# GENERAL DECIDABILITY RESULTS FOR ASYNCHRONOUS SHARED-MEMORY PROGRAMS: HIGHER-ORDER AND BEYOND 

RUPAK MAJUMDAR, RAMANATHAN S. THINNIYAM, AND GEORG ZETZSCHE

Max Planck Institute for Software Systems (MPI-SWS), Paul-Ehrlich-Straße, Building G26, Kaiserslautern, Germany 67663.<br>e-mail address: \{rupak,thinniyam,georg\}@mpi-sws.org


#### Abstract

The model of asynchronous programming arises in many contexts, from lowlevel systems software to high-level web programming. We take a language-theoretic perspective and show general decidability and undecidability results for asynchronous programs that capture all known results as well as show decidability of new and important classes. As a main consequence, we show decidability of safety, termination and boundedness verification for higher-order asynchronous programs-such as OCaml programs using Lwtand undecidability of liveness verification already for order-2 asynchronous programs. We show that under mild assumptions, surprisingly, safety and termination verification of asynchronous programs with handlers from a language class are decidable iff emptiness is decidable for the underlying language class. Moreover, we show that configuration reachability and liveness (fair termination) verification are equivalent, and decidability of these problems implies decidability of the well-known "equal-letters" problem on languages. Our results close the decidability frontier for asynchronous programs.


## 1. Introduction

Asynchronous programming is a common way to manage concurrent requests in a system. In this style of programming, rather than waiting for a time-consuming operation to complete, the programmer can make asynchronous procedure calls which are stored in a task buffer pending later execution. Each asynchronous procedure, or handler, is a sequential program. When run, it can change the global shared state of the program, make internal synchronous procedure calls, and post further instances of handlers to the task buffer. A scheduler repeatedly and non-deterministically picks pending handler instances from the task buffer and executes their code atomically to completion. Asynchronous programs appear in many domains, such as operating system kernel code, web programming, or user applications on mobile platforms. This style of programming is supported natively or through libraries for most programming environments. The interleaving of different handlers hides latencies

[^0][^1]of long-running operations: the program can process a different handler while waiting for an external operation to finish. However, asynchronous scheduling of tasks introduces non-determinism in the system, making it difficult to reason about correctness.

An asynchronous program is finite-data if all program variables range over finite domains. Finite-data programs are still infinite state transition systems: the task buffer can contain an unbounded number of pending instances and the sequential machine implementing an individual handler can have unboundedly large state (e.g., if the handler is given as a recursive program, the stack can grow unboundedly). Nevertheless, verification problems for finite-data programs have been shown to be decidable for several kinds of handlers [SV06, JM07, CV07, GM12]. Several algorithmic approaches have been studied, which tailor to (i) the kinds of permitted handler programs and (ii) the properties that are checked.

State of the art. We briefly survey the existing approaches and what is known about the decidability frontier. The Parikh approach applies to (first-order) recursive handler programs. Here, the decision problems for asynchronous programs are reduced to decision problems over Petri nets [GM12]. The key insight is that since handlers are executed atomically, the order in which a handler posts tasks to the buffer is irrelevant. Therefore, instead of considering the sequential order of posted tasks along an execution, one can equivalently consider its Parikh image. Thus, when handlers are given as pushdown systems, the behaviors of an asynchronous program can be represented by a (polynomial sized) Petri net. Using the Parikh approach, safety (formulated as reachability of a global state), termination (whether all executions terminate), and boundedness (whether there is an a priori upper bound on the task buffer) are all decidable for asynchronous programs with recursive handlers, by reduction to corresponding problems on Petri nets [SV06, GM12]. Configuration reachability (reachability of a specific global state and task buffer configuration), fair termination (termination under a fair scheduler), and fair non-starvation (every pending handler instance is eventually executed) are also decidable, by separate ad hoc reductions to Petri net reachability [GM12]. A "reverse reduction" shows that Petri nets can be simulated by polynomial-sized asynchronous programs (already with finite-data handlers).

In the downclosure approach, one replaces each handler with a finite-data program that is equivalent up to "losing" handlers in the task buffer. Of course, this requires that one can compute equivalent finite-data programs for a given class of handler programs. This has been applied to checking safety for recursive handler programs [ABQ09]. Finally, a bespoke rank-based approach has been applied to checking safety when handlers can perform restricted higher-order recursion [CV07].

Contribution. Instead of studying individual kinds of handler programs, we consider asynchronous programs in a general language-theoretic framework. The class of handler programs is given as a language class $\mathcal{C}$ : An asynchronous program over a language class $\mathcal{C}$ is one where each handler defines a language from $\mathcal{C}$ over the alphabet of handler names, as well as a transformer over the global state. This view leads to general results: we can obtain simple characterizations of which classes of handler programs permit decidability. For example, we do not need the technical assumptions of computability of equivalent finite-data programs from the Parikh and the downclosure approach.

Our first result shows that, under a mild language-theoretic assumption, safety and termination are decidable if and only if the underlying language class $\mathcal{C}$ has a decidable
emptiness problem. ${ }^{1}$ Similarly, we show that boundedness is decidable iff finiteness is decidable for the language class $\mathcal{C}$. These results are the best possible: decidability of emptiness (resp., finiteness) is a requirement for safety and termination verification already for verifying the safety or termination (resp., boundedness) of one sequential handler call. As corollaries, we get new decidability results for all these problems for asynchronous programs over higher-order recursion schemes, which form the language-theoretic basis for programming in higher-order functional languages such as OCaml [Kob09, Ong15], as well as other language classes (lossy channel languages, Petri net languages, etc.).

Second, we show that configuration reachability, fair termination, and fair starvation are mutually reducible; thus, decidability of any one of them implies decidability of all of them. We also show decidability of these problems implies the decidability of a well-known combinatorial problem on languages: given a language over the alphabet $\{\mathrm{a}, \mathrm{b}\}$, decide if it contains a word with an equal number of as and bs. Viewed contrapositively, we conclude that all these decision problems are undecidable already for asynchronous programs over order-2 pushdown languages, since the equal-letters problem is undecidable for this class.

Together, our results "close" the decidability frontier for asynchronous programs, by demonstrating reducibilities between decision problems heretofore studied separately and connecting decision problems on asynchronous programs with decision problems on the underlying language classes of their handlers.

While our algorithms do not assume that downclosures are effectively computable, we use downclosures to prove their correctness. We show that the safety, termination, and boundedness problems are invariant under taking downclosures of runs; this corresponds to taking downclosures of the languages of handlers.

The observation that safety, termination, and boundedness depend only on the downclosure suggests a possible route to implementation. If there is an effective procedure to compute the downclosure for class $\mathcal{C}$, then a direct verification algorithm would replace all handlers by their (regular) downclosures, and invoke existing decision procedures for this class. Thus, we get a direct algorithm based on downclosure constructions for higher-order recursion schemes, using the string of celebrated recent results on effectively computing the downclosure of word schemes [Zet15, HKO16, CPSW16].

We find our general decidability result for asynchronous programs to be surprising. Already for regular languages, the complexity of safety verification jumps from NL (NFA emptiness) to EXPSPACE (Petri net coverability): asynchronous programs are far more expressive than individual handler languages. It is therefore unexpected that safety and termination remain decidable whenever they are decidable for individual handler languages.

## 2. Preliminaries

Basic Definitions. We assume familiarity with basic definitions of automata theory (see, e.g., [HMU07, Sip97]). If $w$ is a word over $\Sigma$ and $\Gamma \subseteq \Sigma$ is a subalphabet, then the projection of $w$ onto $\Gamma$, denoted $\operatorname{Proj}_{\Gamma}(w)$, is the word obtained by erasing from $w$ each symbol that does not belong to $\Gamma$. For a language $L \subseteq \Sigma^{*}$, define $\operatorname{Proj}_{\Gamma}(L)=\left\{\operatorname{Proj}_{\Gamma}(w) \mid w \in L\right\}$. The subword order $\sqsubseteq$ on $\Sigma^{*}$ is defined as $w \sqsubseteq w^{\prime}$ for $w, w^{\prime} \in \Sigma^{*}$ if $w$ can be obtained

[^2]from $w^{\prime}$ by deleting some letters from $w^{\prime}$. For example, $a b b a \sqsubseteq b a b a b a$ but $a b b a ~ \sharp b a a b a$. The downclosure $\downarrow w$ with respect to the subword order, of a word $w \in \Sigma^{*}$, is defined as $\downarrow w:=\left\{w^{\prime} \in \Sigma^{*} \mid w^{\prime} \sqsubseteq w\right\}$. The downclosure $\downarrow L$ of a language $L \subseteq \Sigma^{*}$ is given by $\downarrow L:=\left\{w^{\prime} \in \Sigma^{*} \mid \exists w \in L: w^{\prime} \sqsubseteq w\right\}$. Recall that the downclosure $\downarrow L$ of any language $L$ is a regular language [Hai69].

A multiset $\mathbf{m}: \Sigma \rightarrow \mathbb{N}$ over $\Sigma$ maps each symbol of $\Sigma$ to a natural number. Let $\mathbb{M}[\Sigma]$ be the set of all multisets over $\Sigma$. We treat sets as a special case of multisets where each element is mapped onto 0 or 1 . For example, we write $\mathbf{m}=\llbracket a, a, c \rrbracket$ for the multiset $\mathbf{m} \in \mathbb{M}[\{\mathrm{a}, \mathrm{b}, \mathrm{c}, \mathrm{d}\}]$ with $\mathbf{m}(\mathrm{a})=2, \mathbf{m}(\mathrm{~b})=\mathbf{m}(\mathrm{d})=0$, and $\mathbf{m}(\mathrm{c})=1$. We also use the notation $|\mathbf{m}|=\sum_{\sigma \in \Sigma} \mathbf{m}(\sigma)$. The Parikh image $\operatorname{Parikh}(w) \in \mathbb{M}[\Sigma]$ of a word $w \in \Sigma^{*}$ is the multiset such that $\operatorname{Parikh}(w)(\mathrm{a})$ is the number of times a occurs in $w$.

Given two multisets $\mathbf{m}, \mathbf{m}^{\prime} \in \mathbb{M}[\Sigma]$ we define the multiset $\mathbf{m} \oplus \mathbf{m}^{\prime} \in \mathbb{M}[\Sigma]$ for which, for all $\mathrm{a} \in \Sigma$, we have $\left(\mathbf{m} \oplus \mathbf{m}^{\prime}\right)(\mathrm{a})=\mathbf{m}(\mathrm{a})+\mathbf{m}^{\prime}(\mathrm{a})$. We also define the natural order $\preceq$ on $\mathbb{M}[\Sigma]$ as follows: $\mathbf{m} \preceq \mathbf{m}^{\prime}$ iff there exists $\mathbf{m}^{\Delta} \in \mathbb{M}[\Sigma]$ such that $\mathbf{m} \oplus \mathbf{m}^{\Delta}=\mathbf{m}^{\prime}$. We also define $\mathbf{m}^{\prime} \ominus \mathbf{m}$ for $\mathbf{m} \preceq \mathbf{m}^{\prime}$ analogously: for all $\mathrm{a} \in \Sigma$, we have $\left(\mathbf{m} \ominus \mathbf{m}^{\prime}\right)(\mathrm{a})=\mathbf{m}(\mathrm{a})-\mathbf{m}^{\prime}(\mathrm{a})$. For $\Gamma \subseteq \Sigma$ we regard $\mathbf{m} \in \mathbb{M}[\Gamma]$ as a multiset of $\mathbb{M}[\Sigma]$ where undefined values are sent to 0 .

