton (FSA)* $\mathcal{A} = \langle Q, \Sigma, \hookrightarrow \rangle$ over a finite alphabet Σ consists of a finite set Q of *states*, along with a set $\hookrightarrow \subseteq Q \times \Sigma \times Q$ of *transitions*. Given sets $Q_0, Q_f \subseteq Q$ of initial and accepting states, we denote by $\mathcal{A}(Q_0, Q_f)$ the set of Σ-words labelling runs of $\mathcal{A}$ which begin in some initial state $q_0 \in Q_0$ and terminate in an accepting state $q_f \in Q_f$. The *state-reachability problem* for finite-state automata is to decide, given an automaton $\mathcal{A}$ with states Q, and $Q_0, Q_f \subseteq Q$, whether $\mathcal{A}(Q_0, Q_f) \neq \emptyset$.

Context-Free Grammars A *context-free grammar (CFG)* $\mathcal{G} = \langle V, \Sigma, \hookrightarrow \rangle$ over a finite alphabet Σ consists of a finite set V of *variables*, along with a finite set $\hookrightarrow \subseteq V \times (V \cup \Sigma)^*$ of *productions*. Given a set $V_0 \subseteq V$, we denote by $\mathcal{G}(V_0)$ the set of Σ-words derived by $\mathcal{G}$ from some initial non-terminal $v_0 \in V_0$.

Pushdown Automata A *pushdown automaton (PDA)* $\mathcal{A} = \langle Q, \Sigma, \Gamma, \hookrightarrow \rangle$ over a finite alphabet Σ consist of a finite set Q of *states*, along with a finite *stack alphabet* Γ, and a set $\hookrightarrow \subseteq Q \times \Gamma \times \Sigma \times \Gamma^* \times Q$ of *transitions*. A configuration qw is a state $q \in Q$ paired with a stack-symbol sequence $w \in \Gamma^*$. Given two sets $G_0, G_f \subseteq Q \times \Gamma^*$ of configurations, we denote by $\mathcal{A}(G_0, G_f)$ the set of Σ-words labelling runs of $\mathcal{A}$ which begin in an initial configuration $q_0 w_0 \in G_0$ and terminate in an accepting configuration $q_f w_f \in G_f$. The *state-reachability* problem for pushdown automata is to decide, given an automaton $\mathcal{A}$ with states Q and stack-alphabet Γ, along with two finite configuration sets $G_0, G_f \subseteq Q \times \Gamma^*$, whether $\mathcal{A}(G_0, G_f) \neq \emptyset$.

Vector Addition Systems A *vector addition system (VAS)* $\mathcal{A} = \langle Q, \hookrightarrow \rangle$ of dimension $k \in \mathbb{N}$ is a finite set Q of *states*, along with a finite set $\hookrightarrow \subset Q \times \mathbb{N}^k \times \mathbb{N}^k \times Q$ of *transitions*. A VAS configuration $q\vec{n}$ is a state $q \in Q$ paired with a vector $\vec{n} \in \mathbb{N}^k$. Given two sets $N_0, N_f \subseteq Q \times \mathbb{N}^k$ of configurations, we denote by $\mathcal{A}(N_0, N_f)$ the set of runs of $\mathcal{A}$ which begin in an initial configuration $q_0\vec{n}_0 \in N_0$ and terminate in an accepting configuration $q_f\vec{n}_f \in N_f$. *The configuration-reachability problem* (resp., *the state-reachability problem*) for vector addition systems is to determine, given a k-dimension system $\mathcal{A}$ with states Q, along with two finite configuration sets $N_0, N_f \subseteq Q \times \mathbb{N}^k$, whether $\mathcal{A}(N_0, N_f) \neq \emptyset$ (resp., whether $\mathcal{A}(N_0, N_f') \neq \emptyset$, where $N_f' = \{q\vec{n}' : \exists \vec{n}.\ q\vec{n} \in N_f \text{ and } \vec{n} \leq \vec{n}'\}$ is the *upward-closure* of N_f.)

Turing Machines A *Turing machine (TM)* $\mathcal{A} = \langle Q, \Sigma, \hookrightarrow \rangle$ over a finite alphabet Σ consists of a finite set Q of *states*, along with a set $\hookrightarrow \subseteq Q \times \Sigma \times \{\mathrm{L}, \mathrm{R}\} \times \Sigma \times Q$ of *transitions*. A configuration $\langle q, w_1, w_2 \rangle$ is a state $q \in Q$ along with two words $w_1, w_2 \in \Sigma^*$. Given sets $Q_0, Q_f \subseteq Q$ of initial and accepting states, we denote by $\mathcal{A}(Q_0, Q_f)$ the set of Σ-words w such that $\mathcal{A}$ has a run which begins in an initial configuration $\langle q_0, \varepsilon, w \rangle$ and terminates in an accepting configuration $\langle q_f, w_1, w_2 \rangle$, for some $q_0 \in Q_0$, $q_f \in Q_f$, and $w_1, w_2 \in \Sigma^*$. The *state-reachability problem* for Turing machines is to decide, given a machine $\mathcal{A}$ with states Q, and $Q_0, Q_f \subseteq Q$, whether $\mathcal{A}(Q_0, Q_f) \neq \emptyset$.

B.2 Results for Task-Passing Programs

Theorem 1. *The state-reachability problem for n-region finite-value task-passing parallel programs is undecidable for*

(a) non-recursive programs with $n > 1$, and
(b) recursive programs with $n > 0$.

We prove (a) and (b) separately, both by reduction from the language emptiness problem for Turing machines.

a. By reduction from the language emptiness problem for Turing machines, let $\mathcal{A} = \langle Q, \Sigma, \delta, q_0, q_f \rangle$ be a Turing machine with $\delta = \{d_1, \ldots, d_j\}$, and $w \in \Sigma^k$ an input tape. We define the language $L(\mathcal{A})$ as the set of words $w_1 w_2 \in \Sigma^*$ of reachable accepting-state configurations $\langle q_f, w_1, w_2 \rangle$. We assume, without loss of generality, that upon entering the accepting state q_f, $\mathcal{A}$ performs a sequence of left-moves until reaching the end of the tape; i.e., $\mathcal{A}$ contains the transition $q_f \overset{a/a,\mathrm{L}}{\hookrightarrow} q_f$ for all $a \in \Sigma$. We define a task-passing program $P_\mathcal{A}$ with two regions $\mathtt{r}_L$ and $\mathtt{r}_R$, and one return-value handler `d`, along with an initial procedure given by

```
proc main ()
    var state: Q
    var sym: Σ

    while ⋆ do post r_R ← p ⋆ d;
    post r_R ← p w(k) d;
    post r_R ← p w(k − 1) d;
    ...;
    post r_R ← p w(2) d;

    state := q_0;
    sym := w(1);

    while ⋆ do
        if ⋆ then s_1
        else if ⋆ then s_2
        ...
        else if ⋆ then s_j;

    // check: is state = q_f reachable here?
    return
```

and an auxiliary non-recursive procedure `p` given by

```
proc p (var sym: Σ)
    return sym
```

where each transition $d_i \in \delta$ has a corresponding statement s_i defined as follows. For right-moving transitions $d_i = q \overset{a/b,\mathrm{R}}{\hookrightarrow} q'$, we define s_i as

```
assume state = q;
assume sym = a;
post r_L ← p b d;
state := q′;
ewait r_R // overwrites sym
```

For right-moving transitions $d_i = q \overset{a/b,\mathrm{L}}{\hookrightarrow} q'$, we define s_i as

```
assume state = q;
assume sym = a;
post r_R ← p b d;
state := q′;
ewait r_L // overwrites sym
```

where $\mathtt{d}(a)$ assigns a to `sym`.

By connecting the configurations of $\langle q, w_1, w_2 \rangle$ of $\mathcal{A}$ to the chain of tasks in region $\mathtt{r}_L$—corresponding to the cells of w_1—and the chain of tasks in region $\mathtt{r}_R$—corresponding to the cells of w_2—it is routine to show that $P_\mathcal{A}$ faithfully simulates precisely the runs of $\mathcal{A}$. As we assume $\mathcal{A}$ moves to the left upon encountering the accepting state q_f, we need only check reachability of a valuation q_f to `state` at the end of the `main` procedure to know whether or not $\mathcal{A}$ has an accepting run.

Proposition 1. *$L(\mathcal{A}) \neq \emptyset$ if and only if* `state` $= q_f$ *holds in a valuation reachable at the end of the* `main` *procedure of $P_\mathcal{A}$.*

Thus deciding state-reachability of $P_{\mathcal{A}}$ solves language emptiness for $\mathcal{A}$. □ □

Using only a single region, it will not be possible to create two independent, unbounded task chains. However, if the program is allowed to be recursive, we can leverage the unbounded procedure stack as an additional, independent, unbounded data structure.

b. By reduction from the language emptiness problem for Turing machines, let $\mathcal{A} = \langle Q, \Sigma, \delta, q_0, q_f \rangle$ be a Turing machine with $\delta = \{d_1, \ldots, d_j\}$, and $w \in \Sigma^k$ an input tape. We define the language $L(\mathcal{A})$ as the set of words $w_1 w_2 \in \Sigma^*$ of reachable accepting-state configurations $\langle q_f, w_1, w_2 \rangle$. We assume, without loss of generality, that upon entering the accepting state q_f, $\mathcal{A}$ performs a sequence of left-moves until reaching the end of the tape; i.e., $\mathcal{A}$ contains the transition $q_f \overset{a/a,\mathrm{L}}{\longrightarrow} q_f$ for all $a \in \Sigma$. We define a single-region task-passing program $P_{\mathcal{A}}$ with a single return-value handler `d`, along with an initial procedure given by

```
proc main ()
    var q_cur, q_R: Q
    var sym_R: Σ

    while ⋆ do post r ← p ⋆ d;
    post r ← p w(k) d;
    post r ← p w(k − 1) d;
    ...;
    post r ← p w(1) d;

    ewait r;
    assume q_R = q0;

    // check: is q_cur = qf reachable here?
    return
```

and an auxiliary recursive procedure `p` given by

```
proc p (var sym: Σ)
    var q_cur, q_init, q_R: Q
    var sym_R: Σ

    q_init := ⋆;
    q_cur := q_init;

    while ⋆ do
        if ⋆ then s1
        else if ⋆ then s2
        ...
        else if ⋆ then sj
```

where each transition $d_i \in \delta$ has a corresponding statement s_i defined as follows. For right-moving transitions $d_i = q \overset{a/b,\mathrm{R}}{\longrightarrow} q'$, we define s_i as

```
assume q_cur = q;
assume sym = a;
sym := b;

ewait r;
// At this point q_R, q_cur, and sym_R
//   have been overwritten by the initial-
//   and current-state valuations, and the
//   symbol stored in the right-neighbor
//   who has just moved left.
assume q_R = q′;
post r ← p sym_R d
```

where $\mathtt{d}(q, q', a)$ assigns q to `q_R`, q' to `q_cur`, and a to `sym_R`. Our program thus simulates right moves by awaiting a pending task representing the right neighbor of the current task. For left-moving transitions $d_i = q \overset{a/b,\mathrm{L}}{\longrightarrow} q'$, we define s_i as

```
assume q_cur = q;
assume sym = a;
return (q_init, q′, b);
```

Our program thus simulates left moves by returning to the awaiting task, who promptly recreates its right-neighbor by posting a new task to replace it.

By connecting the configurations of $\langle q, w_1, w_2 \rangle$ of $\mathcal{A}$ to the chain of awaiting tasks—corresponding to the cells of w_1—and the chain of posted tasks—corresponding to the cells of w_2—it is routine to show that $P_{\mathcal{A}}$ faithfully simulates precisely the runs of $\mathcal{A}$. As we assume $\mathcal{A}$ moves to the left upon encountering the accepting state q_f, we need only check reachability of a valuation q_f to `q_cur` at the end of the `main` procedure to know whether or not $\mathcal{A}$ has an accepting run.

Proposition 2. *$L(\mathcal{A}) \neq \emptyset$ if and only if* `q_cur` $= q_f$ *holds in a valuation reachable at the end of the* `main` *procedure of $P_{\mathcal{A}}$.*

Thus deciding state-reachability of $P_{\mathcal{A}}$ solves language emptiness for $\mathcal{A}$. □ □

Theorem ??. *The state-reachability problem for single-region non-recursive finite-value task-passing parallel programs is PTIME-complete for fixed task-depth, and EXPTIME in the task-depth.*

Proof. Let P be a non-aliasing single-wait finite-value single-region non-recursive task-passing parallel program with finite sets of procedures Procs, values Vals, regions Regs, and return-value handlers Rets, and let $\ell \in \mathsf{Vals}$ be a target reachable value. Furthermore, we assume P is non-recursive, which implies there is a maximum task-depth $N \in \mathbb{N}$—i.e., N is the maximum length of a sequence $p_0 p_1 \ldots \in \mathsf{Procs}^*$ such that each p_i contains a post to p_{i+1}.