Language Classes and Full Trios. A language class is a collection of languages, together with some finite representation. Examples are the regular (e.g. represented by finite automata) or the context-free languages (e.g. represented by pushdown automata (PDA)). A relatively weak and reasonable assumption on a language class is that it is a full trio, that is, it is closed under each of the following operations: taking intersection with a regular language, taking homomorphic images, and taking inverse homomorphic images. Equivalently, a language class is a full trio iff it is closed under rational transductions [Ber79] as we explain below.

We assume that all full trios $\mathcal{C}$ considered in this paper are effective: Given a language $L$ from $\mathcal{C}$, a regular language $R$, and a homomorphism $h$, we can compute a representation of the languages $L \cap R, h(L)$, and $h^{-1}(L)$ in $\mathcal{C}$.

Many classes of languages studied in formal language theory form effective full trios. Examples include the regular and the context-free languages [HMU07], the indexed languages [Aho68, DG86], the languages of higher-order pushdown automata [Mas74], higherorder recursion schemes (HORS) [Dam82, HMOS08], Petri nets (both coverability and reachability languages) [Gre78, Jan79], and lossy channel systems (see Section 4.1). (While HORS are usually viewed as representing a tree or collection of trees, one can also view them as representing a word language, as we explain in Section 5.)

Informally, a language class defined by non-deterministic devices with a finite-state control that allows $\varepsilon$-transitions and imposes no restriction between input letter and performed configuration changes (such as non-deterministic pushdown automata) is always a full trio: The three operations above can be realized by simple modifications of the finite-state control. The deterministic context-free languages are a class that is not a full trio.

An asynchronous transducer $\mathcal{T}$ is a tuple $\mathcal{T}=\left(Q, \Gamma, \Sigma, E, q_{0}, F\right)$ with a finite set of states $Q$, finite output alphabet $\Gamma$, finite input alphabet $\Sigma$, a finite set of edges $E \subseteq Q \times \Gamma^{*} \times \Sigma^{*} \times Q$, initial state $q_{0} \in Q$ and set of final states $F \subseteq Q$. We write $p \xrightarrow{v \mid u} q$ if $(p, v, u, q) \in E$ and the machine reads $u$ in state $p$, outputs $v$ and moves to state $q$. We also write $p \xrightarrow{w \mid w^{\prime}} q$ if there are states $q_{0}, q_{1}, \ldots, q_{n}$ and words $u_{1}, u_{2}, \ldots, u_{n}, v_{1}, v_{2}, \ldots, v_{n}$ such that $p=q_{0}, q=q_{n}$, $w^{\prime}=u_{1} u_{2} \cdots u_{n}, w=v_{1} v_{2} \cdots v_{n}$ and $q_{i-1} \xrightarrow{v_{i} \mid u_{i}} q_{i}$ for all $0 \leq i \leq n$.

The transduction $T \subseteq \Gamma^{*} \times \Sigma^{*}$ generated by the transducer $\mathcal{T}$ is the set of tuples $(v, u) \in \Gamma^{*} \times \Sigma^{*}$ such that $q_{0} \xrightarrow{v \mid u} q_{f}$ for some $q_{f} \in F$. Given a language $L \subseteq \Sigma^{*}$, we define $T L:=\left\{v \in \Gamma^{*} \mid \exists u \in L:(v, u) \in T\right\}$. A transduction $T \subseteq \Gamma^{*} \times \Sigma^{*}$ is rational if it is generated by some asynchronous transducer.

A language class $\mathcal{C}$ is a full trio if and only if it is closed under each of the following operations:

- intersection with a regular language,
- taking homomorphic images, and
- taking inverse homomorphic images.

It is well known that these two concepts are equivalent:
Theorem 2.1 (Berstel [Ber79]). A language class is closed under rational transductions if and only if it is a full trio.

Asynchronous Programs: A Language-Theoretic View. We use a language-theoretic model for asynchronous shared-memory programs.
Definition 2.2. Let $\mathcal{C}$ be an (effective) full trio. An asynchronous program (AP) over $\mathcal{C}$ is a tuple $\mathfrak{P}=\left(D, \Sigma,\left(L_{c}\right)_{c \in \mathfrak{C}}, d_{0}, \mathbf{m}_{0}\right)$, where $D$ is a finite set of global states, $\Sigma$ is an alphabet of handler names, $\left(L_{c}\right)_{c \in \mathfrak{C}}$ is a family of languages over $\Sigma$ from $\mathcal{C}$, one for each $c \in \mathfrak{C}$ where $\mathfrak{C}=D \times \Sigma \times D$ is the set of contexts, $d_{0} \in D$ is the initial state, and $\mathbf{m}_{0} \in \mathbb{M}[\Sigma]$ is a multiset of initial pending handler instances.

A configuration $(d, \mathbf{m}) \in D \times \mathbb{M}[\Sigma]$ of $\mathfrak{P}$ consists of a global state $d$ and a multiset $\mathbf{m}$ of pending handler instances. For a configuration $c$, we write $c . d$ and $c . \mathbf{m}$ for the global state and the multiset in the configuration respectively. The initial configuration $c_{0}$ of $\mathfrak{P}$ is given by $c_{0} \cdot d=d_{0}$ and $c_{0} \cdot \mathbf{m}=\mathbf{m}_{0}$. The semantics of $\mathfrak{P}$ is a labeled transition system over the set of configurations, with the transition relation $\xrightarrow{\boldsymbol{\sigma}} \subseteq(D \times \mathbb{M}[\Sigma]) \times(D \times \mathbb{M}[\Sigma])$ given by

$$
(d, \mathbf{m} \oplus \llbracket \sigma \rrbracket) \xrightarrow{\sigma}\left(d^{\prime}, \mathbf{m} \oplus \mathbf{m}^{\prime}\right) \quad \text { iff } \quad \exists w \in L_{d \sigma d^{\prime}}: \operatorname{Parikh}(w)=\mathbf{m}^{\prime}
$$

Moreover, we write $(d, \mathbf{m}) \rightarrow\left(d^{\prime}, \mathbf{m}^{\prime}\right)$ if there exists a $\sigma \in \Sigma$ with $(d, \mathbf{m}) \xrightarrow{\sigma}\left(d^{\prime}, \mathbf{m}^{\prime}\right)$. We use $\rightarrow^{*}$ to denote the reflexive transitive closure of the relation $\rightarrow$. A configuration $c$ is said to be reachable in $\mathfrak{P}$ if $\left(d_{0}, \mathbf{m}_{0}\right) \rightarrow^{*} c$.

Intuitively, the set $\Sigma$ of handler names specifies a finite set of procedures that can be invoked asynchronously. The shared state takes values in $D$. When a handler is called asynchronously, it gets added to a bag of pending handler calls (the multiset $\mathbf{m}$ in a configuration). The language $L_{d \sigma d^{\prime}}$ captures the effect of executing an instance of $\sigma$ starting from the global state $d$, such that on termination, the global state is $d^{\prime}$. Each word $w \in L_{d \sigma d^{\prime}}$ captures a possible sequence of handlers posted during the execution.

Suppose the current configuration is $(d, \mathbf{m})$. A non-deterministic scheduler picks one of the outstanding handlers $\sigma \in \mathbf{m}$ and executes it. Executing $\sigma$ corresponds to picking one of the languages $L_{d \sigma d^{\prime}}$ and some word $w \in L_{d \sigma d^{\prime}}$. Upon execution of $\sigma$, the new configuration has global state $d^{\prime}$ and the new bag of pending calls is obtained by taking $\mathbf{m}$, removing an instance of $\sigma$ from it, and adding the Parikh image of $w$ to it. This reflects the current set of pending handler calls - the old ones (minus an instance of $\sigma$ ) together with the new ones added by executing $\sigma$. Note that a handler is executed atomically; thus, we atomically update the global state and the effect of executing the handler.

```
global var turn = ref 0 and x = ref 0;
let rec s1 () = if * then begin post a; s1(); post b end
let rec s2 () = if * then begin post a; s2(); post b end else post b
let a () = if !turn = 0 then begin turn := 1; x := !x + 1 end else post a
let b () = if !turn = 1 then begin turn := 0; x := !x - 1 end else post b
let s3 () = post s3; post s3
global var t = ref 0;
let c () = if !t = 0 then t := 1 else post c
let d () = if !t = 1 then t := 2 else post d
let f () = if !t = 2 then t := 0 else post f
let cc x = post c; x
let dd x = post d; x
let ff x = post f; x
let id x = x
let h g y = cc (g (dd y))
let rec produce g x = if * then produce (h g) (ff x) else g x
let s4 () = produce id ()
```

Figure 1: Examples of asynchronous programs

Let us see some examples of asynchronous programs. It is convenient to present these examples in a programming language syntax, and to allow each handler to have internal actions that perform local tests and updates to the global state. As we describe informally below, and formally in Appendix A, when $\mathcal{C}$ is a full trio, internal actions can be "compiled away" by taking an intersection with a regular language of internal actions and projecting the internal actions away. Thus, we use our simpler model throughout.

Examples. Figure 1 shows some simple examples of asynchronous programs in an OCamllike syntax. Consider first the asynchronous program in lines $1-5$. The alphabet of handlers is $s 1$, $s 2$, a, and b. The global states correspond to possible valuations to the global variables turn and x ; assuming turn is a Boolean and x takes values in $\mathbb{N}$, we have that $D=$ $\{0,1\} \times\{0,1, \omega\}$, where $\omega$ abstracts all values other than $\{0,1\}$. Since s1 and s2 do not touch any variables, for $d, d^{\prime} \in D$, we have $L_{d, \mathbf{s} 1, d}=\left\{\mathrm{a}^{n} \mathrm{~b}^{n} \mid n \geq 0\right\}, L_{d, \mathbf{s} 2, d}=\left\{\mathrm{a}^{n} \mathrm{~b}^{n+1} \mid n \geq 0\right\}$, and $L_{d, \mathbf{s} 1, d^{\prime}}=L_{d, \mathbf{s} 2, d^{\prime}}=\emptyset$ if $d^{\prime} \neq d$.

For the languages corresponding to a and b , we use syntactic sugar in the form of internal actions; these are local tests and updates to the global state. In our example, we have $L_{(0,0), \mathrm{a},(1,1)}=\{\varepsilon\}, L_{(1, x), \mathrm{a},(1, x)}=\{\mathrm{a}\}$ for all values of $x$, and similarly for b . The meaning is that, starting from a global state $(0,0)$, executing the handler will lead to the global state $(1,1)$ and no handlers will be posted, whereas starting from a global state in which turn is 1 , executing the handler will keep the global state unchanged but post an instance of a. Note that all the languages are context-free.

Consider an execution of the program from the initial configuration $((0,0), \llbracket s 1 \rrbracket)$. The execution of s1 puts $n$ as and $n$ bs into the bag, for some $n \geq 0$. The global variable turn is used to ensure that the handlers a and b alternately update x . When turn is 0 , the handler for a increments x and sets turn to 1 , otherwise it re-posts itself for a future execution. Likewise, when turn is 1 , the handler for b decrements x and sets turn back to 0 , otherwise
it re-posts itself for a future execution. As a result, the variable x never grows beyond 1. Thus, the program satisfies the safety property that no execution sets x to $\omega$.

It is possible that the execution goes on forever: for example, if s1 posts an a and a b , and thereafter only b is chosen by the scheduler. This is not an "interesting" infinite execution as it is not fair to the pending a. In the case of a fair scheduler, which eventually always picks an instance of every pending task, the program terminates: eventually all the as and bs are consumed when they are scheduled in alternation. However, if instead we started with 【s2】, the program will not terminate even under a fair scheduler: the last remaining b will not be paired and will keep executing and re-posting itself forever.

Now consider the execution of s3. It has an infinite fair run, where the scheduler picks an instance of s3 at each step. However, the number of pending instances grows without bound. We shall study the boundedness problem, which checks if the bag can become unbounded along some run. We also study a stronger notion of fair termination, called fair non-starvation, which asks that every instance of a posted handler is executed under any fair scheduler. The execution of s3 is indeed fair, but there can be a specific instance of s3 that is never picked: we say s3 can starve an instance.

The program in lines $9-20$ is higher-order (produce and h take functions as arguments). The language of $s 4$ is the set $\left\{\mathrm{c}^{n} \mathrm{~d}^{n} \mathrm{f}^{n} \mid n \geq 0\right\}$, that is, it posts an equal number of cs , ds, and fs. It is an indexed language; we shall see (Section 5) how this and other higher-order programs can be represented using higher-order recursion schemes (HORS). Note that the OCaml types of produce : $(0 \rightarrow 0) \rightarrow 0 \rightarrow 0$ and $\mathrm{h}:(\mathrm{o} \rightarrow \mathrm{o}) \rightarrow 0 \rightarrow 0$ are higher-order.

The program is similar to the first: the handlers $c, d$, and $f$ execute in "round robin" fashion using the global state $t$ to find their turns. Again, we use internal actions to update the global state for readability. We ask the same decision questions as before: does the program ever reach a specific global state and does the program have an infinite (fair) run? We shall see later that safety and termination questions remain decidable, whereas fair termination does not.

## 3. Decision Problems on Asynchronous Programs

We now describe decision problems on runs of asynchronous programs.
Runs, preruns, and downclosures. A prerun of an $\mathrm{AP} \mathfrak{P}=\left(D, \Sigma,\left(L_{c}\right)_{c \in \mathfrak{C}}, d_{0}, \mathbf{m}_{0}\right)$ is a finite or infinite sequence $\rho=\left(e_{0}, \mathbf{n}_{0}\right), \sigma_{1},\left(e_{1}, \mathbf{n}_{1}\right), \sigma_{2}, \ldots$ of alternating elements of tuples $\left(e_{i}, \mathbf{n}_{i}\right) \in D \times \mathbb{M}[\Sigma]$ and symbols $\sigma_{i} \in \Sigma$. The set of preruns of $\mathfrak{P}$ will be denoted Preruns $(\mathfrak{P})$. Note that if two asynchronous programs $\mathfrak{P}$ and $\mathfrak{P}^{\prime}$ have the same $D$ and $\Sigma$, then $\operatorname{Preruns}(\mathfrak{P})=\operatorname{Preruns}\left(\mathfrak{P}^{\prime}\right)$. The length, denoted $|\rho|$, of a finite prerun $\rho$ is the number of configurations in $\rho$. The $i^{\text {th }}$ configuration of a prerun $\rho$ will be denoted $\rho(i)$.

We define an order $\unlhd$ on preruns as follows: For preruns $\rho=\left(e_{0}, \mathbf{n}_{0}\right), \sigma_{1},\left(e_{1}, \mathbf{n}_{1}\right), \sigma_{2}, \ldots$ and $\rho^{\prime}=\left(e_{0}^{\prime}, \mathbf{n}_{0}^{\prime}\right), \sigma_{1}^{\prime},\left(e_{1}^{\prime}, \mathbf{n}_{1}^{\prime}\right), \sigma_{2}^{\prime}, \ldots$, we define $\rho \unlhd \rho^{\prime}$ if $|\rho|=\left|\rho^{\prime}\right|$ and $e_{i}=e_{i}^{\prime}$, $\sigma_{i}=\sigma_{i}^{\prime}$ and $\mathbf{n}_{i} \preceq \mathbf{n}_{i}^{\prime}$ for each $i \geq 0$. The downclosure $\downarrow R$ of a set $R$ of preruns of $\mathfrak{P}$ is defined as $\downarrow R=\left\{\rho \in \operatorname{Preruns}(\mathfrak{P}) \mid \exists \rho^{\prime} \in R . \rho \unlhd \rho^{\prime}\right\}$.

A run of an $\mathrm{AP} \mathfrak{P}=\left(D, \Sigma,\left(L_{c}\right)_{c \in \mathfrak{C}}, d_{0}, \mathbf{m}_{0}\right)$ is a prerun $\rho=\left(d_{0}, \mathbf{m}_{0}\right), \sigma_{1},\left(d_{1}, \mathbf{m}_{1}\right), \sigma_{2}, \ldots$ starting with the initial configuration $\left(d_{0}, \mathbf{m}_{0}\right)$, where for each $i \geq 0$, we have $\left(d_{i}, \mathbf{m}_{i}\right) \xrightarrow{\sigma_{i+1}}$ $\left(d_{i+1}, \mathbf{m}_{i+1}\right)$. The set of runs of $\mathfrak{P}$ is denoted $\operatorname{Runs}(\mathfrak{P})$. Finally, $\downarrow \operatorname{Runs}(\mathfrak{P})$ is the downward closure of Runs $(\mathfrak{P})$ with respect to $\unlhd$.

An infinite run $c_{0} \xrightarrow{\sigma_{0}} c_{1} \xrightarrow{\sigma_{1}} \ldots$ is called fair if for all $i \geq 0$, if $\sigma \in c_{i}$. $\mathbf{m}$ then there is some $j \geq i$ such that $c_{j} \xrightarrow{\sigma} c_{j+1}$. That is, whenever an instance of a handler $\sigma$ is posted, some instance of $\sigma$ is executed later. Fairness does not guarantee that every specific instance of a handler is executed eventually. We say that an infinite fair run starves a handler $\sigma$ if there exists an index $J \geq 0$ such that for each $j \geq J$, we have (i) $c_{j} . \mathrm{m}(\sigma) \geq 1$ and (ii) whenever $c_{j} \xrightarrow{\sigma} c_{j+1}$, we have $c_{j} . \mathbf{m}(\sigma) \geq 2$. In this case, even if the run is fair, a specific instance of $\sigma$ may never be executed.

Now we give the definitions of the various decision problems.
Definition 3.1 (Properties of finite runs). The Safety (Global state reachability) problem asks, given an asynchronous program $\mathfrak{P}$ and a global state $d_{f} \in D$, is there a reachable configuration $c$ such that $c . d=d_{f}$ ? If so, $d_{f}$ is said to be reachable (in $\mathfrak{P}$ ) and unreachable otherwise. The Boundedness (of the task buffer) problem asks, given an asynchronous program $\mathfrak{P}$, is there an $N \in \mathbb{N}$ such that for every reachable configuration $c$, we have $|c . \mathbf{m}| \leq N$ ? If so, the asynchronous program $\mathfrak{P}$ is bounded; otherwise it is unbounded. The Configuration reachability problem asks, given an asynchronous program $\mathfrak{P}$ and a configuration $c$, is $c$ reachable?
Definition 3.2 (Properties of infinite runs). All the following problems take as input an asynchronous program $\mathfrak{P}$. The Termination problem asks if all runs of $\mathfrak{P}$ are finite. The Fair Non-termination problem asks if $\mathfrak{P}$ has some fair infinite run. The Fair Starvation problem asks if $\mathfrak{P}$ has some fair run that starves some handler.

Our main result in this section shows that many properties of an asynchronous program $\mathfrak{P}$ only depend on the downclosure $\downarrow \operatorname{Runs}(\mathfrak{P})$ of the set $\operatorname{Runs}(\mathfrak{P})$ of runs of the program $\mathfrak{P}$. The proof is by induction on the length of runs. Please see Appendix B for details. For any AP $\mathfrak{P}=\left(D, \Sigma,\left(L_{c}\right)_{c \in \mathfrak{C}}, d_{0}, \mathbf{m}_{0}\right)$, we define the AP $\downarrow \mathfrak{P}=\left(D, \Sigma,\left(\downarrow L_{c}\right)_{c \in \mathfrak{C}}, d_{0}, \mathbf{m}_{0}\right)$, where $\downarrow L_{c}$ is the downclosure of the language $L_{c}$ under the subword order.
Proposition 3.3. Let $\mathfrak{P}=\left(D, \Sigma,\left(L_{c}\right)_{c \in \mathfrak{C}}, d_{0}, \mathbf{m}_{0}\right)$ be an asynchronous program. Then $\downarrow \operatorname{Runs}(\downarrow \mathfrak{P})=\downarrow \operatorname{Runs}(\mathfrak{P})$. In particular, the following holds. (1) For every $d \in D, \mathfrak{P}$ can reach $d$ if and only if $\downarrow \mathfrak{P}$ can reach $d$. (2) $\mathfrak{P}$ is terminating if and only if $\downarrow \mathfrak{P}$ is terminating. (3) $\mathfrak{P}$ is bounded if and only if $\downarrow \mathfrak{P}$ is bounded.

Intuitively, safety, termination, and boundedness is preserved when the multiset of pending handler instances is "lossy": posted handlers can get lost. This corresponds to these handlers never being added to the task buffer. However, if a run demonstrates reachability of a global state, or non-termination, or unboundedness, in the lossy version, it corresponds also to a run in the original problem (and conversely).

In contrast, simple examples show that configuration reachability, fair termination, and fair non-starvation properties are not preserved under downclosures.

## 4. General Decidability Results

In this section, we characterize those full trios $\mathcal{C}$ for which particular problems for asynchronous programs over $\mathcal{C}$ are decidable. Our decision procedures will use the following theorem, summarizing the results from [GM12], as a subprocedure.
Theorem 4.1 [GM12]. Safety, boundedness, configuration reachability, termination, fair non-termination, and fair non-starvation are decidable for asynchronous programs over regular languages.

```
Algorithm 1: Checking Safety
    Input: Asynchronous program \(\mathfrak{P}=\left(D, \Sigma,\left(L_{c}\right)_{c \in \mathfrak{C}}, d_{0}, \mathbf{m}_{0}\right)\) over \(\mathcal{C}\), state \(d \in D\)
    run concurrently
        begin /* find a safe overapproximation */
                foreach tuple \(\left(A_{c}\right)_{c \in \mathfrak{C}}\) of regular languages \(A_{c} \subseteq \Sigma^{*}\) do
                    if \(L_{c} \cap\left(\Sigma^{*} \backslash A_{c}\right)=\emptyset\) for each \(c \in \mathfrak{C}\) then
                            if \(\mathfrak{P}^{\prime}=\left(D, \Sigma,\left(A_{c}\right)_{c \in \mathfrak{C}}, d_{0}, \mathbf{m}_{0}\right)\) does not reach \(d\) then
                                return \(d\) is not reachable.
        begin /* find a run reaching \(d\) */
            foreach prerun \(\rho\) of \(\mathfrak{P}\) do
            if \(\rho\) is a run that reaches \(d\) then
                \(\llcorner\) return \(d\) reachable.
```

4.1. Safety and termination. Our first main result concerns the problems of safety and termination.

Theorem 4.2. Let $\mathcal{C}$ be a full trio. The following are equivalent:
(i) Safety is decidable for asynchronous programs over $\mathcal{C}$.
(ii) Termination is decidable for asynchronous programs over $\mathcal{C}$.
(iii) Emptiness is decidable for $\mathcal{C}$.