We construct a pushdown automaton $\mathcal{A}_P = \langle Q, \Sigma, \Gamma, \delta, q_0, \gamma_0, Q_f \rangle$ which accepts a run when the control-state reaches some $q_f \in Q_f$. We define the states of $\mathcal{A}_P$ to be N-bounded sequences of program control locations and return handlers:

$$Q \overset{\text{def}}{=} \mathsf{Locs} \times (\mathsf{Rets} \times \mathsf{Locs})^{<N}.$$

In this way a state $m_0 d_1 m_1 d_2 m_2 \ldots d_i m_i \in Q$ represents a computation of P in which each control location m_{j-1} $(0 < j \leq i)$ is of a task posted with return-value handler d_j by a task in control location m_j. Note that this representation is only possible since we know the task-depth is bounded by N. Given this state-representation, we define the transition relation δ of $\mathcal{A}_P$ as follows:

Intra-task transitions For each intra-task transition of Figure 2 from control location $m_1 \in \mathsf{Locs}$ to $m_2 \in \mathsf{Locs}$, we add the transition

$$m_1 \longhookrightarrow m_2.$$

POST For each transition of the statement **post** `r` $\leftarrow p\ v\ d$ from control location $m_1 \in \mathsf{Locs}$ to $m_2 \in \mathsf{Locs}$, we add a transition which transfers control directly to p, recording the return location in the finite store,

$$m_1 \vec{m} \longhookrightarrow m' d m_2 \vec{m},$$

where m' is the entry location of procedure p with argument v.

$\exists$**WAIT** For each transition of the statement $\mathbf{ewait}\ \mathtt{r}$ from control location $m_1 \in \mathsf{Locs}$ to $m_2 \in \mathsf{Locs}$, we add a transition which simply pops the pair $\langle v, d\rangle$ from the top of the pushdown stack, and applies the return-value handler,

$$m_1\vec{m} \overset{\text{pop}\langle v,d\rangle}{\hookrightarrow} m_2'\vec{m},$$

where m_2' is the control state reached after applying the return-value handler d to the return value v at control state m_2.

RETURN For each transition of the statement $\mathbf{return}\ v$ from control location $m' \in \mathsf{Locs}$, we add a transition which pushes the return value and return-value handler for the current task onto the pushdown stack, to be later consumed by a subsequent $\mathbf{ewait}$ statement,

$$mdm_2\vec{m} \overset{\text{push}\langle v,d\rangle}{\hookrightarrow} m_2\vec{m}.$$

Proposition 3. *$L(\mathcal{A}_P) \neq \emptyset$ if and only if ℓ is reachable in P.*

As $|Q|$ is $\mathcal{O}((|\mathsf{Locs}| \cdot |\mathsf{Rets}|)^N)$ and $|\Gamma|$ is $\mathcal{O}(|\mathsf{Vals}| \cdot |\mathsf{Rets}|)$, the size of $\mathcal{A}_P$ is polynomial in P. Since language emptiness is decidable in polynomial time for pushdown automata, our procedure gives a polynomial-time algorithm for state-reachability when N is fixed, though exponential in N. □ □

B.3 Results for Single-Wait Programs

Theorem 2. *The state-reachability problem for non-aliasing single-wait finite-value finite-region programs is PTIME-complete when the number of regions is fixed, and EXPTIME-complete in the number of regions.*

ToDo: change the following proof from local-scope to unscoped

Proof. Let P be a non-aliasing local-scope single-wait finite-value program with regions $\mathtt{r}_1, \ldots, \mathtt{r}_n$. We define a *sequential* finite-value program P_s by a code-to-code translation of P. We extend each procedure declaration $\mathbf{proc}\ p\ \mathtt{(var\ l:}\ T\mathtt{)}\ s$ with additional procedure-local variables `rg`, `rg'`, and `rv`,

```
proc p (var l: T)
    var rg[n]: R := [ ⊥; ..; ⊥ ]
    var rg'[n]: R
    var rv: T
    s
```

where R is a type containing $\bot$, and values of the record type

```
{ prc: Procs, arg: Vals, rh: Rets }.
```

Note that R is a finite-type since Procs, Vals, and Rets are finite sets. We translate each statement $\mathbf{return}\ e$ into $\mathbf{return}\ \mathtt{(rg,}\ e\mathtt{)}$, each statement $\mathbf{post}\ \mathtt{r}_i \leftarrow p\ e\ d$ into the assignment

```
rg[i] := { prc = p, arg = e, rh = d }
```

and each statement $\mathbf{ewait}\ \mathtt{r}_i$ into the statement

```
assume rg[i] ≠ ⊥;
call (rg',rv) := rg[i].prc rg[i].arg;
l := rg[i].rh;
rg[i] := ⊥;
for j := 1 to n do
    if rg'[j] ≠ ⊥ then rg[j] := rg'[j]
```

where we assume each $d \in \mathsf{Rets}$ is given by an expression in which `rv` is a free variable. Note that for local-scope programs, the `rg'` array will always be equal to `[ ⊥; ..; ⊥ ]` and can be safely omitted from the translation.

Since regions do not alias, it is not hard to show that the state-reachability problem for the resulting sequential program P_s is equivalent to the state-reachability problem for P. Furthermore, the size of P_s is polynomial in P, while the number of variables in P_s increases by n . Thus our state-reachability problem is PTIME-complete for fixed n since the state-reachability for sequential programs is [7, 33]. When the number n of regions is not fixed, this state-reachability problem becomes EXPTIME-complete, due to the logarithmic encoding of the program values into the n extra variables. □ □

Theorem 3. *The state-reachability problem for local-scope single-wait finite-value finite-region programs is EXPSPACE-complete.*

We show an equivalence between the state-reachability problems of local-scope single-wait recursively parallel programs and vector addition systems (VASS, introduced in Appendix C)—i.e., we show the problems are polynomial-time reducible to each other. EXPSPACE-completeness follows since state-reachability in VASS is known to be EXPSPACE-complete.

Lemma 5. *The state-reachability problem for local-scope single-wait finite-value finite-region programs is polynomial-time reducible to the state-reachability problem for vector addition systems with states (VASS).*

Proof. **ToDo: fix this proof; it was just ported from a later section...**

Let P be a program with finite sets of procedures Procs, values Vals, regions Regs, and return-value handlers Rets, and let $\ell \in \mathsf{Vals}$ be a target reachable value from an initial configuration $c_0 = \langle \ell_0, s, \emptyset\rangle$. We construct a recursive vector addition system $\mathcal{A}_P = \langle Q, q_0, \delta\rangle$ and a set of target states Q_f such that such that ℓ is reachable in P from c_0 if and only if some $q_f \in Q_f$ is reachable in $\mathcal{A}_P$.

To construct the state-space Q of $\mathcal{A}_P$, we assume an explicit representation Locs of P's control locations (and the control locations of the initial statement s). Furthermore, we assume each control-location $m \in \mathsf{Locs}$ also contains a component containing a local-variable valuation. We define $Q \overset{\text{def}}{=} \mathsf{Locs} \cup (\mathsf{Locs} \times \mathsf{Tasks})$ such that each $m \in Q \cap \mathsf{Locs}$ is a *control state* and each $\langle m, p, v, d\rangle \in Q \cap (\mathsf{Locs} \times \mathsf{Tasks})$ is an *auxiliary state* used in the translation of $\mathbf{ewait}$ statements, since we will need two RVASS transitions to model $\mathbf{ewait}$. We set q_0 to the unique state given by the initial control location with the local-variable valuation ℓ_0.

To model the region containers $f : \mathsf{Regs} \to \mathbb{M}[\mathsf{Tasks}]$ we will use a set of counters, one per region-task pair $\langle r, p, v, d\rangle \in \mathsf{Regs} \times \mathsf{Tasks}$. Let $n = |\mathsf{Regs} \times \mathsf{Tasks}|$, and fix an enumeration $\mathsf{cn} : \mathsf{Regs} \times \mathsf{Tasks} \to n$. In this way, a frame $\langle \ell, s, f\rangle$ of P is represented as a configuration $\langle q, \vec{n}\rangle$ of $\mathcal{A}_P$, where q encodes the control state (including ℓ and s), and $\vec{n}$ encodes the region valuation f, where $\vec{n}(\mathsf{cn}(r, p, v, d)) = f(r)(\langle p, v, d\rangle)$ for each $r \in \mathsf{Regs}$ and $\langle p, v, d\rangle \in \mathsf{Tasks}$. Let $\vec{n}_i$ denote the unit vector of dimension i, i.e., $\vec{n}_i(i) = 1$ and $\vec{n}_i(j) = 0$ for $j \neq i$. Given this mapping between configurations, the construction of δ is straightforward; we outline below.

Inter-task transitions For each intra-task transition of Figure 2 from control location $m_1 \in \mathsf{Locs}$ to $m_2 \in \mathsf{Locs}$ we add a transition

$$m_1 \overset{\mathbf{0}}{\hookrightarrow} m_2.$$

POST For each transition of the statement $\mathbf{post}\ r \leftarrow p\ v\ d$ from control location $m_1 \in \mathsf{Locs}$ to $m_2 \in \mathsf{Locs}$, we add a transition which simply increments the counter $i = \mathsf{cn}(r, p, v, d)$,

$$m_1 \overset{\mathbf{0}\vec{n}_i}{\hookrightarrow} m_2.$$

$\exists$**WAIT** For each transition of the statement $\mathbf{ewait}\ r$ from control location $m_1 \in \mathsf{Locs}$, we add ...

a transition from m_1 to $m_2 \in \mathsf{Locs}$ to simultaneously decrement counter $i = \mathsf{cn}(r, p, v, d)$ and apply the return-value handler d,

$$m_1 \xrightarrow{\vec{n}_i \mathbf{0}} m_2 \tag{1}$$

where m_0 is the entry control location for procedure p with argument v, m_f is an exit control location for procedure p with return-value v',

ToDo: which was computed as a summary ...

and m_2 is the control state reached after applying the return-value handler d to the return value v' at control-state m_1.

The correspondence between P and $\mathcal{A}_P$ is easily shown by induction over the transitions between configurations of P and $\mathcal{A}_P$. It follows that ℓ is reachable in P if and only if there exists some reachable state $m \in \mathsf{Locs}$ of $\mathcal{A}_P$, where the local-variable valuation embedded into m is equal to ℓ. Since there are finitely-many such $m \in \mathsf{Locs}$, we can reduce reachability in P to a finite number of reachability queries in $\mathcal{A}_P$. As the local-state valuations are made explicit in the representation of $\mathcal{A}_P$, the size of $\mathcal{A}_P$ becomes $\mathcal{O}(|P| \cdot |\mathsf{Vals}|)$. □ □

Lemma 6. *The state-reachability problem for vector addition systems with states (VASS) is polynomial-time reducible to the state-reachability problem for local-scope single-wait finite-value finite-region programs.*

Proof. Let $k \in \mathbb{N}$, and let $\mathcal{A} = \langle Q, q_0, \delta \rangle$ be a k-dimension VASS, and let q_f be a target reachable state. We construct a single-wait program $P_{\mathcal{A}}$ and a target valuation v_f such that q_f is reachable in $\mathcal{A}$ if and only if v_f is reachable in $P_{\mathcal{A}}$.

The program $P_{\mathcal{A}}$ contains only two procedures: an initial procedure `main` and a dummy procedure `p` which will be posted (resp., awaited) for each addition (resp., subtraction) performed in $\mathcal{A}$. Accordingly, the region-set $\mathsf{Regs} = \{\mathtt{r}_1, \ldots, \mathtt{r}_k\}$ of $P_{\mathcal{A}}$ contains a region $\mathtt{r}_i$ per vector component. The program's local variable `l` is used to store the control-state of $\mathcal{A}$, and we set $\mathsf{Vals} = Q$. Finally, let $\mathsf{Rets} = \{\mathtt{d}_{const}\}$, where $\mathtt{d}_{const}(v) \stackrel{\text{def}}{=} \mathtt{l}$; i.e., $\mathtt{d}_{const}$ is the return-value handler which ignores the return value, keeping the local valuation intact.

We simulate the transitions of $\mathcal{A}$ by awaiting a task from each region $\mathtt{r}_i$ once per decrement to the ith vector component, and subsequently posting a task to each region $\mathtt{r}_i$ once per increment to the ith vector component. Thus for each transition $d_j = q \xrightarrow{\vec{n}_1 \vec{n}_2} q'$, we define the statement s_j given by

$\mathbf{assume}\ \mathtt{l} = q$

$\underbrace{\mathbf{ewait}\ \mathtt{r}_1\mathtt{;}\ \ldots\mathtt{;}\ \mathbf{ewait}\ \mathtt{r}_1}_{\vec{n}_1(1)\text{ times}}\mathtt{;}\ \ldots\mathtt{;}\ \underbrace{\mathbf{ewait}\ \mathtt{r}_k\mathtt{;}\ \ldots\mathtt{;}\ \mathbf{ewait}\ \mathtt{r}_k}_{\vec{n}_1(k)\text{ times}}\mathtt{;}$

$\underbrace{\mathbf{post}\ \mathtt{r}_1 \leftarrow \mathtt{p} \star \mathtt{d}_{const}\mathtt{;}\ \ldots\mathtt{;}\ \mathbf{post}\ \mathtt{r}_1 \leftarrow \mathtt{p} \star \mathtt{d}_{const}}_{\vec{n}_2(1)\text{ times}}\mathtt{;}$

$\ldots\mathtt{;}$

$\underbrace{\mathbf{post}\ \mathtt{r}_k \leftarrow \mathtt{p} \star \mathtt{d}_{const}\mathtt{;}\ \ldots\mathtt{;}\ \mathbf{post}\ \mathtt{r}_k \leftarrow \mathtt{p} \star \mathtt{d}_{const}}_{\vec{n}_2(k)\text{ times}}\mathtt{;}$

$\mathtt{l}\ \mathtt{:=}\ q'.$

Finally, the initial procedure is given by

```
proc main ()
    l := q_0;
    while ⋆ do
        if ⋆ then s_1
        else if ⋆ then s_2
        ...
        else if ⋆ then s_{|δ|}.
```

Note the correspondence between configurations of $\mathcal{A}$ and $P_{\mathcal{A}}$. Each configuration $\langle q, \vec{n} \rangle$ of $\mathcal{A}$ maps directly to a configuration $\langle q, s, f \rangle$ of $P_{\mathcal{A}}$, where s is the loop statement of the initial procedure, and $|f(\mathtt{r}_i)| = \vec{n}(i)$. Given this correspondence, it follows easily that the "state" q_f is reachable in $\mathcal{A}$ if and only if the "valuation" q_f is reachable in $P_{\mathcal{A}}$. As there are $\mathcal{O}(|\mathcal{A}|)$ statements in $P_{\mathcal{A}}$ per transition of $\mathcal{A}$, the size of $P_{\mathcal{A}}$ is $\mathcal{O}(|\mathcal{A}|^2)$. □ □

Theorem 4. *The state-reachability problem for global-scope single-wait finite-value finite-region programs is EXPSPACE-complete.*

To proceed we show an equivalence between the state-reachability problems of global-scope single-wait recursively parallel programs and vector addition systems (VASS, introduced in Appendix C)—i.e., we show the problems are polynomial-time reducible to each other. EXPSPACE-completeness follows since state-reachability in VASS is known to be EXPSPACE-complete.