Proof. We begin with "(i) $\Rightarrow$ (iii)". Let $K \subseteq \Sigma^{*}$ be given. We construct

$$
\mathfrak{P}=\left(D, \Sigma,\left(L_{c}\right)_{c \in \mathfrak{C}}, d_{0}, \mathbf{m}_{0}\right)
$$

such that $\mathbf{m}_{0}=\llbracket \sigma \rrbracket, D=\left\{d_{0}, d_{1}\right\}, L_{d_{0}, \sigma, d_{1}}=K$ and $L_{c}=\emptyset$ for $c \neq\left(d_{0}, \sigma, d_{1}\right)$. We see that $\mathfrak{P}$ can reach $d_{1}$ iff $K$ is non-empty. Next we show "(ii) $\Rightarrow$ (iii)". Consider the alphabet $\Gamma=(\Sigma \cup\{\varepsilon\}) \times\{0,1\}$ and the homomorphisms $g: \Gamma^{*} \rightarrow \Sigma^{*}$ and $h: \Gamma^{*} \rightarrow\{\sigma\}^{*}$, where for $x \in \Sigma \cup\{\varepsilon\}$, we have $g((x, i))=x$ for $i \in\{0,1\}, h((x, 1))=\sigma$, and $h((x, 0))=\varepsilon$. If $R \subseteq \Gamma^{*}$ is the regular set of words in which exactly one position belongs to the subalphabet $(\Sigma \cup\{\varepsilon\}) \times\{1\}$, then the language $K^{\prime}:=h\left(g^{-1}(K) \cap R\right)$ belongs to $\mathcal{C}$. Note that $K^{\prime}$ is $\emptyset$ or $\{\sigma\}$, depending on whether $K$ is empty or not. We construct $\mathfrak{P}=\left(D,\{\sigma\},\left(L_{c}\right)_{c \in \mathfrak{C}}, d_{0}, \mathbf{m}_{0}\right)$ with $D=\left\{d_{0}\right\}, \mathbf{m}_{0}=\llbracket \sigma \rrbracket, L_{d_{0}, \sigma, d_{0}}=K^{\prime}$ and all languages $L_{c}=\emptyset$ for $c \neq\left(d_{0}, \sigma, d_{0}\right)$. Then $\mathfrak{P}$ is terminating iff $K$ is empty.

To prove "(iii) $\Rightarrow$ (i)", we design an algorithm deciding safety assuming decidability of emptiness. Given asynchronous program $\mathfrak{P}$ and state $d$ as input, the algorithm consists of two semi-decision procedures: one which searches for a run of $\mathfrak{P}$ reaching the state $d$, and the second which enumerates regular overapproximations $\mathfrak{P}^{\prime}$ of $\mathfrak{P}$ and checks the safety of $\mathfrak{P}^{\prime}$ using Theorem 4.1. Each $\mathfrak{P}^{\prime}$ consists of a regular language $A_{c}$ overapproximating $L_{c}$ for each context $c$ of $\mathfrak{P}$. We use decidability of emptiness to check that $L_{c} \cap\left(\Sigma^{*} \backslash A_{c}\right)=\emptyset$ to ensure that $\mathfrak{P}^{\prime}$ is indeed an overapproximation.

Algorithm 1 clearly gives a correct answer if it terminates. Hence, we only have to argue that it always does terminate. Of course, if $d$ is reachable, the first semi-decision procedure will terminate. In the other case, termination is due to the regularity of downclosures: if $d$ is not reachable in $\mathfrak{P}$, then Proposition 3.3 tells us that $\downarrow \mathfrak{P}$ cannot reach $d$ either. But $\downarrow \mathfrak{P}$ is an asynchronous program over regular languages; this means there exists a safe regular overapproximation and the second semi-decision procedure terminates.

```
Algorithm 2: Checking Termination
    Input: Asynchronous program }\mathfrak{P}=(D,\Sigma,(\mp@subsup{L}{c}{}\mp@subsup{)}{c\in\mathfrak{C}}{},\mp@subsup{d}{0}{},\mp@subsup{\mathbf{m}}{0}{})\mathrm{ over }\mathcal{C
    run concurrently
        begin /* find a terminating overapproximation */
                foreach tuple ( }\mp@subsup{A}{c}{}\mp@subsup{)}{c\in\mathbb{C}}{}\mathrm{ of regular languages }\mp@subsup{A}{c}{}\subseteq\mp@subsup{\Sigma}{}{*}\mathrm{ do
                    if LL
                        if }\mp@subsup{\mathfrak{P}}{}{\prime}=(D,\Sigma,(\mp@subsup{A}{c}{}\mp@subsup{)}{c\in\mathfrak{C}}{},\mp@subsup{d}{0}{},\mp@subsup{\mathbf{m}}{0}{})\mathrm{ terminates then
                                return \mathfrak{P}\mathrm{ terminates.}
        begin /* find a self-covering run */
            foreach prerun \rho of }\mathfrak{P}\mathrm{ do
                if \rho is a self-covering run then
                    return }\mathfrak{P}\mathrm{ does not terminate.
```

To prove "(iii) $\Rightarrow$ (ii)", we adopt a similar method as for safety. Algorithm 2 for termination consists of two semi-decision procedures. By standard well-quasi-ordering arguments, an infinite run of an asynchronous program $\mathfrak{P}$ is witnessed by a finite self-covering run. The first semi-decision procedure enumerates finite self-covering runs (trying to show non-termination). The second procedure enumerates regular asynchronous programs $\mathfrak{P}^{\prime}$ that overapproximate $\mathfrak{P}$. As before, to check termination of $\mathfrak{P}^{\prime}$, it applies the procedure from Theorem 4.1. Clearly, the algorithm's answer is always correct. Moreover, it gives an answer for every input. If $\mathfrak{P}$ does not terminate, it will find a self-covering sequence. If $\mathfrak{P}$ does terminate, then Proposition 3.3 tells us that $\downarrow \mathfrak{F}$ is a terminating finite-state overapproximation. This implies that the second procedure will terminate in that case.

Let us point out a particular example. The class $\mathcal{L}$ of languages of lossy channel systems is defined like the class of languages of Well-Structured Transition Systems (WSTS) with upward-closed sets of accepting configurations as in [GRB07], except that we only consider lossy channel systems [ABJ98] instead of arbitrary WSTS. Then $\mathcal{L}$ forms a full trio with decidable emptiness. Although downclosures of lossy channel languages are not effectively computable (an easy consequence of [May03]), our algorithm employs Theorem 4.2 to decide safety and termination.

### 4.2. Boundedness.

Theorem 4.3. Let $\mathcal{C}$ be a full trio. The following are equivalent:
(i) Boundedness is decidable for asynchronous programs over $\mathcal{C}$.
(ii) Finiteness is decidable for $\mathcal{C}$.

Proof. Clearly, the construction for "(i) $\Rightarrow$ (iii)" of Theorem 4.2 also works for "(i) $\Rightarrow$ (ii)": $\mathfrak{P}$ is unbounded iff $K$ is infinite.

For the converse, we first note that if finiteness is decidable for $\mathcal{C}$ then so is emptiness. Given $L \subseteq \Sigma^{*}$ from $\mathcal{C}$, consider the homomorphism $h:(\Sigma \cup\{\lambda\})^{*} \rightarrow \Sigma^{*}$ with $h(\mathrm{a})=\mathrm{a}$ for every a $\in \Sigma$ and $h(\lambda)=\varepsilon$. Then $h^{-1}(L)$ belongs to $\mathcal{C}$ and $h^{-1}(L)$ is finite if and only if $L$ is empty: in the inverse homomorphism, $\lambda$ can be arbitrarily inserted in any word. By Theorem 4.2, this implies that we can also decide safety. As a consequence of considering only full trios, it is easy to see that the problem of context reachability reduces to safety: a

```
Algorithm 3: Checking Boundedness
    Input: Asynchronous program \(\mathfrak{P}=\left(D, \Sigma,\left(L_{c}\right)_{c \in \mathfrak{C}}, d_{0}, \mathbf{m}_{0}\right)\) over \(\mathcal{C}\)
    \(F \leftarrow \emptyset, I \leftarrow \emptyset ;\)
    foreach context \(c=\left(d, \sigma, d^{\prime}\right) \in \mathfrak{C}\) do
        if \(L_{c}\) is infinite then
            if \(c\) is reachable in \(\mathfrak{P}\) then /* using algorithm for safety */
                return \(\mathfrak{P}\) is unbounded.
            \(I \leftarrow I \cup\{c\}\)
        else
            \(F \leftarrow F \cup\{c\}\)
    foreach context \(c \in F\) do
        foreach finite set \(A \subseteq \Sigma^{*}\) do /* find a finite \(A\) with \(L_{c}=A * /\)
            if \(L_{c} \cap\left(\Sigma^{*} \backslash A\right)=\emptyset\) and \(L_{c} \cap\{w\} \neq \emptyset\) for each \(w \in A\) then
                \(A_{c} \leftarrow A ;\)
                break; /* end inner foreach */
    foreach context \(c \in I\) do
        \(A_{c} \leftarrow \emptyset\)
    if \(\mathfrak{P}^{\prime}=\left(D, \Sigma,\left(A_{c}\right)_{c \in \mathfrak{C}}, d_{0}, \mathbf{m}_{0}\right)\) is bounded then
        return \(\mathfrak{P}\) is bounded.
    else
        return \(\mathfrak{P}\) is unbounded.
```

context $c=\left(d, \sigma, d^{\prime}\right) \in \mathfrak{C}$ is reachable in $\mathfrak{P}$ if there is a reachable configuration $(d, \mathbf{m})$ in $\mathfrak{P}$ with $\mathbf{m}(\sigma) \geq 1$. This latter condition can be checked by moving from $d$ to a special state via a $\sigma$ transition.

We now explain Algorithm 3 for deciding boundedness of a given aysnchronous program $\mathfrak{P}=\left(D, \Sigma,\left(L_{c}\right)_{c \in \mathfrak{C}}, d_{0}, \mathbf{m}_{0}\right)$. For every context $c$, we first check if $L_{c}$ is infinite (feasible by assumption). This paritions the set of contexts of $\mathfrak{P}$ into sets $I$ and $F$ which are the contexts for which the corresponding language $L_{c}$ is infinite and finite respectively. If any context in $I$ is reachable, then $\mathfrak{P}$ is unbounded. Otherwise, all the reachable contexts have a finite language. For every finite language $L_{c}$ for some $c \in F$, we explicitly find all the members of $L_{c}$. This is possible because any finite set $A$ can be checked with $L_{c}$ for equality. $L_{c} \subseteq A$ can be checked by testing whether $L_{c} \cap\left(\Sigma^{*} \backslash A\right)=\emptyset$ and $L_{c} \cap\left(\Sigma^{*} \backslash A\right)$ effectively belongs to $\mathcal{C}$. On the other hand, checking $A \subseteq L_{c}$ just means checking whether $L_{c} \cap\{w\} \neq \emptyset$ for each $w \in A$, which can be done the same way. We can now construct asynchronous program $\mathfrak{P}^{\prime}$ which replaces all languages for contexts in $I$ by $\emptyset$ and replaces those corresponding to $F$ by the explicit description. Clearly $\mathfrak{P}^{\prime}$ is bounded iff $\mathfrak{P}$ is bounded (since no contexts from $I$ are reachable) and the former can be decided by Theorem 4.1.

We observe that boundedness is strictly harder than safety or termination: There are full trios for which emptiness is decidable, but finiteness is undecidable, such as the languages of reset vector addition systems [DFS98] (see [TZ19] for a definition of the language class) and languages of lossy channel systems.
4.3. Configuration reachability and liveness properties. Theorems 4.2 and 4.3 completely characterize for which full trios safety, termination, and boundedness are decidable.

We turn to configuration reachability, fair termination, and fair starvation. We suspect that it is unlikely that there is a simple characterization of those language classes for which the latter problems are decidable. However, we show that they are decidable only for a limited range of infinite-state systems. To this end, we prove that decidability of any of these problems implies (a) decidability of the others as well, and also implies (b) the decidability of a simple combinatorial problem (called the $Z$-intersection problem) that is known to be undecidable for many expressive classes of languages.

Let $Z \subseteq\{\mathrm{a}, \mathrm{b}\}^{*}$ be the language $Z=\left\{\left.w \in\{\mathrm{a}, \mathrm{b}\}^{*}| | w\right|_{\mathrm{a}}=|w|_{\mathrm{b}}\right\}$. The $Z$-intersection problem for a language class $\mathcal{C}$ asks, given a language $K \subseteq\{\mathrm{a}, \mathrm{b}\}^{*}$ from $\mathcal{C}$, whether $K \cap Z \neq \emptyset$. Informally, $Z$ is the language of all words with an equal number of as and bs and the $Z$ intersection problem asks if there is a word in $K$ with an equal number of as and bs.
Theorem 4.4. Let $\mathcal{C}$ be a full trio. The following statements are equivalent:
(i) Configuration reachability is decidable for asynchronous programs over $\mathcal{C}$.
(ii) Fair termination is decidable for asynchronous programs over $\mathcal{C}$.
(iii) Fair starvation is decidable for asynchronous programs over $\mathcal{C}$.

Moreover, if decidability holds, then $Z$-intersection is decidable for $\mathcal{C}$.
Proof. We prove Theorem 4.4 by providing reductions among the three problems and showing that $Z$-intersection reduces to configuration reachability. We use diagrams similar to automata to describe asynchronous programs. Here, circles represent global states of the program and we draw an edge

$$
\text { (d) } \xrightarrow{\sigma \mid L}\left(d^{\prime}\right)
$$

in case we have $L_{d, \sigma, d^{\prime}}=L$ in our asynchronous program $\mathfrak{P}$. Furthermore, we have $L_{d, \sigma, d^{\prime}}=\emptyset$ whenever there is no edge that specifies otherwise. To simplify notation, we draw an edge (d) $\xrightarrow{w \mid L} d^{\prime}$ in an asynchronous program for a word $w \in \Sigma^{*}, w=\sigma_{1} \ldots \sigma_{n}$ with $\sigma_{1}, \ldots, \sigma_{n} \in \Sigma$, to symbolize a sequence of states

which removes $\llbracket \sigma_{1}, \ldots, \sigma_{n} \rrbracket$ from the task buffer and posts a multiset of handlers specified by the language $L$.
Proof of " $(i i) \Rightarrow(i)$ " Given an asynchronous program $\mathfrak{P}=\left(D, \Sigma,\left(L_{c}\right)_{c \in \mathfrak{C}}, d_{0}, \mathbf{m}_{0}\right)$ and a configuration $\left(d_{f}, \mathbf{m}_{f}\right) \in D \times \mathbb{M}[\Sigma]$, we construct asynchronous program $\mathfrak{P}^{\prime}$ as follows. Let $\mathbf{z}$ be a fresh letter and let $\mathbf{m}_{f}=\llbracket \sigma_{1}, \ldots, \sigma_{n} \rrbracket$. We obtain $\mathfrak{P}^{\prime}$ from $\mathfrak{P}$ by adding a new state $d_{f}^{\prime}$ and including the following edges:


Starting from $\left(d_{0}, \mathbf{m}_{0} \oplus \llbracket z \rrbracket\right)$, the program $\mathfrak{P}^{\prime}$ has a fair infinite run iff $\left(d_{f}, \mathbf{m}_{f}\right)$ is reachable in $\mathfrak{P}$. The 'if' direction is obvious. Conversely, $z$ has to be executed in any fair run $\rho$ of $\mathfrak{P}^{\prime}$ which implies that $d_{f}^{\prime}$ is reached by $\mathfrak{P}^{\prime}$ in $\rho$. Since only z can be executed at $d_{f}^{\prime}$ in $\rho$, this means that the multiset is exactly $\mathbf{m}_{f}$ when $d_{f}$ is reached during $\rho$. Clearly this initial segment of $\rho$ corresponds to a run of $\mathfrak{P}$ which reaches the target configuration.
Proof of " $(i i i) \Rightarrow(i i)$ " We construct $\mathfrak{P}^{\prime}=\left(D, \Sigma^{\prime},\left(L_{c}^{\prime}\right)_{c \in \mathfrak{C}^{\prime}}, d_{0}, \mathbf{m}_{0}^{\prime}\right)$ given

$$
\mathfrak{P}=\left(D, \Sigma,\left(L_{c}\right)_{c \in \mathfrak{C}}, d_{0}, \mathbf{m}_{0}\right)
$$

over $\mathcal{C}$ as follows. Let $\Sigma^{\prime}=\Sigma \cup\{\mathbf{s}\}$, where $\mathbf{s}$ is a fresh handler. Replace each edge

$$
(d) \xrightarrow{\sigma \mid L}\left(d^{\prime}\right) \text { by }(d) \xrightarrow{\sigma \mid L \cup L \mathbf{s}}\left(d^{\prime}\right) \curvearrowright \mathbf{s} \mid \varepsilon
$$

at every state $d \in D$. Moreover, we set $\mathbf{m}_{0}^{\prime}=\mathbf{m}_{0} \oplus \llbracket \mathbf{s} \rrbracket$. Then $\mathfrak{P}^{\prime}$ has an infinite fair run that starves some handler if and only if $\mathfrak{P}$ has an infinite fair run. From an infinite fair run $\rho$ of $\mathfrak{P}$, we obtain an infinite fair run of $\mathfrak{P}^{\prime}$ which starves $\mathbf{s}$, by producing $\mathbf{s}$ while simulating $\rho$ and consuming it in the loop. Conversely, from an infinite fair run $\rho^{\prime}$ of $\mathfrak{P}^{\prime}$ which starves some $\tau$, we obtain an infinite fair run $\rho$ of $\mathfrak{P}$ by omitting all productions and consumptions of s and removing two extra instances of s from all configurations.
Proof of " $(i) \Rightarrow($ iii $)$ " From $\mathfrak{P}=\left(D, \Sigma,\left(L_{c}\right)_{c \in \mathfrak{C}}, d_{0}, \mathbf{m}_{0}\right)$ over $\mathcal{C}$, for each subset $\Gamma \subseteq \Sigma$ and $\tau \in \Sigma$, we construct an asynchronous program $\mathfrak{P}_{\Gamma, \tau}=\left(D^{\prime}, \Sigma^{\prime},\left(L_{c}\right)_{c \in \mathfrak{C}^{\prime}}, d_{0}^{\prime}, \mathbf{m}_{0}^{\prime}\right)$ over $\mathcal{C}$ such that a particular configuration is reachable in $\mathfrak{P}_{\Gamma, \tau}$ if and only if $\mathfrak{P}$ has a fair infinite run $\rho_{\Gamma, \tau}$, where $\Gamma$ is the set of handlers that is executed infinitely often in $\rho_{\Gamma, \tau}$ and $\rho_{\Gamma, \tau}$ starves $\tau$. Since there are only finitely many choices for $\Gamma$ and $\tau$, decidability of configuration reachability implies decidability of fair starvation. The idea is that a run $\rho_{\Gamma, \tau}$ exists if and only if there exists a run

$$
\begin{equation*}
\left(d_{0}, \mathbf{m}_{0}\right) \xrightarrow{\sigma_{1}} \cdots \xrightarrow{\sigma_{n}}\left(d_{n}, \mathbf{m}_{n}\right)=\left(e_{0}, \mathbf{n}_{0}\right) \xrightarrow{\gamma_{1}}\left(e_{1}, \mathbf{n}_{1}\right) \xrightarrow{\gamma_{2}} \cdots \xrightarrow{\gamma_{k}}\left(e_{k}, \mathbf{n}_{k}\right), \tag{4.1}
\end{equation*}
$$

where $\bigcup_{i=1}^{k}\left\{\gamma_{i}\right\}=\Gamma$, for each $1 \leq i \leq k \mathbf{n}_{i} \in \mathbb{M}[\Gamma], \mathbf{m}_{n} \preceq \mathbf{n}_{k}$, and for each $i \in\{1, \ldots, k\}$ with $\gamma_{i}=\tau$, we have $\mathbf{n}_{i-1}(\tau) \geq 2$. In such a run, we call $\left(d_{0}, \mathbf{m}_{0}\right) \xrightarrow{\sigma_{1}} \cdots \xrightarrow{\sigma_{n}}\left(d_{n}, \mathbf{m}_{n}\right)$ its first phase and $\left(e_{0}, \mathbf{n}_{0}\right) \xrightarrow{\gamma_{1}} \cdots \xrightarrow{\gamma_{k}}\left(e_{k}, \mathbf{n}_{k}\right)$ its second phase.

Let us explain how $\mathfrak{P}_{\Gamma, \tau}$ reflects the existence of a run as in Eq. (4.1). The set $\Sigma^{\prime}$ of handlers of $\mathfrak{P}_{\Gamma, \tau}$ includes $\Sigma, \bar{\Sigma}$ and $\hat{\Sigma}$, where $\bar{\Sigma}=\{\bar{\sigma} \mid \sigma \in \Sigma\}$ and $\hat{\Sigma}=\{\hat{\sigma} \mid \sigma \in \Sigma\}$ are disjoint copies of $\Sigma$. This means, a multiset $\mathbb{M}\left[\Sigma^{\prime}\right]$ contains multisets $\mathbf{m}^{\prime}=\mathbf{m} \oplus \overline{\mathbf{m}} \oplus \hat{\mathbf{m}}$ with $\mathbf{m} \in \mathbb{M}[\Sigma], \overline{\mathbf{m}} \in \mathbb{M}[\bar{\Sigma}]$, and $\hat{\mathbf{m}} \in \mathbb{M}[\hat{\Sigma}]$. A run of $\mathfrak{P}_{\Gamma, \tau}$ simulates the two phases of $\rho$. While simulating the first phase, $\mathfrak{P}_{\Gamma, \tau}$ keeps two copies of the task buffer, $\mathbf{m}$ and $\overline{\mathbf{m}}$. The copying is easily accomplished by a homomorphism with $\sigma \mapsto \sigma \bar{\sigma}$ for each $\sigma \in \Sigma$. At some point, $\mathfrak{P}_{\Gamma, \tau}$ switches into simulating the second phase. There, $\overline{\mathbf{m}}$ remains unchanged, so that it stores the value of $\mathbf{m}_{n}$ in Eq. (4.1) and can be used in the end to make sure that $\mathbf{m}_{n} \preceq \mathbf{n}_{k}$.

Hence, in the second phase, $\mathfrak{P}_{\Gamma, \tau}$ works, like $\mathfrak{P}$, only with $\Sigma$. However, whenever a handler $\sigma \in \Sigma$ is executed, it also produces a task $\hat{\sigma}$. These handlers are used at the end to make sure that every $\gamma \in \Gamma$ has been executed at least once in the second phase. Also, whenever $\tau$ is executed, $\mathfrak{P}_{\Gamma, \tau}$ checks that at least two instances of $\tau$ are present in the task buffer, thereby ensuring that $\tau$ is starved.

In the end, a distinguished final state allows $\mathfrak{P}_{\Gamma, \tau}$ to execute handlers in $\Gamma$ and $\bar{\Gamma}$ simultaneously to make sure that $\mathbf{m}_{n} \preceq \mathbf{n}_{k}$. In its final state, $\mathfrak{P}_{\Gamma, \tau}$ can execute handlers $\hat{\gamma} \in \hat{\Gamma}$ and $\gamma \in \Gamma$ (without creating new handlers). In the final configuration, there can be no $\hat{\sigma}$ with $\sigma \in \Sigma \backslash \Gamma$, and there has to be exactly one $\hat{\gamma}$ for each $\gamma \in \Gamma$. This guarantees that (i) each handler in $\Gamma$ is executed at least once during the second phase, (ii) every handler executed in the second phase is from $\Gamma$, and (iii) $\mathbf{m}_{n}$ contains only handlers from $\Gamma$ (because handlers from $\bar{\Sigma}$ cannot be executed in the second phase).