Lemma 7. *The state-reachability problem for vector addition systems with states (VASS) is polynomial-time reducible to the state-reachability problem for global-scope single-wait finite-value finite-region programs*

Proof. As the program $P_{\mathcal{A}}$ constructed in Lemma 6 from a given VASS $\mathcal{A}$ only uses the $\mathbf{ewait}$ statement in the initial procedure, $P_{\mathcal{A}}$ is also a global-scope program. □ □

Lemma 8. *The state-reachability problem for global-scope single-wait finite-value finite-region programs is polynomial-time reducible to the state-reachability problem for vector addition systems with states (VASS).*

Proof. Let P be a program with finite sets of procedures Procs, values Vals, regions Regs, and return-value handlers Rets, and let $\ell \in \mathsf{Vals}$ be a target reachable value. We construct a vector addition system with states $\mathcal{A}_P = \langle Q, q_0, \delta \rangle$ and a set of target states Q_f such that ℓ is reachable in P if and only if some $q_f \in Q_f$ is reachable in Q_f.

To construct the state-space Q of $\mathcal{A}_P$, we assume an explicit representation Locs of the control locations of P's initial procedure.

Furthermore, we assume each control location $m \in \mathsf{Locs}$ also contains a component $\mathsf{val}(m)$ containing a local-variable valuation.

ToDo: define Σ, and the widgets

We define $Q \stackrel{\text{def}}{=} \mathsf{Locs} \cup \{Q_a : a \in \Sigma\}$, and we set q_0 to the unique control state given by the initial control location with the local-variable valuation ℓ_0.

To model the region containers $f : \mathsf{Regs} \to \mathbb{M}[\mathsf{Tasks}]$ we use a set of counters, one per region-task pair $\langle r, p, v, d \rangle \in \mathsf{Regs} \times \mathsf{Tasks}$. Let $n = |\mathsf{Regs} \times \mathsf{Tasks}|$, and fix an enumeration $\mathsf{cn} : \mathsf{Regs} \times \mathsf{Tasks} \to n$. In this way, a frame $\langle \ell, s, f \rangle$ of P is represented as a configuration $\langle q, \vec{n} \rangle$ of $\mathcal{A}_P$, where q encodes the control state (including ℓ and s), and $\vec{n}$ encodes the region valuation f, where $\vec{n}(\mathsf{cn}(r, p, v, d)) = f(r)(\langle p, v, d \rangle)$ for each $r \in \mathsf{Regs}$ and $\langle p, v, d \rangle \in \mathsf{Tasks}$. Let $\vec{n}_i$ denote the unit vector of dimension i, i.e., $\vec{n}(i) = 1$ and $\vec{n}(j) = 0$ for $j \neq i$. Given this mapping between configurations, the construction of δ is straightforward; we outline below.

Proposition 4 (Ganty and Majumdar [14])**.** *For every semi-sequential program P over alphabet Σ, there exists a VASS $\mathcal{A}_P$ with a state q_f such that for all $w \in \Sigma^*$, $w \in L(P)$ if and only if some $\langle q_f, \vec{n} \rangle$ is reachable in $\mathcal{A}_P$ with $\vec{n}(a) = \Pi(w)(a)$ for all $a \in \Sigma$. Furthermore, the size of $\mathcal{A}_P$ is polynomial in the size of P.*

Task-components For each task $a = \langle p, v, d\rangle \in \Sigma$, let P_a be the semi-sequential program defined by $p \in \mathsf{Procs}$ at initial valuation $v \in \mathsf{Vals}$, and let $\mathcal{A}_a = \langle Q_a, q_{a,0}, \delta_a\rangle$ and $q_{a,f}$ be the VASS and VASS state given by Proposition 4. We include each $\mathcal{A}_a$ as a subcomponent of $\mathcal{A}_P$, which is accessible from $\mathcal{A}_P$ by the transitions added in the translation of the $\mathbf{ewait}$ statement.

Inter-task transitions For each inter-task transition of Figure 2 from control location $m_1 \in \mathsf{Locs}$ to $m_2 \in \mathsf{Locs}$ we add a transition

$$m_1 \xhookrightarrow{\mathbf{0}\mathbf{0}} m_2.$$

POST For each transition of the statement $\mathbf{post}\ r \leftarrow p\ v\ d$ from control location $m_1 \in \mathsf{Locs}$ to $m_2 \in \mathsf{Locs}$, we add a transition which simply increments the counter $i = \langle r, p, v, d\rangle$,

$$m_1 \xhookrightarrow{\mathbf{0}\vec{n}_i} m_2.$$

∃WAIT For each transition of the statement $\mathbf{ewait}\ r$ from control location $m_1 \in \mathsf{Locs}$, we add two transitions for each task $a = \langle p, v, d\rangle$: one to decrement counter $i = \langle r, p, v, d\rangle$ and pass control to the VASS $\mathcal{A}_a$, and one to return control from $\mathcal{A}_a$, and apply the return-value handler d,

$$m_1 \xhookrightarrow{\vec{n}_i\mathbf{0}} q_{a,0} \qquad q_{a,f} \xhookrightarrow{\mathbf{0}\mathbf{0}} m_2$$

where m_2 is the control state reached after applying the return-value handler d to the program value at $q_{a,f}$ at control state m_1.

... □

Theorem 5. *The state-reachability problem for single-wait finite-value finite-region programs is decidable.*

To proceed we show an equivalence between the state-reachability problems of single-wait recursively parallel programs and recursive vector addition systems (RVASS, introduced in Appendix C)—i.e., we show the problems are polynomial-time reducible to each other. Decidability follows since, as we show in Theorem 8 (of Appendix C.2), RVASS state-reachability is decidable.

Lemma 9. *The state-reachability problem for single-wait finite-value finite-region programs P over values* Vals *is reducible to the state-reachability problem for recursive vector addition systems in time $\mathcal{O}(|P| \cdot |\mathsf{Vals}|)$.*

Proof. Let P be a program with finite sets of procedures Procs, values Vals, regions Regs, and return-value handlers Rets, and let $\ell \in \mathsf{Vals}$ be a target reachable value from an initial configuration $c_0 = \langle \ell_0, s, \emptyset\rangle$. We construct a recursive vector addition system $\mathcal{A}_P = \langle Q, q_0, \delta\rangle$ and a set of target states Q_f such that such that ℓ is reachable in P from c_0 if and only if some $q_f \in Q_f$ is reachable in $\mathcal{A}_P$.

To construct the state-space Q of $\mathcal{A}_P$, we assume an explicit representation Locs of P's control locations (and the control locations of the initial statement s). Furthermore, we assume each control-location $m \in \mathsf{Locs}$ also contains a component containing a local-variable valuation. We define $Q \stackrel{\text{def}}{=} \mathsf{Locs} \cup (\mathsf{Locs} \times \mathsf{Tasks})$ such that each $m \in Q \cap \mathsf{Locs}$ is a *control state* and each $\langle m, p, v, d\rangle \in Q \cap (\mathsf{Locs} \times \mathsf{Tasks})$ is an *auxiliary state* used in the translation of $\mathbf{ewait}$ statements, since we will need two RVASS transitions to model $\mathbf{ewait}$. We set q_0 to the unique state given by the initial control location with the local-variable valuation ℓ_0.

To model the region containers $f : \mathsf{Regs} \to \mathbb{M}[\mathsf{Tasks}]$ we will use a set of counters, one per region-task pair $\langle r, p, v, d\rangle \in \mathsf{Regs} \times \mathsf{Tasks}$. Let $n = |\mathsf{Regs} \times \mathsf{Tasks}|$, and fix an enumeration $\mathsf{cn} : \mathsf{Regs} \times \mathsf{Tasks} \to n$. In this way, a frame $\langle \ell, s, f\rangle$ of P is represented as a configuration $\langle q, \vec{n}\rangle$ of $\mathcal{A}_P$, where q encodes the control state (including ℓ and s), and $\vec{n}$ encodes the region valuation f, where $\vec{n}(\mathsf{cn}(r, p, v, d)) = f(r)(\langle p, v, d\rangle)$ for each $r \in \mathsf{Regs}$ and $\langle p, v, d\rangle \in \mathsf{Tasks}$. Let $\vec{n}_i$ denote the unit vector of dimension i, i.e., $\vec{n}_i(i) = 1$ and $\vec{n}_i(j) = 0$ for $j \neq i$. Given this mapping between configurations, the construction of δ is straightforward; we outline below.

Inter-task transitions For each intra-task transition of Figure 2 from control location $m_1 \in \mathsf{Locs}$ to $m_2 \in \mathsf{Locs}$ we add an additive transition

$$m_1 \xhookrightarrow{\mathbf{0}\mathbf{0}} m_2.$$

POST For each transition of the statement $\mathbf{post}\ r \leftarrow p\ v\ d$ from control location $m_1 \in \mathsf{Locs}$ to $m_2 \in \mathsf{Locs}$, we add an additive transition which simply increments the counter $i = \mathsf{cn}(r, p, v, d)$,

$$m_1 \xhookrightarrow{\mathbf{0}\vec{n}_i} m_2.$$

∃WAIT For each transition of the statement $\mathbf{ewait}\ r$ from control location $m_1 \in \mathsf{Locs}$, we add two transitions: one recrusive transition to perform a recursive call of some task $\langle p, v, d\rangle$ in region r, and one additive transition from m_1 to $m_2 \in \mathsf{Locs}$ to simultaneously decrement counter $i = \mathsf{cn}(r, p, v, d)$ and apply the return-value handler d,

$$m_1 \xhookrightarrow{\vec{n}_i\mathbf{0}} \langle m_2, p, v', d\rangle \xhookrightarrow{m_0 m_f} m_2$$

where m_0 is the entry control location for procedure p with argument v, m_f is an exit control location for procedure p with return-value v', and m_2 is the control state reached after applying the return-value handler d to the return value v' at control-state m_1. (Note that the order of these transitions is important, since counter i must be checked *before* control, and possibly more instances of task $\langle p, v, d\rangle$ in region r, returns from procedure p.)

The correspondence between P and $\mathcal{A}_P$ is easily shown by induction over the transitions between configurations of P and $\mathcal{A}_P$. It follows that ℓ is reachable in P if and only if there exists some reachable state $m \in \mathsf{Locs}$ of $\mathcal{A}_P$, where the local-variable valuation embedded into m is equal to ℓ. Since there are finitely-many such $m \in \mathsf{Locs}$, we can reduce reachability in P to a finite number of reachability queries in $\mathcal{A}_P$. As the local-state valuations are made explicit in the representation of $\mathcal{A}_P$, the size of $\mathcal{A}_P$ becomes $\mathcal{O}(|P| \cdot |\mathsf{Vals}|)$. □ □

Lemma 10. *The state-reachability problem for recursive vector addition systems $\mathcal{A}$ is reducible to the state-reachability problem for single-wait finite-value finite-region programs in time $\mathcal{O}(|\mathcal{A}|^2)$.*

Proof. Let $k \in \mathbb{N}$, and let $\mathcal{A} = \langle Q, q_0, \delta_1 \uplus \delta_2\rangle$ be a RVASS over k-length vectors with additive transitions δ_1 and recursive transitions δ_2, and let q_f be a target reachable state. We construct a single-wait program $P_{\mathcal{A}}$ and a target valuation v_f such that q_f is reachable in $\mathcal{A}$ if and only if v_f is reachable in $P_{\mathcal{A}}$.

The program $P_{\mathcal{A}}$ contains two types of procedures: a set of recursive procedure $\{p_q : q \in Q\}$ whose invocations will correspond to recursive transitions in $\mathcal{A}$, and a *dummy procedure* p_{D} which will be posted (resp., awaited) for each addition (resp., subtraction) performed in $\mathcal{A}$. Accordingly, the region-set $\mathsf{Regs} = \{\mathtt{r}_1, \ldots, \mathtt{r}_k, \mathtt{r}_{call}\}$ of P contains a region $\mathtt{r}_i$ per vector component, and a *call region* $\mathtt{r}_{call}$. As the program's local variable ℓ is used to store the control-state of $\mathcal{A}$, we set $\mathsf{Vals} = Q$. Finally, let $\mathsf{Rets} = \{\mathtt{d}_{const}\}$, where $\mathtt{d}_{const}(v) \stackrel{\text{def}}{=} \ell$; i.e., $\mathtt{d}_{const}$ is the return-value handler which ignores the return value, keeping the local valuation intact.

The top-level statement for the dummy procedure p_0 is simply $\mathbf{return}\ \ \star$; the top-level statement for the other procedures p_q for $q \in Q$ will simulate all transitions of $\mathcal{A}$ and return only when the control-state reaches q. Let $\delta = \{d_1, \ldots, d_n\}$. We define s_i for each $d_i \in \delta$ as follows. We simulate recursive transitions by

calling a procedure which may only return upon reaching q_2. For each transition $d_i = q \xrightarrow{q_1 q_2} q'$, s_i is given by

```
assume l = q;
call l := p_{q_2} q_1;
l := q'.
```

We simulate the additive transitions by awaiting a task from each region $\mathtt{r}_i$ once per decrement to the ith vector component, and subsequently posting a task to each region $\mathtt{r}_i$ once per increment to the ith vector component. For each transition $d_i = q \xrightarrow{\vec{n}_1 \vec{n}_2} q'$, s_i is given by

```
assume l = q
ewait r_1; ...; ewait r_1 ; ...; ewait r_k; ...; ewait r_k ;
   (r_1 times)                    (r_k times)
   post r_1 ← p_0 ⋆ d_const; ...; post r_1 ← p_0 ⋆ d_const ;
   (n_2(1) times)
   ...;
   post r_k ← p_0 ⋆ d_const; ...; post r_k ← p_0 ⋆ d_const ;
   (n_2(k) times)
l := q'.
```

Finally, the top-level statement for procedure p_q is

```
while ⋆ do
    if l = q and ⋆ then return ⋆
    else if ⋆ then s_1
    else if ⋆ then s_2
    ...
    else if ⋆ then s_n
    else skip.
```

Note the correspondence between configurations of $\mathcal{A}$ and $P_{\mathcal{A}}$. Each frame $\langle q, \vec{n} \rangle$ of $\mathcal{A}$ maps directly to a frame $\langle q, s, f \rangle$ of $P_{\mathcal{A}}$, where s is the top-level statement of some procedure $p_{q'}$, and $|f(\mathtt{r}_i)| = \vec{n}(i)$; this correspondence indeed extends directly to the configurations of $\mathcal{A}$ and $P_{\mathcal{A}}$. Given this correspondence, it follows easily that the "state" q_f is reachable in $\mathcal{A}$ if and only if the "valuation" q_f is reachable in $P_{\mathcal{A}}$. As there are $\mathcal{O}(|Q|)$ statements in $P_{\mathcal{A}}$ per transition of $\mathcal{A}$, the size of $P_{\mathcal{A}}$ is $\mathcal{O}(|\mathcal{A}|^2)$. □ □

B.4 Results for Multi-Wait Programs

Theorem ??. *The state-reachability problem for simple-handler multi-wait finite-value finite-region programs is PTIME-complete when the number of regions is fixed, and EXPTIME-complete in the number of regions and return-value handlers.*

Proof. Let P be a simple-handler multi-wait finite-value finite-region program with regions $\mathtt{r}_1, \ldots, \mathtt{r}_n$ and return handlers $\mathtt{d}_1, \ldots, \mathtt{d}_k$. For simplicity, suppose additionally that P is a local-scope program. We define a sequential finite-value program P_{s} by a code-to-code translation of P. First we extend each procedure declaration **proc** p `(var l:` T`)` s with additional procedure-local variables `sum`, `rh`, and `rv`,

```
proc p (var l: T)
    var sum[n]: { first: T, last: T }
    var rh[n][k]: 𝔹
    var rv: T
    for i := 1 to n do
        let guess = ⋆ in
            sum[i] := { first = guess; last =
                guess };
        rh[i] := [ false; ..; false ];
    s
```

Then we translate each statement **post** $\mathtt{r}_i \leftarrow p\ e\ \mathtt{d}_j$ into the call

```
if rh[i][j] and ⋆ then
    skip
else
    call rv := p e;
    sum[i].last := d_j;
    rh[i][j] := true
```

and each statement **await** $\mathtt{r}_i$ into the statement

```
assume sum[i].first = l;
l := sum[i].last;
let guess = ⋆ in
    sum[i] := { first = guess; last = guess };
rh[i] := [ false; ..; false ]
```

where we assume each $\mathtt{d}_j \in \mathsf{Rets}$ is given by an expression in which `rv` is a free variable. Note that in the translation P_{s}, the value of `sum[`i`].first` is forced to be equal to the current local valuation at the point where the **await** $\mathtt{r}_i$ statement is encountered; in this way, even though each return-value handler is eagerly applied to an initially guessed local valuation, we are retroactively guaranteed that the valuation is a valid one.

By the simple-handler hypothesis, return-value handlers are commutative among each other, and idempotent. As the order of applying different return-value handlers at an **await** statement is thus unimportant, we may eagerly apply each handler at the program point where its task is posted. To simulate the various orderings of multiple occurrences of the same handler, we must simply decide whether the last-encountered instance is to be the finally applied one; idempotence guarantees that considering only the last suffices. We use the variable `rh[`i`][`j`]` to ensure that at least one instance of $\mathtt{d}_j$ has been applied if one was posted.

It is not hard to see that the state-reachability problem for the resulting sequential program P_{s} is equivalent to state-reachability of P. Furthermore, the size of P_{s} is polynomial in P, while the number of variables in P_{s} increases by nk. Thus our state-reachability problem is PTIME-complete for fixed n and k, since the state-reachability problem for sequential programs is [7, 33]. When the numbers n of regions or k of handlers is not fixed, this state-reachability problem becomes EXPTIME-complete, due to the logarithmic encoding of the program values into the $n * k$ extra variables.

Although we have showed the result only for local-scope programs, the extension to un-scoped programs is straightforward. Essentially, each task must also *return* their `sum` and `rh` variables, which are composed with the calling-task's `sum` and `rh` variables, e.g., by checking the caller's `sum[`i`].last` is equal to the callee's `sum[`i`].first`, and updating the caller's `sum[`i`].last` to the callee's `sum[`i`].last`. □ □

Theorem 6. *The state-reachability problem for local-scope multi-wait single-region finite-value programs is NP-complete.*

We show NP-hardness in Lemma 11 by a reduction from circuit satisfiability [28], and membership in NP in Lemma 12 by a procedure which solves a polynomial number of polynomial-sized integer linear programs.

Lemma 11. *The circuit satisfiability problem [28] is polynomial-time reducible to the state-reachability problem for local-scope multi-wait single-region finite-value programs.*

Proof. Let C be a Boolean circuit with wires W, gates G, inputs I, and an output wire $w_0 \in W$. Without loss of generality, assume that each gate $g \in G$ is connected to exactly two input wires and

two output wires, and that each input $h \in I$ is connected to exactly two wires. The circuit satisfiability problem asks if there exists a valuation to the inputs I which makes the value of write w_0 true.

We construct a multi-wait single-region finite-value program P_C as follows. Let `Wire` be the type defined as

```
type Wire = { id: W, active: 𝔹, val: 𝔹 }
```

and define a procedure for writing a value to a wire,

```
proc set (var id: W, val: 𝔹)
    var fst, snd: Wire

    if ⋆ then
        fst.id := id;
        fst.val := val;
        fst.active := true
    else
        snd.id := id;
        snd.val := val;
        snd.active := true;

    return (fst,snd)
```

which takes a value to be written and returns two output wires (one of which is written to), and a procedure for reading the value of a wire,

```
proc get (var id: W, fst, snd: Wire)
    var val: 𝔹

    if ⋆ then
        assume fst.active and fst.id = id;
        val := fst.val;
        fst.active := false;
    else
        assume snd.active and snd.id = id;
        val := snd.val;
        snd.active := false;

    return (val,fst,snd).
```

which takes two wires `fst` and `snd`, reads a value from one of them, and returns the same (but mutated) wires, along with the value read. For each gate $g \in G$ connected to input wires w_1, w_2, output wires w_3, w_4, and computing a function $f : \mathbb{B} \to \mathbb{B}$, we declare a procedure,

```
proc p_g (var val: 𝔹)
    var fst0, snd0, fst, snd: Wire
    var a, b, c: 𝔹

    fst := fst0;
    snd := snd0;
    call (a,fst,snd) := get(w_1,fst,snd);
    call (b,fst,snd) := get(w_2,fst,snd);
    c := f(a,b);
    assume c = val;
    return (fst0,snd0,fst,snd).
```

Finally, the initial procedure posts two instances of `set` per input $h \in I$, and two instances of `set` per gate $g \in G$, along with one instance of $\mathtt{p}_g$, then waits until every task is consumed in some sequence,

```
proc init ()
    var fst, snd: Wire
    var val: 𝔹
    var done: 𝔹

    fst.active := false;
    snd.active := false;
    done := false;

    // input h_1
    val := ⋆;
    post r ← set(w_{h_1,1},val) d_w;
    post r ← set(w_{h_1,2},val) d_w;

    // input h_2
    val := ⋆;
    post r ← set(w_{h_2,1},val) d_w;
    post r ← set(w_{h_2,2},val) d_w;

    ...;

    // gate g_1
    val := ⋆;
    post r ← p_{g_1}(val) d_rw;
    post r ← set(w_{g_1,3},val) d_w;
    post r ← set(w_{g_1,4},val) d_w;

    ...;

    await r;
    done := true,
```

where $w_{h_i,j}$ (resp., $w_{g_i,j}$) denotes the jth wire of input h_i (resp., gate g_i). The return handler $\mathtt{d}_w(\mathtt{f},\mathtt{s})$ assigns `f` to `fst` and `s` to `snd`, and $\mathtt{d}_{rw}(\mathtt{f0},\mathtt{s0},\mathtt{f},\mathtt{s})$ ensures `f0 = fst` and `s0 = snd`[3], and assigns `f` to `fst` and `s` to `snd`.

The program P_C simulates C by evaluating each gate $g \in G$ one-by-one at the **await** statement, based on an ordering such that g's input wires are active exactly when the task of procedure $\mathtt{p}_g$ is consumed. This is possible since the *setting* of each input wire $w \in W$ of g is also a pending task (of procedure `set`), which in turn can be scheduled immediately before $\mathtt{p}_g$. Such an execution is guaranteed to be explored since every possible ordering of pending task consumption is considered at the **await** statement.

We then ask if there is a reachable state in which

```
fst.id = w_0 and fst.val = true and done =
    true
```

and if so, it must be the case that C is satisfiable. Inversely, if C is satisfiable then there must exist an execution to … since every possible circuit evaluation order is considered. □ □

Lemma 12. *The state-reachability problem for local-scope multi-wait single-region finite-value programs P over values* Vals *and return-value handlers* Rets *is reducible to solving a $\mathcal{O}(|P|^3 \cdot |\mathsf{Vals}|^3 \cdot |\mathsf{Rets}|)$-length series of integer linear programs, each of size $\mathcal{O}(|P|^5 \cdot |\mathsf{Vals}|^5 \cdot |\mathsf{Rets}|)$.*

Proof. Let P be a program with finite sets of procedures Procs, values Vals, and return-value handlers Rets, and let $\ell \in \mathsf{Vals}$ be a target reachable value from an initial configuration $c_0 = \langle \ell_0, s, \emptyset \rangle$. We construct two sequences $\mathcal{A}^{\mathrm{s}}_1, \mathcal{A}^{\mathrm{s}}_2, \ldots$ and $\mathcal{A}^{\mathrm{t}}_1, \mathcal{A}^{\mathrm{t}}_2, \ldots$ of finite-state automata. Intuitively, each $\mathcal{A}^{\mathrm{s}}_i$ will be a *sync-point summary automaton*, characterizing pairs of program states reachable between two consecutive **await** statements; each $\mathcal{A}^{\mathrm{t}}_i$ will be a *task-summary automaton*, characterizing pairs of program states reachable between the entry and the exit of each task's procedure.

[3] We can block executions by allowing return handlers to be partial functions.

To construct the state-space Q of these automata, we assume an explicit representation Locs of P's control locations (and the control locations of the initial statement s). Furthermore, we assume each control location $m \in \mathsf{Locs}$ contains a component $\mathsf{val}(m)$ containing a local-variable valuation, in addition to the program-counter location $\mathsf{ctrl}(m)$. We say a control location $m \in \mathsf{Locs}$ is a *sync(hronization) point of* $p \in \mathsf{Procs}$ if $\mathsf{ctrl}(m)$ is either the entry location of p, or $\mathsf{ctrl}(m)$ is the target of an $\mathbf{await}$ statement. A pair $m_1, m_2 \in \mathsf{Locs}$ of consecutive[4] synchronization points of p is called a *sync(hronization)-point pair of* p. Note that in a multi-wait single-region program, between each sync-point pair $\langle m_1, m_2 \rangle$, execution proceeds from m_1 first by a phase of sequential and task-posting transitions until an $\mathbf{await}$ statement incident on m_2 is reached. At this point a second phase of consuming the posted tasks begins; execution proceeds to m_2 only after all of the tasks posted during the first phase have been consumed. Furthermore, we may assume, without loss of generality, that each procedure ends with a sync point, i.e., the final statement is $\mathbf{await;\ return}\ v$, for some $v \in \mathsf{Vals}$, since P is a local-scope program.