Let us now describe $\mathfrak{P}_{\Gamma, \tau}$ in detail. We have $\Sigma^{\prime}=\Sigma \cup \bar{\Sigma} \cup \hat{\Gamma} \cup\{\mathbf{z}\}$, where $\mathbf{z}$ is a fresh letter. The set of states is $D^{\prime}=D \cup \tilde{D} \cup\left\{d_{f}\right\}$, where $\tilde{D}=\{\tilde{d} \mid d \in D\}$ is a disjoint copy of $D$ for simulating the second phase, and $d_{f}$ is a fresh state. Moreover, we change the languages
$L_{c}$ in the following way:

$$
\text { (d) } \xrightarrow{\sigma \mid L}\left(d^{\prime}\right) \rightsquigarrow(d) \xrightarrow{\sigma \bar{\sigma} \mid h^{ \pm}(L)}\left(d^{\prime}\right)
$$

where $h^{ \pm}: \Sigma^{*} \rightarrow(\Sigma \cup \bar{\Sigma})^{*}$ is the homomorphism with $\sigma \mapsto \sigma \bar{\sigma}$ for every $\sigma \in \Sigma$. This means, for every context $c=\left(d, \sigma, d^{\prime}\right)$, we have $d \xrightarrow{\sigma \bar{\sigma} \mid h^{ \pm}(L)} d^{\prime}$ in $\mathfrak{P}_{\Gamma, \tau}$. Note that since $\mathcal{C}$ is a full trio, it is in particular closed under homomorphisms. Hence, $h^{ \pm}(L)$ belongs to $\mathcal{C}$. Moreover, $\mathfrak{P}_{\Gamma, \tau}$ can spontaneously switch to simulating the second phase: For each $d \in D$, we have

$$
\text { (d) } \xrightarrow{z \mid\{z\}} \text { ( }(\hat{d})
$$

Here, the handler z merely allows us to move to state $\tilde{d}$ without interfering with the other handlers. In the second phase, $\mathfrak{P}_{\Gamma, \tau}$ simulates $\mathfrak{P}$ slightly differently. We perform the following replacement:

$$
(d) \xrightarrow{\sigma \mid L}\left(d^{\prime}\right) \rightsquigarrow \begin{cases}(\hat{d}) \xrightarrow{\sigma \mid L \hat{\sigma}} \longrightarrow  \tag{4.2}\\ (\hat{d}) \xrightarrow{\tau \tau \mid L \tau \hat{\tau}} \xrightarrow{\left(d^{\prime}\right)} & \text { if } \sigma \neq \tau \\ \left(d^{\prime}\right) & \text { if } \sigma=\tau\end{cases}
$$

Note that since $\mathcal{C}$ is a full trio, the languages $L v=\{u v \mid u \in L\}$ belong to $\mathcal{C}$ for every word $v$. Finally, $\mathfrak{P}_{\Gamma, \tau}$ can spontaneously switch to its distinguished state $d_{f}$, so that we have for every $\tilde{d} \in \tilde{D}$ and every $\gamma \in \Gamma$ :


The initial configuration of $\mathfrak{P}_{\Gamma, \tau}$ is $\left(d_{0}, \mathbf{m}_{0}^{\prime}\right)$, where $\mathbf{m}_{0}^{\prime}=\mathbf{m}_{0} \oplus \overline{\mathbf{m}}_{0} \oplus \llbracket z \rrbracket$, where $\overline{\mathbf{m}}_{0}$ is obtained from $\mathbf{m}_{0}$ by replacing every occurrence of $\sigma \in \Sigma$ in $\mathbf{m}_{0}$ with $\bar{\sigma}$. The final configuration is $\left(d_{f}, \mathbf{m}_{f}\right)$, where $\mathbf{m}_{f} \in \mathbb{M}[\hat{\Gamma}]$ is the multiset with $\mathbf{m}_{f}(\hat{\gamma})=1$ for every $\gamma \in \Gamma$. It is now straightforward to check that $\left(d_{f}, \mathbf{m}_{f}\right)$ is reachable in $\mathfrak{P}_{\Gamma, \tau}$ if and only if $\mathfrak{P}$ has an infinite fair run that starves $\tau$.
Decidability of $Z$-intersection To complete the proof of Theorem 4.4, we reduce $Z$-intersection to configuration reachability. Given $K \subseteq\{\mathrm{a}, \mathrm{b}\}^{*}$ from $\mathcal{C}$, we construct the asynchronous program $\mathfrak{P}=\left(D, \Sigma,\left(L_{c}\right)_{c \in \mathfrak{C}}, d_{0}, \mathbf{m}_{0}\right)$ over $\mathcal{C}$ where $D=\left\{d_{0}, 0,1\right\}, \Sigma=\{\mathrm{a}, \mathrm{b}, \mathrm{c}\}$, by including the following edges:


The initial task buffer is $\mathbf{m}_{0}=\llbracket \subset \rrbracket$. Then clearly, the configuration $(0, \llbracket \rrbracket)$ is reachable in $\mathfrak{P}$ if and only if $K \cap Z \neq \emptyset$.

If the construction seems abstract, recall the example from Section 2: the procedure s1 plays the role of $K$ and generates strings from its language in $\{\mathrm{a}, \mathrm{b}\}^{*}$. The procedures a and b take turns to ensure there is an equal number of them; the states 0 and 1 are values of turn.

Theorem 4.4 is useful in the contrapositive to show undecidability. For example, one can show undecidability of $Z$-intersection for languages of lossy channel systems (see Section 4.1): One expresses reachability in a non-lossy FIFO system by making sure that the numbers of enqueue- and dequeue-operations match. Thus, for asynchronous programs over lossy channel systems, the problems of Theorem 4.4 are undecidable. We also use Theorem 4.4 in Section 5 to conclude undecidability for higher-order asynchronous programs, already at order 2.

## 5. Higher-Order Asynchronous Programs

We apply our general decidability results to asynchronous programs over (deterministic) higher-order recursion schemes (HORS). Kobayashi [Kob09] has shown how higher-order functional programs can be modeled using HORS. In his setting, a program contains instructions that access certain resources. The path language of the HORS produced by Kobayashi is the set of possible sequences of instructions. For us, the input program contains post instructions (which post handlers i.e. add tasks to the task buffer) and we translate higher-order programs with post instructions into a HORS whose path language is used as the language of handlers.

We recall some definitions from [Kob09]. The set of types is defined by the grammar $A:=\circ \mid A \rightarrow A$. The order $\operatorname{ord}(A)$ of a type $A$ is inductively defined as ord $(0)=0$ and $\operatorname{ord}(A \rightarrow B):=\max (\operatorname{ord}(A)+1, \operatorname{ord}(B))$. The arity of a type is inductively defined by $\operatorname{arity}(\mathrm{o})=0$ and $\operatorname{arity}(A \rightarrow B)=\operatorname{arity}(B)+1$. We assume a countably infinite set Var of typed variables $x: A$. For a set $\Theta$ of typed symbols, the set $\tilde{\Theta}$ of terms generated from $\Theta$ is the least set which contains $\Theta$ such that whenever $s: A \rightarrow B$ and $t: A$ belong to $\tilde{\Theta}$, then also $s t: B$ belongs to $\tilde{\Theta}$. By convention the type $o \rightarrow \ldots(0 \rightarrow(0 \rightarrow 0))$ is written $\circ \rightarrow \ldots \rightarrow \circ \rightarrow 0$ and the term $\left(\left(t_{1} t_{2}\right) t_{3} \cdots\right) t_{n}$ is written $t_{1} t_{2} \cdots t_{n}$. We write $\bar{x}$ for a sequence $\left(x_{1}, x_{2}, \ldots, x_{n}\right)$ of variables.

A higher-order recursion scheme (HORS) is a tuple $\mathscr{S}=(\Sigma, \mathcal{N}, \mathcal{R}, S)$ where $\Sigma$ is a set of typed terminal symbols of types of order 0 or $1, \mathcal{N}$ is a set of typed non-terminal symbols (disjoint from terminal symbols), $S$ : o is the start non-terminal symbol and $\mathcal{R}$ is a set of rewrite rules $F x_{1} x_{2} \cdots x_{n} \rightarrow t$ where $F: A_{1} \rightarrow \cdots \rightarrow A_{n} \rightarrow$ o is a non-terminal in $\mathcal{N}, x_{i}: A_{i}$ for all $i$ are variables and $t: o$ is a term generated from $\Sigma \cup \mathcal{N} \cup\left\{x_{1}, x_{2}, \ldots, x_{n}\right\}$. The order of a HORS is the maximum order of a non-terminal symbol. We define a rewrite relation $\rightarrow$ on terms over $\Sigma \cup \mathcal{N}$ as follows: $F \bar{a} \rightarrow t[\bar{x} / \bar{a}]$ if $F \bar{x} \rightarrow t \in \mathcal{R}$, and if $t \rightarrow t^{\prime}$ then $t s \rightarrow t^{\prime} s$ and $s t \rightarrow s t^{\prime}$. The reflexive, transitive closure of $\rightarrow$ is denoted $\rightarrow^{*}$. A sentential form $t$ of $\mathscr{S}$ is a term over $\Sigma \cup \mathcal{N}$ such that $S \rightarrow^{*} t$.

If $N$ is the maximum arity of a symbol in $\Sigma$, then a (possibly infinite) tree over $\Sigma$ is a partial function $\operatorname{tr}$ from $\{0,1, \ldots, N-1\}^{*}$ to $\Sigma$ that fulfills the following conditions: $\varepsilon \in \operatorname{dom}(\operatorname{tr}), \operatorname{dom}(\operatorname{tr})$ is closed under prefixes, and if $\operatorname{tr}(w)=a$ and $\operatorname{arity}(a)=k$ then $\{j \mid w j \in \operatorname{dom}(\operatorname{tr})\}=\{0,1, \ldots, k-1\}$.

A deterministic HORS is one where there is exactly one rule of the form $F x_{1} x_{2} \cdots x_{n} \rightarrow t$ for every non-terminal $F$. Following [Kob09], we show how a deterministic HORS can be used to represent a higher-order pushdown language arising from a higher-order functional program. Sentential forms can be seen as trees over $\Sigma \cup \mathcal{N}$. A sequence $\Pi$ over $\{0,1, \ldots, N-1\}$ is a path of $t r$ if every finite prefix of $\Pi \in \operatorname{dom}(t r)$ and $\Pi$ is maximal. The set of finite paths in a tree $t r$ will be denoted Paths $(t r)$. Unless otherwise specified, a path is finite in our
context. Associated with any path $\Pi=n_{1}, n_{2}, \ldots, n_{k}$ is the word

$$
w_{\Pi}=\operatorname{tr}\left(n_{1}\right) \operatorname{tr}\left(n_{1} n_{2}\right) \cdots \operatorname{tr}\left(n_{1} n_{2} \cdots n_{k}\right) .
$$

We will also write $w_{\Pi}(\bar{m})$ to denote the letter $\operatorname{tr}(\bar{m})$ where $\bar{m}=n_{1} n_{2} \ldots n_{l}$ for some $l \leq k$. We associate a value tree $\mathcal{T}_{\mathscr{S}}$ associated with a deterministic HORS $\mathscr{S}$ in the following way. For a sentential form $t$ define the finite tree $t^{\perp}$ by case analysis: $t^{\perp}=f$ if $f$ is a terminal symbol and $t^{\perp}=t_{1}^{\perp} t_{2}^{\perp}$ if $t=t_{1} t_{2}$ and $t_{1}^{\perp} \neq \perp$, and $t^{\perp}=\perp$ otherwise.