In the following, we denote the language of an automaton $\mathcal{A} = \langle Q, \Sigma, \delta \rangle$ with states Q, alphabet Σ, and transitions $\delta \subseteq Q \times \Sigma \times Q$ between two states $q_0, q_f \in Q$ by $\mathcal{A}(q_0, q_f) \stackrel{\text{def}}{=} \{w \in \Sigma^* : q_0 \Rightarrow^w_{\delta^*} q_f\}$, i.e., $\mathcal{A}(q_0, q_f)$ is the set of words labelling runs of $\mathcal{A}$ from q_0 to q_f.

Let $Q \stackrel{\text{def}}{=} \mathsf{Locs}$. We model task-posting by labeling the transitions of the automata by tasks, and define the alphabet $\Sigma \stackrel{\text{def}}{=} \mathsf{Tasks} \cup \{\varepsilon\}$. The initial task-summary automaton is $\mathcal{A}^{\mathrm{t}}_0 = \langle Q, \{\varepsilon\}, \emptyset \rangle$ with states Q, alphabet $\{\varepsilon\}$, and the empty set $\emptyset$ of transitions.

Construction of $\mathcal{A}^{\mathrm{s}}_i$ For $i > 0$, we define the i^{th} *sync-point summary automaton*, characterizing state-reachability between sync-point pairs, as

$$\mathcal{A}^{\mathrm{s}}_i = \left\langle Q \cup \bar{Q}, \Sigma, \delta^{\mathrm{s}}_i \right\rangle,$$

where the states Q and $\bar{Q} \stackrel{\text{def}}{=} \{\bar{q} : q \in Q\}$ correspond, resp., to control locations of the first (task-posting) and second (task-consuming) phases, and the transitions $\delta^{\mathrm{s}}_i = \delta^+ \uplus \delta' \uplus \delta^-_i$ are partitioned into first-phase transitions $\delta^+ \subseteq Q \times \Sigma \times Q$, phase-change transitions $\delta' = \{\langle q, \varepsilon, \bar{q} \rangle : q \in Q\}$, and second-phase transitions $\delta^-_i \subseteq \bar{Q} \times \Sigma \times \bar{Q}$.

The relation δ^+ is given directly by the sequential and task-posting transitions of the input program. The relation δ^-_i contains a transition $\langle \bar{q}, a, \bar{q}' \rangle$ summarizing the consumption of the task $a = \langle p, v, d \rangle$ if and only if there exists $q_0, q_f \in Q$ and $v' \in \mathsf{Vals}$ such that

- $\mathsf{ctrl}(q_0)$ is the entry location of p and $\mathsf{val}(q_0) = v$,
- $\mathsf{ctrl}(q_f)$ is an exit location of p and $\mathsf{val}(q_f) = v'$,
- $\mathcal{A}^{\mathrm{t}}_{i-1}(q_0, q_f)$ is non-empty,
- $\mathsf{ctrl}(q) = \mathsf{ctrl}(q')$, and
- $\mathsf{val}(q') = d(v')(\mathsf{val}(q))$; i.e., $\mathsf{val}(q')$ is the value resulting from executing the return-value handler assignment `1 := d(v')` with the value $\mathsf{val}(q)$.

In other words, $\langle \bar{q}, a, \bar{q}' \rangle$ summarizes the effect of consuming task a, based on $\mathcal{A}^{\mathrm{t}}_{i-1}$'s summarization of a, including the local-variable update due to the return-value handler d. In this way, the possible behaviors between sync-points are computed using the thus-far computed (entire) behaviors of each posted task.

Note that not every word of $\mathcal{A}^{\mathrm{s}}_i(q_0, \bar{q}_f)$ represents a valid computation between two consecutive sync points q_0 and q_f, since $\mathcal{A}^{\mathrm{s}}_i$ cannot ensure that each task posted in the first phase is consumed in the second. For $q_0, q_f \in Q$, we say a word $w_1 w_2 \in \mathcal{A}^{\mathrm{s}}_i(q_0, \bar{q}_f)$ is *balanced* if and only if $\Pi(w_1) = \Pi(w_2)$ and there exists $q \in Q$ such that $w_1 \in \mathcal{A}^{\mathrm{s}}_i(q_0, q)$ and $w_2 \in \mathcal{A}^{\mathrm{s}}_i(\bar{q}, \bar{q}_f)$. We say $\mathcal{A}^{\mathrm{s}}_i(q_0, \bar{q}_f)$ *has a balanced run* if some word of $\mathcal{A}^{\mathrm{s}}_i(q_0, \bar{q}_f)$ is balanced. For each sync-point pair $\langle q_0, q_f \rangle$, we can decide whether $\mathcal{A}^{\mathrm{s}}_i(q_0, \bar{q}_f)$ has a balanced run by integer linear programming. In particular, given $\mathcal{A}^{\mathrm{s}}_i$ and $\langle q_0, q_f \rangle$, we construct a integer linear program $\Phi^{\mathrm{s}}_i(q_0, q_f)$ which has a positive integer solution exactly when $\mathcal{A}^{\mathrm{s}}_i(q_0, \bar{q}_f)$ has a balanced run.

Construction of Φ^{s}_i Given the sync-point summary automaton $\mathcal{A}^{\mathrm{s}}_i$ and sync-point pair $q_0, q_f \in Q$, we construct an ILP, denoted $\Phi^{\mathrm{s}}_i(q_0, q_f)$. Fix (finite) enumerations $q_1 q_2 \ldots$, $a_1 a_2 \ldots$, and $d_1 d_2 \ldots$ of the states, symbols, and transitions, resp., of $\mathcal{A}^{\mathrm{s}}_i$; i.e., $Q = \{q_1, q_2, \ldots\}$, $\Sigma = \{a_1, a_2, \ldots\}$, and $\delta = \{d_1, d_2, \ldots\}$. Additionally, assume that $d_j = \langle q_j, \varepsilon, \bar{q}_j \rangle \in \delta'$ for each $q_j \in Q$. We define $\Phi^{\mathrm{s}}_i(q_0, q_f)$ as an integer linear program with $|\delta^{\mathrm{s}}_i|$ *transition occurrence variables*, one d_j for each transition $d_j \in \delta^{\mathrm{s}}_i$, and $|\Sigma| - 1$ *task counter variables*, one a_j for each $a_j \in \Sigma \setminus \{\varepsilon\}$. Then $\Phi^{\mathrm{s}}_i(q_0, q_f)$ contains the following constraints

$$\text{for each } q_k \in Q, \qquad \left(\mathrm{d}_k + \sum_{d_j \in \delta^+(q_k, \cdot, \cdot)} \mathrm{d}_j - \sum_{d_j \in \delta^+(\cdot, \cdot, q_k)} \mathrm{d}_j \right) = \begin{cases} 0 & \text{when } q_k \neq \\ 1 & \text{when } q_k = \end{cases}$$

to ensure each state in the first phase is exited once per entry (except q_0, which is exited one extra time),

$$\text{for each } \bar{q}_k \in \bar{Q}, \qquad \left(\mathrm{d}_k + \sum_{d_j \in \delta^-_i(\cdot, \cdot, \bar{q}_k)} \mathrm{d}_j - \sum_{d_j \in \delta^-_i(\bar{q}_k, \cdot, \cdot)} \mathrm{d}_j \right) = \begin{cases} 0 & \text{when } q_k \neq \\ 1 & \text{when } q_k = \end{cases}$$

to ensure each state in the second phase is exited once per entry (except q_f, which is entered one extra time),

$$\left(\sum_{d_j \in \delta'} \mathrm{d}_j \right) = 1$$

to ensure at most one inter-phase transition is taken, and

$$\text{for each } a_k \in \Sigma, \qquad \left(\sum_{d_j \in \delta^+(\cdot, a_k, \cdot)} \mathrm{d}_j \right) = \mathrm{a}_k = \left(\sum_{d_j \in \delta^-_i(\cdot, a_k, \cdot)} \mathrm{d}_j \right)$$

to ensure that the number of occurrences of each a_k in the first phase is equal to the number of occurrence in the second phase. (Note that the a_j variables are not strictly necessary; they are added only for clarity.) Supposing $d_{j_1} d_{j_2} \ldots$ is a connected sequence of transitions through $\mathcal{A}^{\mathrm{s}}_i$, a corresponding solution to the given set of constraints would set the variables $\mathrm{d}_{j_1}, \mathrm{d}_{j_2}, \ldots$ to positive (non-zero) values corresponding to the number of times each transition is taken in $\mathcal{A}^{\mathrm{s}}_i$. However, supposing there are loops in $\mathcal{A}^{\mathrm{s}}_i$ which are not connected to any of the selected transitions, the given constraints do not prohibit solutions which take each transition of these loops an arbitrary number of times. This is a standard issue with encoding automaton traces which can be addressed by adding a polynomial number of constraints [**?**] to $\Phi^{\mathrm{s}}_i(q_0, q_f)$.

Proposition 5. *$\mathcal{A}^{\mathrm{s}}_i(q_0, \bar{q}_f)$ has a balanced run if and only if $\Phi^{\mathrm{s}}_i(q_0, q_f)$ has a positive integer solution.*

Note that the size of Φ^{s}_i is bounded by $\mathcal{O}(|P|^5 \cdot |\mathsf{Vals}|^5 \cdot |\mathsf{Rets}|)$, since each of $\mathcal{O}(|Q|^2)$-many programs $\Phi^{\mathrm{s}}_i(q, q')$ contains $\mathcal{O}(|\delta^{\mathrm{s}}_i|) = \mathcal{O}(|Q|^2 \cdot |\Sigma|)$ variables and $\mathcal{O}(|Q| + |\Sigma|)$ constraints, where $\mathcal{O}(|Q|) = \mathcal{O}(|P| \cdot |\mathsf{Vals}|)$ and $\mathcal{O}(|\Sigma|) = \mathcal{O}(|P| \cdot |\mathsf{Vals}| \cdot |\mathsf{Rets}|)$.

[4] In some standard representation of p's static control-flow, e.g., in the control-flow graph of p.

Construction of $\mathcal{A}_i^{\mathrm{t}}$ For $i > 0$ we define the i^{th} task-summary automaton, characterizing state-reachability among synchronization points, as

$$\mathcal{A}_i^{\mathrm{t}} = \langle Q, \{\varepsilon\}, \delta_i^{\mathrm{t}} \rangle$$

such that $\langle q, \varepsilon, q' \rangle \in \delta_i^{\mathrm{t}}$ if and only if $\langle q, q' \rangle$ is a sync-point pair, and $\mathcal{A}_i^{\mathrm{s}}(q, q')$ has a balanced run.

Note that there are only finitely-many transitions which can be added over the entire $\mathcal{A}_i^{\mathrm{s}}$ and $\mathcal{A}_i^{\mathrm{t}}$ sequence. It follows that there exists a fixed-point $m \in \mathbb{N}$ of this sequence, and it is not hard to see that $\mathcal{A}_m^{\mathrm{s}}$ and $\mathcal{A}_m^{\mathrm{t}}$ capture every behavior of the input program P.

Proposition 6. *A synchronization point q_f of the initial task is reachable from an initial control location q_0 if and only if $\mathcal{A}_m^{\mathrm{t}}(q_0, q_f)$ is non-empty.*

Though we consider here only state-reachability to a synchronization point contained in the initial task for simplicity, Proposition 6 can indeed be extended to arbitrary control locations of arbitrary tasks. As the set of possible added transitions is bounded by $\mathcal{O}(|Q|^2 \cdot |\Sigma|) = \mathcal{O}(|P|^3 \cdot |\mathsf{Vals}|^3 \cdot |\mathsf{Rets}|)$, our procedure is guaranteed to terminate in polynomial-time. □ □

Theorem 7. *The state-reachability problem for multi-wait finite-value programs is polynomial-time equivalent to the configuration-reachability problem for vector addition systems.*

We demonstrate this equivalence by a polynomial-time reduction in each direction. Though VAS configuration-reachability has been shown decidable [**?**], the precise asymptotic complexity is not known; VAS state-reachability gives an EXPSPACE lower-bound.

Lemma 13. *The configuration-reachability problem for vector addition systems is reducible to the state-reachability problem for multi-wait finite-value programs.*

Proof. Let $\mathcal{A} = \langle Q, \delta \rangle$ be a k-dimension vector addition system with $\delta = \{d_1, \ldots, d_n\}$, and without loss of generality suppose $|Q| = 1$. To decide whether $\mathcal{A}$ has a run from some $\vec{n}_0 \in \mathbb{N}^k$ to $\mathbf{0} \in \mathbb{N}^k$ (again, without loss of generality), we construct a multi-wait program $P_{\mathcal{A}}$ and a local valuation ℓ which is reachable in $P_{\mathcal{A}}$ if and only if $\mathcal{A}(\vec{n}_0, \mathbf{0}) \neq \emptyset$.

We will construct $P_{\mathcal{A}}$ such that the number of pending tasks in a configuration is equal to the sum of vector components in a corresponding configuration of $\mathcal{A}$. We then simulate each step of $\mathcal{A}$, which subtracts $\vec{n}_1 \in \mathbb{N}^k$ and adds $\vec{n}_2 \in \mathbb{N}^k$, by consuming $\sum_i \vec{n}_1(i)$ tasks and posting $\sum_i \vec{n}_2(i)$ tasks, while ensuring each task consumed (resp., posted) corresponds to a subtraction (resp., addition) to the correct vector-component.

For each transition $d_i = \langle q, \vec{n}_1, \vec{n}_2, q \rangle$ we define the sequence $\sigma_i \in [1, k]^*$ of counter decrements as

$$\sigma_i \stackrel{\text{def}}{=} \underbrace{11\ldots11}_{\vec{n}_1(1)\text{ times}}\ \underbrace{22\ldots22}_{\vec{n}_1(2)\text{ times}}\ \cdots\ \underbrace{kk\ldots kk}_{\vec{n}_1(k)\text{ times}}.$$

We assume, without loss of generality, that each transition has a non-zero decrement vector, i.e., $\vec{n}_1 \neq \mathbf{0}$ and thus $|\sigma_i| > 0$. We will use return-value handlers to ensure that a $|\sigma_i|$-length sequence of consecutively-consumed tasks corresponds the decrement of transition d_i. For each $j \in \{1, \ldots, |\sigma_i|\}$, let $\mathtt{d}_{i,j}(v)$ be the return-value handler defined by

```
assume cur_tx = i;
assume cur_pos = j;
if cur_pos = |σ_i| then
    assume v = true;
    cur_tx := ⋆;
    cur_pos := 1
else
    assume v = false;
    cur_pos := cur_pos + 1,
```

which checks that consuming a given task corresponds to a decrement (by one) of the $\sigma_i(j)^{\text{th}}$ component of the decrement vector of $d_i \in \delta$. For each increment vector $\vec{n}$ (i.e., $\langle q, \vec{n}_1, \vec{n}, q \rangle \in \delta$ for some $\vec{n}_1 \in \mathbb{N}^k$), or initial vector $\vec{n} = \vec{n}_0$, we declare the procedure

```
proc inc_n⃗ ()
    for var idx := 1 to k do
        for var cnt := 1 to n⃗(idx) do
            let tx = ⋆
            and pos = ⋆ in
            assume σ_tx(pos) = idx;
            post r ← p_tx ⋆ d_tx,pos.
```

which posts $\vec{n}(m)$ tasks for each $m \in \{1, \ldots, k\}$, to be consumed later by arbitrary positions j of the decrement sequences σ_i (since `pos` is assigned $\star$) of arbitrary transitions d_i (since `tx` is assigned $\star$) such that $\sigma_i(j) = m$—this ensures that the subsequent consumption of a task with handler $\mathtt{d}_{i,j}$ corresponds to decrementing the m^{th} component of $\vec{n}$. To perform the increment of transition $d_i \in \delta$ by vector $\vec{n}_2$, we declare the procedure $\mathtt{p}_i$, which non-deterministically calls $\mathtt{inc}_{\vec{n}_2}$, as

```
proc p_i ()
    if ⋆ then
        call inc_n⃗2 ();
        return true
    else
        return false.
```

Note that the Boolean return value is used by the attached return-value handler $\mathtt{d}_{i,j}$ (for some $j \in \{1, \ldots, |\sigma_i|\}$) to ensure that the increment is only performed once per transition d_i, by the last-consumed task in the $|\sigma_i|$-length sequence.

Finally, the initial procedure `main` simply adds tasks corresponding to the initial vector $\vec{n}_0$ to an initially-empty region container, then loops until every task has been consumed:

```
proc main ()
    var cur_tx :=  ⋆ ;
    var cur_pos = 1;
    var empty := false;
    call inc_n⃗0 ();
    await r;

    // check: is this point reachable?
    empty := true;
    return.
```

Checking that $P_{\mathcal{A}}$ faithfully simulates $\mathcal{A}$ is easily done by noticing the correspondence between configurations $q\vec{n}$ of $\mathcal{A}$ and configurations of $P_{\mathcal{A}}$ with $\sum_i \vec{n}$ pending tasks. Since `empty = true` is only reachable when there are no pending tasks, reachability to `empty = true` implies $\mathbf{0}$ is reachable in $\mathcal{A}$. Furthermore, if $\mathbf{0}$ is reachable in $\mathcal{A}$, a run of $P_{\mathcal{A}}$ will eventually proceed past the **await** statement without pending tasks, setting `empty = true`.

Proposition 7. $\mathcal{A}(\vec{n}_0, \mathbf{0}) \neq \emptyset$ *if and only if* `empty = true` *is reachable in* $P_{\mathcal{A}}$.

Since the size of $P_{\mathcal{A}}$ is polynomial in $\mathcal{A}$, we have a polynomial-time reduction for deciding configuration-reachability in $\mathcal{A}$. □ □

Lemma 14. *The state-reachability problem for multi-wait finite-value programs is reducible to the configuration-reachability problem for vector addition systems.*

Proof. Let P be a multi-wait program, and ℓ a local-state valuation. To decide whether ℓ is reachable in P we will perform a polynomial number (in the size of P) of reachability queries on polynomial-sized (in the size of P) vector addition systems. To simplify the proof, we will assume, without loss of generality, that ℓ is reachable only in the initial procedure of P, immediately following a synchronization point.

Assume P is given by a transition system $\hookrightarrow_P \subseteq \mathsf{Locs} \times \mathsf{Stmts} \times \mathsf{Locs}$ as in Section B.1, and for $p \in \mathsf{Procs}$ let $\mathsf{Syncs}_p \subseteq \mathsf{Locs}$ be the set of synchronization-point targets, or entry points, of procedure p, and $\mathsf{Syncs} = \bigcup_p \mathsf{Syncs}_p$; i.e.,

$$\{m_2 \in \mathsf{Locs} : \langle m_1, \mathbf{await}\ r, m_2\rangle \in \hookrightarrow_P \text{ or } \mathsf{ctrl}(m_2) = \ldots\}.$$

We a construct sequence of graphs $G_0 G_1 \ldots$ and a sequence of vector addition systems $\mathcal{A}_1 \mathcal{A}_2 \ldots$ inductively, to summarize executions of P. Each graph G_i with nodes Syncs contains an edge between each pair of synchronization points which have been found to be pairwise-reachable. The initial graph $G_0 \stackrel{\text{def}}{=} \langle \mathsf{Syncs}, \emptyset\rangle$ contains no edges. Each vector addition system $\mathcal{A}_i$ is constructed using the summaries of G_{i-1}, and each graph G_i for $i > 0$ is constructed using configuration-reachability queries of $\mathcal{A}_i$.

Fix an alphabet $\Sigma = \mathsf{Tasks}$. At each step i, we construct $\mathcal{A}_i$ as follows. First, we notice that the sequence of tasks posted between two consecutive synchronization points of a procedure forms a context-free language over Σ. ...

ToDo: pick it up here.

The task-posting grammar Let $\mathcal{G}_i = \left\langle \mathsf{Locs}^2, \Sigma, \hookrightarrow_{\mathcal{G}_i}\right\rangle$ be the context-free grammar over alphabet $\Sigma = \mathsf{Regs} \times \mathsf{Tasks}$ with variables Locs^2 and the following productions:

1. For each $m_1, m_2, m_3 \in \mathsf{Locs}$ such that P contains an intra-task transition from m_1 to m_2,
$$\langle m_1, m_3\rangle \hookrightarrow_{\mathcal{G}_i} \langle m_2, m_3\rangle .$$
2. For each $m_1, m_2, m_3, m_4, m_5 \in \mathsf{Locs}$ such that P contains a transition $\mathbf{call}$ `l := p v` from m_1 to m_2,
$$\langle m_1, m_3\rangle \hookrightarrow_{\mathcal{G}_i} \langle m_2, m_3\rangle \langle m_4, m_5\rangle ,$$
such that m_3 is the entry location of $p \in \mathsf{Procs}$ with local valuation $v \in \mathsf{Vals}$, and m_4 is the exit location of p which returns the value $v' \in \mathsf{Vals}$ such that `l := v'` updates the local valuation of m_1 to that of m_2.
3. For each location $m \in \mathsf{Locs}$,
$$\langle m, m\rangle \hookrightarrow_{\mathcal{G}_i} \varepsilon.$$
4. For each $m_1, m_2, m_3 \in \mathsf{Locs}$ such that P contains a transition $\mathbf{post}\ r \leftarrow p\ v\ d$ from m_1 to m_2,
$$\langle m_1, m_3\rangle \hookrightarrow_{\mathcal{G}_i} \langle m_2, m_3\rangle \langle r, p, v, d\rangle .$$
5. Finally, for each $m_1, m_2, m_3 \in \mathsf{Locs}$ such that m_1 and m_2 are synchronization point targets, and $m_1 \hookrightarrow_{\mathcal{A}_i} m_2$,
$$\langle m_1, m_3\rangle \hookrightarrow_{\mathcal{G}_i} \langle m_2, m_3\rangle .$$

Intuitively, derivations of $\mathcal{G}_i(\langle m_1, m_2\rangle)$ correspond to sequential executions between entry and exit points, resp., m_1 and m_2, of some procedure $p \in \mathsf{Procs}$, including recursive calls to other procedures (productions groups 1, 2, and 3). The words generated by $\mathcal{G}_i(m_1, m_2)$ form a context-free language summarizing the tasks posted along the execution of p, and the procedures it calls, which are still pending when p returns (production group 4). Note that these pending tasks can only have been posted between the last synchronization point of p before p returns. For this reason we allow $\mathcal{G}_i$ to jump from the entry point of p to the target of any $\mathbf{await}$ transition in p (production group 5), and do not allow $\mathcal{G}_i$ to otherwise proceed past an $\mathbf{await}$ transition. The number of productions of $\mathcal{G}_i$ is $\mathcal{O}(|\mathsf{Locs}|^5)$.

ToDo: clean up the following proposition.

Proposition 8. *For all $m_1, m_2 \in \mathsf{Locs}$, $\mathcal{G}_i(\langle m_1, m_2\rangle)$ is Parikh-equivalent to the region map reached by an execution of P from m_1 to m_2.*

The task-posting VAS

Proposition 9 ([14])**.** *Given any context-free grammar $\mathcal{G}$ with variables V and $v \in V$, there exists a vector addition system $\mathcal{A}$ of size polynomial in $\mathcal{G}$ with states Q and $Q_0, Q_f \subseteq Q$ such that $\mathcal{A}(Q_0\mathbf{0}, Q_f\vec{n}) \neq \emptyset$ if and only if $\vec{n} \in \Pi(\mathcal{G}(v))$.*

For each $v \in \mathsf{Locs}^2$, let $\mathcal{A}^i_{\mathcal{G}}(Q^v_1, Q^v_2)$ be the vector addition system for $\mathcal{G}_i(v)$ given by Proposition 9.

The task-post-and-consume system We construct a vector addition system $\mathcal{A}_i = \langle Q, \hookrightarrow_{\mathcal{A}_i}\rangle$ such that $Q = \mathsf{Locs} \cup Q^i_{\mathcal{G}} \cup (Q^i_{\mathcal{G}} \times \mathsf{Locs} \times \mathsf{Rets})$ with the following transitions:

1. For each $m \in \mathsf{Locs}$, $d \in \mathsf{Rets}$, and $q_1, q_2 \in Q^i_{\mathcal{G}}$ such that $q_1 \xhookrightarrow{\vec{n}_1\vec{n}_2}_{\mathcal{A}^i_{\mathcal{G}}} q_2$
$$\langle q_1, m, d\rangle \xhookrightarrow{\vec{n}_1\vec{n}_2}_{\mathcal{A}_i} \langle q_2, m, d\rangle .$$
2. For each sync-point pair $\langle m_1, m_2\rangle \in \mathsf{Locs}^2$ of P and $Q_1, Q_2 \subseteq Q^i_{\mathcal{G}}$ such that $\mathcal{A}^i_{\mathcal{G}}(Q_1, Q_2)$ is the ... of $\mathcal{G}(\langle m_1, m_2\rangle)$, and $q_1 \in Q_1$, $q_2 \in Q_2$,
$$m_1 \xhookrightarrow{\mathbf{00}}_{\mathcal{A}_i} q_1 \quad \text{and} \quad q_2 \xhookrightarrow{\mathbf{00}}_{\mathcal{A}_i} m_2.$$
3. For each $m_1, m_2, m_3, m_4 \in \mathsf{Locs}$, $t \in \mathsf{Regs} \times \mathsf{Tasks}$, and $v' \in \mathsf{Vals}$ such that
   - $t = \langle r, p, v, d\rangle$,
   - $\mathsf{ctrl}(m_1) = \mathsf{ctrl}(m_2)$ and $\mathsf{val}(m_2) = d(v')(\mathsf{val}(m_1))$,
   - m_3 is the entry location of p with valuation v,
   - m_4 is the exit location of p with valuation v',
   - $\mathcal{A}^i_{\mathcal{G}}(Q_1, Q_2)$ is the ... of $\mathcal{G}(\langle m_3, m_4\rangle)$,
   - $q_1 \in Q_1$ and $q_2 \in Q_2$,

   $$m_1 \xhookrightarrow{\vec{n}_t\mathbf{0}}_{\mathcal{A}_i} \langle q_1, \langle m_1, d\rangle\rangle \quad \text{and} \quad \langle q_2, \langle m_1, d\rangle\rangle \xhookrightarrow{\mathbf{00}}_{\mathcal{A}_i} m_2.$$

Intuitively, runs of $\mathcal{A}_i(m_1, m_2)$ correspond to executions between two adjacent synchronization points m_1 and m_2 of some procedure $p_0 \in \mathsf{Procs}$, which begin by a derivation of $\mathcal{G}_i(\langle m_1, m_2\rangle)$, summarizing the tasks posted between m_1 and m_2 (transition groups 1 and 2). Technically, $\mathcal{A}_i$ transfers control at location m_1 to the VAS $\mathcal{A}^i_{\mathcal{G}}$ which computes the Parikh image of $\mathcal{G}_i(\langle m_1, m_2\rangle)$, and returns control to m_2 when $\mathcal{A}^i_{\mathcal{G}}$ finishes. In the second phase (transition group 3), $\mathcal{A}_i$ repeatedly selects pending tasks $\langle p, v, d\rangle$ with executions from $m_3 \in \mathsf{Locs}$ to $m_4 \in \mathsf{Locs}$, and again transfers control at location m_2 to the VAS $\mathcal{A}^i_{\mathcal{G}}$ which computes the Parikh image of $\mathcal{G}_i(\langle m_3, m_4\rangle)$, and returns control again to m_2, updating p_0's current valuation using the return handler d of the posted task.