Intuitively, the tree $t^{\perp}$ is obtained by replacing any subterm whose root is a non-terminal by the nullary symbol $\perp$. We define the partial order $\leq_{t}$ on $\operatorname{dom}(\Sigma) \cup\{\perp\}$ by $\perp \leq_{t} a$ for all $a \in \operatorname{dom}(\Sigma)$, which is extended to trees by

$$
t \leq_{t} s \Longleftrightarrow \forall \bar{n} \in \operatorname{dom}(t): \bar{n} \in \operatorname{dom}(s) \wedge t(\bar{n}) \leq_{t} s(\bar{n})
$$

The value tree generated by a HORS $\mathscr{S}$ is denoted $\mathcal{T}_{\mathscr{S}}$ and is obtained by repeated application of the rewrite rules to the start symbol $S$. To make this formal, we write $\operatorname{lub}(T)$ for the least upper bound of a collection $T$ of trees with respect to the order $\leq_{t}$. Then we set $\mathcal{T}_{\mathscr{S}}:=\operatorname{lub}\left(\left\{t^{\perp} \mid S \rightarrow^{*} t\right\}\right)$. Any HORS for which the value tree $\mathcal{T}_{\mathscr{S}}$ does not contain any $\perp$ symbol is called productive. All HORS dealt with in this paper are assumed to be productive. Checking if a given HORS is productive is decidable, see, e.g., [Gre16].

Let $\Sigma_{1}:=\{a \in \Sigma \mid \operatorname{arity}(a)=1\}$. The path language $\mathcal{L}_{\mathrm{p}}(\mathscr{S})$ of a deterministic HORS $\mathscr{S}$ is defined as $\left\{\operatorname{Proj}_{\Sigma_{1}}\left(w_{\Pi}\right) \mid \Pi \in \operatorname{Paths}\left(\mathcal{T}_{\mathscr{S}}\right)\right\}$. The tree language $\mathcal{L}_{\mathrm{t}}(\mathscr{S})$ associated with a (nondeterministic) HORS is the set of sentential forms of $\mathscr{S}$ which contain only labels from $\Sigma$.

The deterministic HORS corresponding to the higher-order function s4 from Figure 1 is given by $\mathscr{S}=(\Sigma, \mathcal{N}, \mathcal{R}, S)$, where

$$
\begin{aligned}
\Sigma & =\{\mathrm{br}: \mathrm{o} \rightarrow \mathrm{o} \rightarrow \mathrm{o}, \mathrm{c}, \mathrm{~d}, \mathrm{f}: \mathrm{o} \rightarrow \mathrm{o}, \mathrm{e}: \mathrm{o}\} \\
\mathcal{N} & =\{S: \mathrm{o}, F:(\mathrm{o} \rightarrow \mathrm{o}) \rightarrow \mathrm{o} \rightarrow \mathrm{o}, H:(\mathrm{o} \rightarrow \mathrm{o}) \rightarrow \mathrm{o} \rightarrow \mathrm{o}, I: \mathrm{o} \rightarrow \mathrm{o}\} \\
\mathcal{R} & =\{S \rightarrow F I \mathrm{e}, I x \rightarrow x, F G x \rightarrow \mathrm{br}(F(H G)(\mathrm{f} x))(G x), H G x \rightarrow \mathrm{c}(G(\mathrm{~d} x))\}
\end{aligned}
$$

The path language $\mathcal{L}_{\mathrm{p}}(\mathscr{S})=\left\{\mathrm{c}^{n} \mathrm{~d}^{n} \mathrm{f}^{n} \mid n \geq 0\right\}$. To see this, apply the reduction rules to get the value tree $\mathcal{T}_{\mathscr{S}}$ shown on the right:

$$
\begin{aligned}
S & \rightarrow F I \mathrm{e} \rightarrow \mathrm{br}(F(H I)(\mathrm{fe}))(I \mathrm{e}) \\
& \rightarrow \mathrm{br}(F(H I)(\mathrm{fe})) \mathrm{e} \\
& \rightarrow \mathrm{br}(\mathrm{br}(F(H(H I))(\mathrm{f}(\mathrm{fe})))(H I(\mathrm{fe}))) \mathrm{e} \\
& \rightarrow \mathrm{br}(\mathrm{br}(F(H(H I))(\mathrm{f}(\mathrm{fe})))(\mathrm{c}(I(\mathrm{~d}(\mathrm{fe}))))) \mathrm{e} \\
& \rightarrow \mathrm{br}(\mathrm{br}(F(H(H I))(\mathrm{f}(\mathrm{fe})))(\mathrm{c}(\mathrm{~d}(\mathrm{fe})))) \mathrm{e} \\
& \rightarrow \cdots
\end{aligned}
$$



A HORS $\mathscr{S}$ is called a word scheme if it has exactly one nullary terminal symbol e and all other terminal symbols $\tilde{\Sigma}$ are of arity one. The word language $\mathcal{L}_{\mathrm{w}}(\mathscr{S}) \subseteq \tilde{\Sigma}^{*}$ defined by $\mathscr{S}$ is $\mathcal{L}_{\mathrm{w}}(\mathscr{S})=\left\{a_{1} a_{2} \cdots a_{n} \mid\left(a_{1}\left(a_{2} \cdots\left(a_{n}(\mathrm{e})\right) \cdots\right)\right) \in \mathcal{L}_{\mathrm{t}}(\mathscr{S})\right\}$. We denote by $\mathcal{H}$ the class of languages $\mathcal{L}_{\mathrm{w}}(\mathscr{S})$ that occur as the word language of a higher-order recursion scheme $\mathscr{S}$. Note that path languages and languages of word schemes are both word languages
over the set $\tilde{\Sigma}$ of unary symbols considered as letters. They are connected by the following proposition, a proof of which is given in Appendix C. ${ }^{2}$

Proposition 5.1. For every order-n $\operatorname{HORS} \mathscr{S}=(\Sigma, \mathcal{N}, S, \mathcal{R})$ there exists an order-n word scheme $\mathscr{S}^{\prime}=\left(\Sigma^{\prime}, \mathcal{N}^{\prime}, S^{\prime}, \mathcal{R}^{\prime}\right)$ such that $\mathcal{L}_{\mathrm{p}}(\mathscr{S})=\mathcal{L}_{\mathrm{w}}\left(\mathscr{S}^{\prime}\right)$.

A consequence of [Kob09] and Prop. 5.1 is that the "post" language of higher-order functional programs (i.e. the language corresponding to sequences of newly spawned tasks) can be modeled as the language of a word scheme. Hence, we define an asynchronous program over HORS as an asynchronous program over the language class $\mathcal{H}$ and we can use the following results on word schemes.

Theorem 5.2. HORS and word schemes form effective full trios [CPSW16]. Emptiness $[\mathrm{KO11]}$ and finiteness [Par17] of order-n word schemes are ( $n-1$ )-EXPTIME-complete.

Now Theorems 4.2 and 4.3, together with Theorem 5.1 imply the decidability results in Corollary 5.3. The undecidability result is a consequence of Theorem 4.4 and the undecidability of the $Z$-intersection problem for indexed languages or equivalently, order-2 pushdown automata as shown in [Zet15]. Order-2 pushdown automata can be effectively turned into order-2 OI grammars [DG86], which in turn can be translated into order-2 word schemes [Dam82]. A direct (and independent) proof of undecidability of the $Z$-intersection problem for order-2 word schemes has appeared in [Kob19, Theorem 4].

Corollary 5.3. For asynchronous programs over HORS: (1) Safety, termination, and boundedness are decidable. (2) Configuration reachability, fair termination, and fair starvation are undecidable already at order-2.

## 6. A Direct Algorithm and Complexity Analysis

We say that downclosures are computable for a language class $\mathcal{C}$ if for a given description of a language $L$ in $\mathcal{C}$, one can compute an automaton for the regular language $\downarrow L$. A consequence of Proposition 3.3 and Theorem 4.1 is that if one can compute downclosures for some language class, then one can avoid the enumerative approaches of Section 4 and get a "direct algorithm." The direct algorithm replaces each handler by its downclosure and then invokes the decision procedure summarized in Theorem 4.1.
6.1. Higher-Order Recursion Schemes. The direct algorithm for asynchronous programs over HORS relies on the recent breakthrough results on computing downclosures.
Theorem 6.1 [Zet15, HKO16, CPSW16]. Downclosures are effectively computable for $\mathcal{H}$.
Unfortunately, current techniques do not yet provide a complexity upper bound based on the above theorem. To understand why, we describe the current algorithm for computing downclosures. In [Zet15], it was shown that in a full trio $\mathcal{C}$, downclosures are computable if and only if the diagonal problem for $\mathcal{C}$ is decidable. The latter asks, given a language $L \subseteq \Sigma^{*}$, whether for every $k \in \mathbb{N}$, there is a word $w \in L$ with $|w|_{\sigma} \geq k$ for every $\sigma \in \Sigma$. The diagonal problem was then shown to be decidable for higher-order pushdown automata [HKO16] and

[^3]then for word schemes [CPSW16]. However, the algorithm from [Zet15] to compute downclosures using an oracle for the diagonal problem employs enumeration to compute a downclosure automaton. Thus, the enumeration is hidden inside the downclosure computation.

We conjecture that downclosures can in fact be computed in elementary time (for word schemes of fixed order). This would imply an elementary time procedure for asynchronous programs over HORS as well.
6.2. Context Free Languages. For handlers over context-free languages, given e.g., as pushdown automata, Ganty and Majumdar show an EXPSPACE upper bound. Precisely, the algorithm of [GM12] constructs for each handler a polynomial-size Petri net with certain guarantees (forming a so-called adequate family of Petri nets) that accepts a Parikh equivalent language. These Petri nets are then used to construct a larger Petri net, polynomial in the size of the asynchronous program and the adequate family of Petri nets, in which the respective property (safety, boundedness, or termination) can be phrased as a query decidable in EXPSPACE. A natural question is whether there is a downclosure-based algorithm with the same asymptotic complexity.

Our goal is to replace the Parikh-equivalent Petri nets with Petri nets recognizing the downclosure of a language. It is an easy consequence of Proposition 3.3 that the resulting Petri nets can be used in place of the adequate families of Petri nets in the procedures for safety, termination, and boundedness of [GM12]. Thus, if we can ensure these Petri nets are polynomial in the size of the handler, we get an EXPSPACE upper bound. Unfortunately, a finite automaton for $\downarrow L$ may require exponentially many states in the pushdown automaton [BLS15]. Thus a naive approach gives a 2EXPSPACE upper bound.

We show here that for each context-free language $L$, one can construct in polynomial time a 1 -bounded Petri net accepting $\downarrow L$. (Recall that a Petri net is 1 -bounded if every reachable marking has at most one token in each place.) This is a language theoretic result of independent interest. When used in the construction of [GM12], this matches the EXPSPACE upper bound.

As a byproduct, our translation yields a simple direct construction of a finite automaton for $\downarrow L$ when $L$ is given as a pushdown automaton. This is of independent interest because earlier constructions of $\downarrow L$ always start from a context-free grammar and produce (of necessity!) exponentially large NFAs [vL78, Cou91, BLS15].