Proposition 10. *$\mathcal{A}_i$ computes only executions allowed by P.*

Finally, let $G_i = \langle \mathsf{Locs}, E\rangle$ be the graph over nodes Locs which has an edge $\langle m_1, m_2\rangle \in E$ if and only if m_1 and m_2 are adjacent synchronization points in P, and $\mathcal{A}_i(m_1, m_2) \neq \emptyset$.

Proposition 11. *Let $m_1 \in \mathsf{Locs}$ and $m_2 \in \mathsf{Locs}$ be synchronization points of some procedure $p \in \mathsf{Procs}$ of P. Then m_2 is reachable from m_1 in P if and only if m_2 is reachable from m_1 in G_ω.*

Finally, since the number of possible edges to add in the G_i sequence is bounded by Locs^2, our construction terminates in at most Locs^2 steps.

Proposition 12. *The sequence $G_0, G_1, \ldots$ has a fixed-point.*

Since the size of each G_i and $\mathcal{A}_i$ is polynomial in P, and the G_i sequence converges in a polynomial number of steps, we have a polynomial-time reduction for deciding state-reachability in P by configuration-reachability in each $\mathcal{A}_i$. □ □

C. Recursive Vector Addition Systems

In this section we extend the model of vector addition systems with states (VASS) to include recursive edges, summarizing an initialized run on an independent vector valuation, in addition to the usual additive edges.

C.1 Formal Model

Fix $k \in \mathbb{N}$. A *recursive vector addition system (RVASS)* $\mathcal{A} = \langle Q, q_0, \delta \rangle$ *of dimension* k is a finite set Q of states, along with an initial state $q_0 \in Q$, and a finite set $\delta = \delta_1 \uplus \delta_2$ of transitions partitioned into *additive* transitions $\delta_1 \subset Q \times \mathbb{N}^k \times \mathbb{N}^k \times Q$ and *recursive* transitions $\delta_2 \subseteq Q \times Q \times Q \times Q$. We write

$$q \xhookrightarrow{\vec{n}_1 \vec{n}_2} q' \qquad \text{when } \langle q, \vec{n}_1, \vec{n}_2, q' \rangle \in \delta, \text{ and}$$
$$q \xhookrightarrow{q_1 q_2} q' \qquad \text{when } \langle q, q_1, q_2, q' \rangle \in \delta.$$

A *(non-recursive) vector addition system (with states) (VASS)* is a recursive vector addition system $\langle Q, q_0, \delta \rangle$ such that δ contains only additive transitions.

An *(RVASS) frame* $\langle q, \vec{n} \rangle$ is a state $q \in Q$ along with a vector $\vec{n} \in \mathbb{N}^k$, and an *(RVASS) configuration* $c \in (Q \times \mathbb{N}^k)^*$ is a sequence of frames representing a stack of non-recursive sub-computations. The transition relation $\to^{\text{rvas}}$ for recursive vector addition systems is defined in Figure 9. The ADDITIVE rule updates the top frame $\langle q, \vec{n} \rangle$ by subtracting the vector $\vec{n}_1$ from $\vec{n}$, adding the vector $\vec{n}_2$ to the result, and updating the control state to q'. The CALL rule pushes on the frame-stack a new frame $\langle q_1, \mathbf{0} \rangle$ from which the RETURN rule will eventually pop at some point when the control state is q_2; when this happens, the vector $\vec{n}_1$ of the popped frame is added to the vector $\vec{n}_2$ of the frame below. We describe an application of the CALL (resp., RETURN) rule as a *call (resp., return) transition*.

The configuration $\langle q_0, \mathbf{0} \rangle$ is called *initial*, and an *execution of a RVASS $\mathcal{A}$ (from c_0 to c_j)* is a configuration sequence $c_0 c_1 \ldots c_j$ where

- c_0 is initial, and
- $c_i \to^{\text{rvas}} c_{i+1}$ for $0 \leq i < j$.

We say a configuration $\langle q, \vec{n} \rangle\, c$ (alternatively, the frame $\langle q, \vec{n} \rangle$, or the state q) is *reachable in $\mathcal{A}$ (from c_0)* when there exists an execution of $\mathcal{A}$ from c_0 to $\langle q, \vec{n} \rangle\, c$. The *state-reachability problem* is to decide whether a given state q is reachable.

C.2 RVASS State-reachability

Though decidability has been shown for a restricted class of RVASS where the vector-component of the top frame of a configuration is not added to the vector-component of the frame below on a return transition [17]—i.e., where the consequent to the RETURN rule is $\langle q_2, \vec{n}_1 \rangle\, \langle q, \vec{n}_2 \rangle\, c \to \langle q', \vec{n}_2 \rangle\, c$—decidability of state-reachability for RVASS in general has not been shown before. To begin, we recall a few notions and results about non-recursive vector addition systems.

Lemma 15 (Abdulla et al. [1])**.** *The state-reachability problem for (non-recursive) vector addition systems is decidable.*

A *preorder* $\preceq$ is a reflexive and transitive binary relation on a set D. We write $d_1 \prec d_2$ to mean $d_1 \preceq d_2$ and $d_2 \not\preceq d_1$. We say $\preceq$ is *well-founded* when there are no infinite sequences $d_1 \succ d_2 \succ \cdots$. A set D_1 is *canonical* when $d_1, d_2 \in D_1$ implies $d_1 \not\preceq d_2$, and a canonical $D_1 \subseteq D_2$ is a *minor set* of D_2 when for every $d_2 \in D_2$ there exists $d_1 \in D_1$ with $d_1 \preceq d_2$. Note that when $\preceq$ is well-founded every $D_1 \subseteq D$ has a (possibly infinite) minor set. A set $D_0 \subseteq D$ is an *ideal* (or *upward closed*) when $d_1 \preceq d_2$ and $d_1 \in D_0$ implies $d_2 \in D_0$. The *upward closure* of a set $D_1 \subseteq D$ is the ideal $D_1{\uparrow} = \{d \in D : \exists d_1 \in D_1.\ d_1 \preceq d\}$ generated by D_1.

A preorder $\preceq$ is a *well-ordering* if for every infinite sequence $d_1 d_2 \ldots \in D^\omega$ there exists $i < j$ such that $d_i \preceq d_j$. Note that when $\preceq$ is a well-ordering, every upward-closed set $D_1 \subseteq D$ has a finite minor set D_0 (and $D_0{\uparrow} = D_1$).

The set of immediate predecessors of a configuration set C of $\mathcal{A}$ is denoted $\text{pre}_{\mathcal{A}}(C) = \{c : \exists c' \in C.\ c \to^{\text{rvas}}_{\mathcal{A}} c'\}$. We define $\text{pre}^j_{\mathcal{A}}(C)$ as $j \in \mathbb{N}$ applications of $\text{pre}_{\mathcal{A}}$, e.g., $\text{pre}^3_{\mathcal{A}}(C) = \text{pre}_{\mathcal{A}} \circ \text{pre}_{\mathcal{A}} \circ \text{pre}_{\mathcal{A}}(C)$, and $\text{pre}^*_{\mathcal{A}}(C) = \bigcup_{j \in \mathbb{N}} \text{pre}^j_{\mathcal{A}}(C)$.

Let $\preceq_{\text{VAS}}$ be the preorder on the configurations of non-recursive VASs defined by $\langle q_1, \vec{n}_1 \rangle \preceq_{\text{VAS}} \langle q_2, \vec{n}_2 \rangle$ if and only if $q_1 = q_2$ and $\vec{n}_1 \leq \vec{n}_2$. It is easy to see that $\preceq_{\text{VAS}}$ is a well-ordering on $Q \times \mathbb{N}^k$, and under $\preceq_{\text{VAS}}$, $\text{pre}_{\mathcal{A}}(C)$, $\text{pre}^j_{\mathcal{A}}(C)$ and $\text{pre}^*_{\mathcal{A}}(C)$ are upward closed sets when C is. Since VASs are *monotonic* (or "upward compatible" [11]) w.r.t. $\preceq_{\text{VAS}}$—i.e., $c_1 \to c_2$ and $c_1 \preceq_{\text{VAS}} c_1'$ implies there exists c_2' such that $c_1' \to c_2'$ and $c_2 \preceq_{\text{VAS}} c_2'$—Lemma 15 follows easily by a fixed-point computation of $\text{pre}^*_{\mathcal{A}}$ from the upward closure of a target configuration $\{\langle q, \mathbf{0} \rangle\}{\uparrow}$.

As is the case for recursive programs, every return transition of an RVASS execution has a previously-occurring corresponding call transition. We say a call transition $q\vec{n}\, c \to q_1 \mathbf{0}\, q\vec{n}\, c$ *matches* a return transition $q_2 \vec{n}'\, q\vec{n}\, c \to q'(\vec{n} \oplus \vec{n}')\, c$ when $\langle q, q_1, q_2, q' \rangle \in \delta_2$. For a configuration sequence $h = c_1 c_2 \ldots c_i$ and $0 < j_1 < j_2 < i$, we say a call transition $c_{j_1} \to c_{j_1+1}$ *matches* a return transition $c_{j_2} \to c_{j_2+1}$ *in* h when there exists $j_1 < k_1 < k_2 < j_2$ such that $c_{k_1} \to c_{k_1+1}$ matches $c_{k_2} \to c_{k_2+1}$ in h whenever there exists a call or return transition $c_k \to c_{k+1}$ for some $j_1 < k < j_2$.

Theorem 8. *The state-reachability problem for recrusive vector addition systems is decidable.*

We prove this by a forward fixed-point algorithm on a converging sequence of vector addition systems (VAS). The initial VAS corresponds to the given RVASS without its recursive transitions. Then, in each step, we add additive transitions summarizing the possible recursive subcomputations of the previous VAS.

To simplify our presentation, we assume a few syntactic shortcuts. We will simply write the state q in a context where a frame is expected to mean $\langle q, \mathbf{0} \rangle$, or $q\vec{n}$ to mean $\langle q, \vec{n} \rangle$. Similarly, we will denote a singleton set of configurations by its element, e.g., $\langle q, \mathbf{0} \rangle$, or $q\mathbf{0}$, or simply q, will mean $\{\langle q, \mathbf{0} \rangle\}$ in contexts where a set is expected. Finally, we assume that arguments to the pre operators are implicitly upward closed. Composing these simplifications, we can write, e.g., $\text{pre}_{\mathcal{A}}(q)$ to mean $\text{pre}_{\mathcal{A}}(\{\langle q, \mathbf{0} \rangle\}{\uparrow})$.

ToDo: update RVASS definition without initial state.

Proof. Let $\mathcal{A} = \langle Q, \delta_1 \uplus \delta_2 \rangle$ be an RVASS and $q_0, q_f \in Q$ a pair of states. To decide whether q_f is coverable in $\mathcal{A}$ from $\langle q_0, \mathbf{0} \rangle$, we define a sequence $\mathcal{A}_0 \mathcal{A}_1 \ldots$ of (non-recursive) VASSs characterizing the backward-reachable states of $\mathcal{A}$ from q_f inductively. For $i \in \mathbb{N}$, let $\mathcal{A}_i \stackrel{\text{def}}{=} \langle Q, \delta_1 \cup \delta_2^i \rangle$, where $\delta_2^0 \stackrel{\text{def}}{=} \delta_{\text{R}}$ is initialized with a set of *return transitions*

$$\delta_{\text{R}} \stackrel{\text{def}}{=} \left\{ q_2 \xhookrightarrow{\mathbf{00}} q' : q \xhookrightarrow{q_1 q_2} q' \in \delta_2 \right\},$$

$$\textsc{Additive}\quad \frac{q \xhookrightarrow{\vec{n}_1 \vec{n}_2} q' \qquad \vec{n} \geq \vec{n}_1}{\langle q, \vec{n} \rangle\, c \xrightarrow{\text{rvas}} \langle q', \vec{n} \ominus \vec{n}_1 \oplus \vec{n}_2 \rangle\, c}$$

$$\textsc{Call}\quad \frac{q \xhookrightarrow{q_1 q_2} q'}{\langle q, \vec{n} \rangle\, c \xrightarrow{\text{rvas}} \langle q_1, \mathbf{0} \rangle\, \langle q, \vec{n} \rangle\, c} \qquad \textsc{Return}\quad \frac{q \xhookrightarrow{q_1 q_2} q'}{\langle q_2, \vec{n}_1 \rangle\, \langle q, \vec{n}_2 \rangle\, c \xrightarrow{\text{rvas}} \langle q', \vec{n}_1 \oplus \vec{n}_2 \rangle\, c}$$

Figure 9. The transition relation for recursive vector addition systems. To simplify presentation, we assume that there is at most one recursive transition originating from each state, i.e., for all $q \in Q$, $\left|\delta_2 \cap (\{q\} \times Q^3)\right| \leq 1$. We denote by $\mathbf{0}$ the vector $\langle 0, 0, \ldots, 0 \rangle$, and by $\oplus$ and $\ominus$ the usual vector addition and subtraction operators.

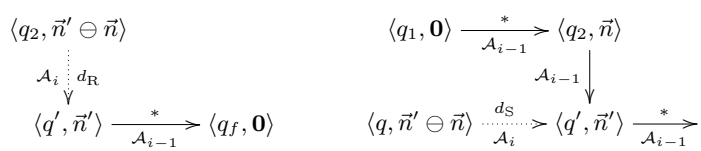

where $d_{\mathrm{R}} = q_2 \xhookrightarrow{\mathbf{0}\vec{n}} q' \in \delta_{\mathrm{R}}^i$ and $d_{\mathrm{S}} = q \xhookrightarrow{\mathbf{0}\vec{n}} q' \in \delta_{\mathrm{S}}^i$

Figure 10. Let $q \xhookrightarrow{q_1 q_2} q'$ be a recursive transition of $\mathcal{A}$. At each step i, we add a return transition d_{R} to $\mathcal{A}_i$ for all $\vec{n} \leq \vec{n}'$ such that $q'\vec{n}'$ can reach $q_f\mathbf{0}$ in $\mathcal{A}_{i-1}$. Similarly, we add a summary transition d_{S} to $\mathcal{A}_i$ for all $\vec{n} \leq \vec{n}'$ such that $q'\vec{n}'$ can reach $q_f\mathbf{0}$, and $q_1\mathbf{0}$ can reach $q_2\vec{n}$ in $\mathcal{A}_{i-1}$.

and $\delta_2^{i+1} \stackrel{\text{def}}{=} \delta_2^i \cup \delta_{\mathrm{S}}^{i+1}$ adds at each step a set of *summary transitions*—to be defined shortly. First, let

$$M_i(q, q', \vec{n}') \stackrel{\text{def}}{=} \text{minor}\, \{\vec{n} \in \mathbb{N}^k : q\vec{n} \in \text{pre}^*_{\mathcal{A}_i}(q'\vec{n}')\}$$

be the minor set of backward-reachable vectors from $q'\vec{n}'$ to q, and

$$C_i(q_1, q_2, N) \stackrel{\text{def}}{=} \text{major}\, \{\vec{n} \in \mathbb{N}^k : \vec{n} \leq N \text{ and } \mathbf{0} \in M_i(q_1, q_2, n)\}$$

be the major set of N-bounded vector contributions from q_1 to q_2. We then define the summary transitions added at step $i+1$ as

$$\delta_{\mathrm{S}}^{i+1} \stackrel{\text{def}}{=} \left\{ q \xhookrightarrow{\mathbf{0}\vec{n}} q' : \begin{array}{l} q \xhookrightarrow{q_1 q_2} q' \in \delta_2 \\ M_i(q', q_f, \mathbf{0}) = N \neq \perp \\ \vec{n} \in C_i(q_1, q_2, N) \end{array} \right\}.$$

(In the above we write $\vec{n} \leq N$ for $\vec{n} \in \mathbb{N}^k$ and $N \subseteq \mathbb{N}^k$ when $\vec{n} \leq \vec{n}'$ for all $\vec{n}' \in N$.) Figure 10 depicts the conditions under which transitions are added.

Proposition 13. *For all $q, q' \in Q$, $\vec{n} \in \mathbb{N}^k$, and $i \in \mathbb{N}$, $M_i(q, q', \vec{n}) \leq M_i(q, q', \vec{n})$.*

ToDo: is that true? maybe incomparables will be added later?

Since the sets of transitions of consecutive $\mathcal{A}_i$s are non-decreasing, the set of (backward) reachable configurations of each $\mathcal{A}_i$ are contained in $\mathcal{A}_{i+1}$.

Proposition 14. $\text{pre}^*_{\mathcal{A}_0}(q_f) \subseteq \text{pre}^*_{\mathcal{A}_1}(q_f) \subseteq \ldots$

Furthermore, since $\leq$ is a well-ordering on $\bigcup_i \delta_2^i$, and each δ_2^{i+1} does not add transitions larger (w.r.t. $\leq$) than those of δ_2^i, our VASS sequence $\mathcal{A}_0\mathcal{A}_1 \ldots$ converges.

Proposition 15. *The sequence $\mathcal{A}_0\mathcal{A}_1 \ldots$ has a fixed-point.*

Finally, we must show that the RVASS state-reachability problem (to q_f) for $\mathcal{A}$ is equivalent to the VAS state-reachability problem (to q_f) for $\mathcal{A}_m$, where f_m is the fixed-point of the sequence $f_0 f_1 \cdots$. Since we construct each $\mathcal{A}_i$ by considering only backward-reachable configurations of $\mathcal{A}$, one direction is simple.

Lemma 16. *For each $i \in \mathbb{N}$, if $\langle q, \vec{n} \rangle \in \text{pre}^*_{\mathcal{A}_i}(q_f)$ then there exists a configuration c such that $\langle q, \vec{n} \rangle\, c \in \text{pre}^*_{\mathcal{A}}(q_f)$.*

To see that every backward-reachable configuration $\langle q, \vec{n} \rangle\, c$ of $\mathcal{A}$ from q_f has a corresponding $\langle q, \vec{n} \rangle$ reached by $\mathcal{A}_m$, we argue that for any sequence of transitions of $\mathcal{A}$ to q_f, $\mathcal{A}_m$ has added return and summary transitions to simulate, respectively, unmatched return transitions, and matching call-return transitions occurring along the path.

Lemma 17. *If $\langle q, \vec{n} \rangle\, c \in \text{pre}^*_{\mathcal{A}}(q_f)$ then $\langle q, \vec{n} \rangle \in \text{pre}^*_{\mathcal{A}_m}(q_f)$.*

Proof. We proceed by deriving a contradiction: if q_f is reachable from $\langle q, \vec{n} \rangle\, c$ in $\mathcal{A}$ while q_f is not reachable from $\langle q, \vec{n} \rangle$ in $\mathcal{A}_m$, the q_f must be reachable from $\langle q, \vec{n} \rangle\, c$ in $\mathcal{A}$ by an execution with an infinite depth of recursive calls.

Fix a minimal $j_0 \in \mathbb{N}$ such that for some c^0 we have $q^0\vec{n}^0c^0 \in \text{pre}^{j_0}_{\mathcal{A}}(q_f)$ while $q^0\vec{n}^0 \notin \text{pre}^{j_0}_{\mathcal{A}_m}(q_f)$. Furthermore, suppose $q^0\vec{n}^0c^0$ is a predecessor of $q\vec{n}c$; since j_0 is minimal $q\vec{n}$ is backward reachable in $\mathcal{A}_m$. On the one hand there can be no additive transition between $q^0\vec{n}^0c^0$ and $q\vec{n}c$ in $\mathcal{A}$; otherwise $\mathcal{A}_m$ would contain the same transition, and $q^0\vec{n}^0$ would be backward reachable in $\mathcal{A}_m$. On the other hand, if $\mathcal{A}$ made a return transition from $q^0\vec{n}^0c^0$ to $q\vec{n}c$ (where $\vec{n}^0 \leq \vec{n}$) then δ_{R}^m would have contained a transition from $q^0\vec{n}^0$ to $q\vec{n}$ (as one for every $\vec{n}^0 \leq \vec{n}$ is added). Thus the only remaining possibility is that $\mathcal{A}$ made a call transition from $q^0\vec{n}^0\, c^0$ to $q_1^0\mathbf{0}\, q^0\vec{n}^0\, c^0$.

Let $\langle q_2^0, \vec{n}_2^0 \rangle\, \langle q^0, \vec{n}^0 \rangle\, c^0 \rightarrow^{\text{rvas}} \langle q'^0, \vec{n}^0 \oplus \vec{n}_2^0 \rangle\, c^0$ be the matching return transition in $\mathcal{A}$; note that $\langle q'^0, \vec{n}^0 \oplus \vec{n}_2^0 \rangle$ must be reachable in $\mathcal{A}_m$ since j_0 is minimal, as well as $\langle q_1^0, \mathbf{0} \rangle$. However, since $\langle q^0, \vec{n}^0 \rangle$ is not backward reachable in $\mathcal{A}_m$, we know $\mathbf{0} \notin f_m(q_1^0, q_2^0, \vec{n}_2^0)$; otherwise δ_{R}^m would contain a transition from $\langle q^0, \vec{n}^0 \rangle$ to $\langle q'^0, \vec{n}^0 \oplus \vec{n}_2^0 \rangle$. Thus $\langle q_2^0, \vec{n}_2^0 \rangle$ is not reachable from $\langle q_1^0, \mathbf{0} \rangle$ in $\mathcal{A}_m$, despite the fact that $\langle q_2^0, \vec{n}_2^0 \rangle\, c^1$ is reachable from $\langle q_1^0, \mathbf{0} \rangle\, c^1$ in $\mathcal{A}$.

Applying the same arguments again, we fix a minimal $j_1 \in \mathbb{N}$ such that $\langle q^1, \vec{n}^1 \rangle\, c^1 \in \text{pre}^{j_1}_{\mathcal{A}}(q_2^0\vec{n}_2^0)$ while $\langle q^1, \vec{n}^1 \rangle \notin \text{pre}^{j_1}_{\mathcal{A}_m}(q_2^0\vec{n}_2^0)$. It is not hard to see that $\langle q^1, \vec{n}^1 \rangle$ must again be the predecessor to a call transition; i.e., $\langle q^1, \vec{n}^1 \rangle\, c^1 \rightarrow^{\text{rvas}} \langle q_1^1, \mathbf{0} \rangle\, \langle q^1, \vec{n}^1 \rangle\, c^1$. As before, let $\langle q_2^1, \vec{n}_2^1 \rangle\, \langle q^1, \vec{n}^1 \rangle\, c^1 \rightarrow^{\text{rvas}} \langle q'^1, \vec{n}^1 \oplus \vec{n}_2^1 \rangle\, c^1$ be the matching return transition in $\mathcal{A}$; note that $\langle q'^1, \vec{n}^1 \oplus \vec{n}_2^1 \rangle$ must be reachable in $\mathcal{A}_m$ since j_1 is minimal, as well as $\langle q_1^1, \mathbf{0} \rangle$. Again, since $\langle q^1, \vec{n}^1 \rangle$ is not backward reachable in $\mathcal{A}_m$, we know

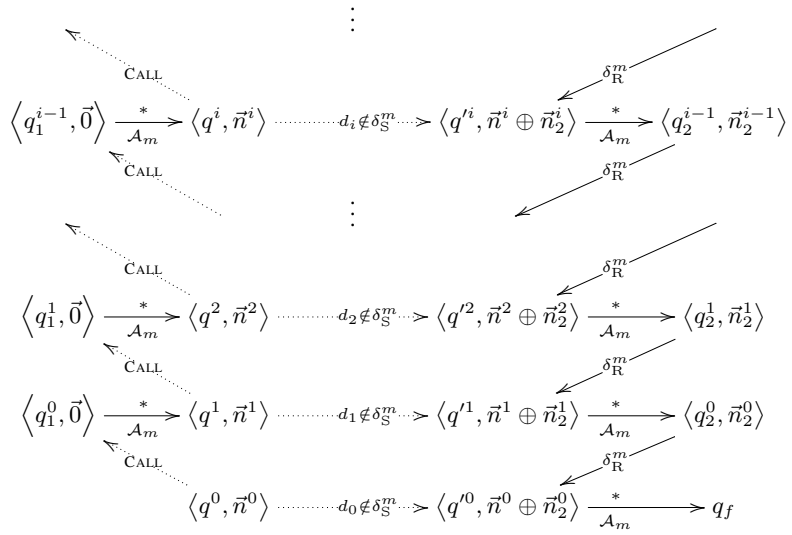

Figure 11. If q_f is reachable from $q^0\vec{n}^0$ in $\mathcal{A}$ but not in $\mathcal{A}_m$, then q_f can only be reachable from $q^0\vec{n}^0$ in $\mathcal{A}$ by an infinite execution; see the proof of Lemma 17.

$\mathbf{0} \notin f_m(q_1^1, q_2^1, \vec{n}_2^1)$; otherwise d_{R}^m would contain a transition from $\langle q^1, \vec{n}^1 \rangle$ to $\langle q'^1, \vec{n}^1 \oplus \vec{n}_2^1 \rangle$. Thus we know $\langle q_2^1, \vec{n}_2^1 \rangle$ is not reachable from $\langle q_1^1, \mathbf{0} \rangle$ in $\mathcal{A}_m$, despite the fact that $\langle q_2^1, \vec{n}_2^1 \rangle\, c^2$ is reachable from $\langle q_1^1, \mathbf{0} \rangle\, c^2$ in $\mathcal{A}$.

Note that in each step $i > 0$, $j_i < j_{i-1}$, since j_i is the length of a sub-execution of a j_{i-1}-length (sub) execution. Repeating the argument recursively ad infimum—i.e., insisting that $\mathcal{A}$ must have made an additional recursive call (see Figure 11)—we must conclude that $j_0 > j_1 > j_2 > \cdots$ is an infinite decreasing sequence of natural numbers, which is of course not possible. □ □

In particular, for any $\vec{n} \in \mathbb{N}^k$, we have that $\langle q_0, \mathbf{0} \rangle$ is backward reachable from $\langle q_f, \vec{n} \rangle$ in $\mathcal{A}_m$ if (and only if by Lemma 16) $\langle q_0, \mathbf{0} \rangle$ is backward reachable from $\langle q_f, \vec{n} \rangle$ in $\mathcal{A}$. □ □