We begin with some preliminary definitions.
Pushdown automata. If $\Gamma$ is an alphabet, we write $\bar{\Gamma}=\{\bar{\gamma} \mid \gamma \in \Gamma\}$. Moreover, if $x=\bar{\gamma}$, then we define $\bar{x}=\gamma$. For a word $v \in(\Gamma \cup \bar{\Gamma})^{*}, v=v_{1} \cdots v_{n}, v_{1}, \ldots, v_{n} \in \Gamma \cup \bar{\Gamma}$, we set $\bar{v}=\overline{v_{n}} \cdots \overline{v_{1}}$. A pushdown automaton is a tuple $\mathcal{A}=\left(Q, \Sigma, \Gamma, E, q_{0}, q_{f}\right)$, where $Q$ is a finite set of states, $\Sigma$ is its input alphabet, $\Gamma$ is its stack alphabet, $E$ is a finite set of edges, $q_{0} \in Q$ is its initial state, and $F \subseteq Q$ is its set of final states. An edge is a four-tuple ( $p, R, v, q$ ), where $p, q \in Q, R \subseteq \Sigma^{*}$ is a regular language, and $v \in \Gamma \cup \bar{\Gamma} \cup\{\varepsilon\}$. We also write

$$
p \xrightarrow{R \mid v} q
$$

to denote an edge $(p, R, v, q) \in E$. Intuitively, it tells us that from state $q$, we can read any word in $R$ as input and modify the stack as specified by $v$. Here, $v \in \Gamma$ means we push $v$ onto the stack. Moreover, $v \in \Gamma, v=\bar{\gamma}$, means we pop $\gamma$ from the stack. Finally, $v=\varepsilon$ means we do not change the stack. Let us make this formal. A configuration of $\mathcal{A}$ is a pair $(q, w)$ with $q \in Q$ and $w \in \Gamma^{*}$. For configurations $(q, w)$ and $\left(q^{\prime}, w^{\prime}\right)$, we write $(q, w) \xrightarrow{u}\left(q^{\prime}, w^{\prime}\right)$ if there
is an edge $\left(q, R, v, q^{\prime}\right)$ in $\mathcal{A}$ such that $u \in R$ and (i) if $v=\varepsilon$, then $w^{\prime}=w$, (ii) if $v \in \Gamma$, then $w^{\prime}=w v$ and (iii) if $v=\bar{\gamma}$ for $\gamma \in \Gamma$, then $w=w^{\prime} \gamma$.

A run in $\mathcal{A}$ is a sequence $\left(q_{0}, w_{0}\right), \ldots,\left(q_{n}, w_{n}\right)$ of configurations and words $u_{1}, \ldots, u_{n} \in$ $\Sigma^{*}$ such that $\left(q_{i-1}, w_{i-1}\right) \xrightarrow{u_{i}}\left(q_{i}, w_{i}\right)$ for $1 \leq i \leq n$. Its length is $n$ and its initial and final configuration are $\left(q_{0}, w_{0}\right)$ and $\left(q_{n}, w_{n}\right)$, respectively. The run is said to read the word $u_{1} \cdots u_{n}$. The stack height of the run is defined as $\max \left\{\left|w_{i}\right| \mid 0 \leq i \leq n\right\}$. We call the run positive (resp. dually positive) if $\left|w_{i}\right| \geq\left|w_{0}\right|$ (resp. $\left.\left|w_{i}\right| \geq\left|w_{n}\right|\right)$ for every $1 \leq i \leq n$, i.e. if the stack never drops below its initial height (resp. is always above its final height).

We write $(q, w) \stackrel{u}{\Rightarrow}\left(q^{\prime}, w^{\prime}\right)$ for configurations $(q, w),\left(q^{\prime}, w^{\prime}\right)$ if there is a run with initial configuration $(q, w)$ and final configuration $\left(q^{\prime}, w^{\prime}\right)$ that reads $u$. If there is such a run with stack height $\leq h$, then we write $(q, w) \stackrel{u}{\Rightarrow}{ }_{h}\left(q^{\prime}, w^{\prime}\right)$. The language accepted by $\mathcal{A}$ is

$$
\mathrm{L}(\mathcal{A})=\left\{u \in \Sigma^{*} \mid\left(q_{0}, \varepsilon\right) \stackrel{u}{\Rightarrow}\left(q_{f}, \varepsilon\right)\right\} .
$$

There is also a language accepted with bounded stack height. For $h \in \mathbb{N}$, we define

$$
\mathrm{L}_{h}(\mathcal{A})=\left\{u \in \Sigma^{*} \mid\left(q_{0}, \varepsilon\right) \stackrel{u}{\Rightarrow}{ }_{h}\left(q_{f}, \varepsilon\right)\right\} .
$$

Thus, $L_{h}(\mathcal{A})$ is the set of words accepted by $\mathcal{A}$ using runs of stack height at most $h$.
In order to exploit the symmetry between forward and backward computations in pushdown automata, we will consider dual pushdown automata. The dual automaton of $\mathcal{A}$, denoted $\overline{\mathcal{A}}$, is obtained from $\mathcal{A}$ by changing each edge $p \xrightarrow{R \mid v} q$ into $q \xrightarrow{R^{\text {rev }} \mid \bar{v}} p$. Then $\mathrm{L}(\overline{\mathcal{A}})=\mathrm{L}(\mathcal{A})^{\text {rev }}$. We will sometimes argue by duality, which is the principle that every statement that is true for any $\mathcal{A}$ is also true for any $\overline{\mathcal{A}}$.

Petri nets. A (labeled) Petri net is a tuple $N=\left(\Sigma, S, T, \partial_{0}, \partial_{1}, \lambda, \mathbf{m}_{0}, \mathbf{m}_{f}\right)$ where $\Sigma$ is its input alphabet, $S$ is a finite set of places, $T$ is a finite set of transitions, $\partial_{0}, \partial_{1}: T \rightarrow \mathbb{M}[S]$ are maps that specify an input marking $\partial_{0}(t)$ and an output marking $\partial_{1}(t)$ for each transition $t \in T, \lambda: T \rightarrow \Sigma \cup\{\varepsilon\}$ assigns labels to transitions, and $\mathbf{m}_{0}, \mathbf{m}_{f}$ are its initial and final marking. More generally, multisets $\mathbf{m} \in \mathbb{M}[S]$ are called markings of $N$.

For markings $\mathbf{m}_{1}, \mathbf{m}_{2} \in \mathbb{M}[S]$ and $a \in \Sigma \cup\{\varepsilon\}$, we write $\mathbf{m}_{1} \xrightarrow{a} \mathbf{m}_{2}$ if there is a transition $t \in T$ with $\lambda(t)=a, \mathbf{m}_{1} \geq \partial_{0}(t)$, and $\mathbf{m}_{2}=\mathbf{m}_{1} \ominus \partial_{0}(t) \oplus \partial_{1}(t)$. Moreover, we write $\mathbf{m}_{1} \stackrel{w}{\Rightarrow} \mathbf{m}_{2}$ if there are $n \in \mathbb{N}, a_{1}, \ldots, a_{n} \in \Sigma \cup\{\varepsilon\}$, and markings $\mathbf{m}_{0}^{\prime}, \ldots, \mathbf{m}_{n}^{\prime}$ such that $w=a_{1} \cdots a_{n}$ and $\mathbf{m}_{1}=\mathbf{m}_{0}^{\prime} \xrightarrow{a_{1}} \mathbf{m}_{1}^{\prime} \xrightarrow{a_{2}} \cdots \xrightarrow{a_{n}} \mathbf{m}_{n}^{\prime}=\mathbf{m}_{2}$. Furthermore, we write $\mathbf{m}_{1} \Rightarrow \mathbf{m}_{2}$ if there exists a $w \in \Sigma^{*}$ with $\mathbf{m}_{1} \stackrel{w}{\Rightarrow} \mathbf{m}_{2}$. The language accepted by $N$ is $\mathrm{L}(N)=\left\{w \in \Sigma^{*} \mid \mathbf{m}_{0} \stackrel{w}{\Rightarrow} \mathbf{m}_{f}\right\}$.

For $k \in \mathbb{N}$, a Petri net $N$ is $k$-bounded if for every marking $\mathbf{m} \in \mathbb{M}[S]$ with $\mathbf{m}_{0} \Rightarrow \mathbf{m}$, we have $|\mathbf{m}| \leq k$.

Our main results are as follows.
Theorem 6.2 (Succinct Downclosures for PDAs). Given a pushdown automaton $\mathcal{A}$, one can construct in polynomial time a pushdown automaton $\hat{\mathcal{A}}$ so that $\downarrow \mathrm{L}(\hat{\mathcal{A}})=\downarrow \mathrm{L}(\mathcal{A})$. Moreover, we have $\mathrm{L}(\hat{\mathcal{A}})=\mathrm{L}_{h}(\hat{\mathcal{A}})$ for some bound $h$ that is polynomial in the size of $\mathcal{A}$.

As a pushdown automaton with bounded stack can be simulated by a 1-bounded Petri net (essentially, by keeping places for each position in the stack), we get the following corollary and also the promised EXPSPACE upper bound.

Corollary 6.3. Given a pushdown automaton $\mathcal{A}$, one can construct in polynomial time a 1-bounded Petri net $N$ with $\mathrm{L}(N)=\downarrow \mathrm{L}(\mathcal{A})$.

We now prove Theorem 6.2. The augmented automaton $\hat{\mathcal{A}}=\left(Q, \Sigma, \hat{\Gamma}, \hat{E}, q_{0}, q_{f}\right)$ is defined as follows. We first compute for any $p, q \in Q$ the subalphabet $\Delta_{p, q}(\mathcal{A}) \subseteq \Sigma$ with

$$
\Delta_{p, q}(\mathcal{A})=\left\{a \in \Sigma\left|\exists u \in M_{p, q}(\mathcal{A}),|u|_{a} \geq 1\right\}\right.
$$

where $M_{p, q}(\mathcal{A})=\left\{u \in \Sigma^{*} \mid \exists v \in \Gamma^{*}:(p, \varepsilon) \stackrel{u}{\Rightarrow}(p, v),(q, v) \Rightarrow(q, \varepsilon)\right\}$. Note that it is easy to construct in polynomial time a pushdown automaton $\mathcal{A}_{p, q}$ for the language $M_{p, q}(\mathcal{A})$ :
$\mathcal{A}_{p, q}$ has a set $Q_{p} \cup Q_{q}$ of states consisting of two disjoint copies of the states of $\mathcal{A}$. The transitions between two states in $Q_{p}$ are inherited from $\mathcal{A}$ while for two states in $Q_{q}$ we replace the input on any transition by $\varepsilon$ but retain the stack operations from $\mathcal{A}$. The start state is $p \in Q_{p}$ and the final state is $q \in Q_{q}$. There is an $\varepsilon$-transition from $p \in Q_{p}$ to $q \in Q_{q}$. This concludes the construction of $\mathcal{A}_{p, q}$.

Since $a \in \Delta_{p, q}(\mathcal{A})$ iff $M_{p, q}(\mathcal{A}) \cap \Sigma^{*} a \Sigma^{*} \neq \emptyset$, we can decide in polynomial time whether a given $a \in \Sigma$ belongs to $\Delta_{p, q}(\mathcal{A})$ by checking the PDA for $M_{p, q}(\mathcal{A}) \cap \Sigma^{*} a \Sigma^{*} \neq \emptyset$ obtained by product construction for emptiness. Thus, we can compute $\Delta_{p, q}(\mathcal{A})$ in polynomial time. We construct $\hat{\mathcal{A}}$ as follows. For any $p, q \in Q$, we introduce a fresh stack symbol $[p, q]$ and then we add edges

$$
\begin{equation*}
p \xrightarrow{\Delta_{p, q}(\mathcal{A})^{*} \mid[p, q]} p, \tag{6.1}
\end{equation*}
$$

$$
q \xrightarrow{\Delta_{q, p}(\overline{\mathcal{A}})^{*} \mid \overline{[p, q]}} q .
$$

The following lemma tells us that $\mathrm{L}(\hat{\mathcal{A}})$ has the same downward closure as $\mathcal{A}$.
Lemma 6.4. $\mathrm{L}(\mathcal{A}) \subseteq \mathrm{L}(\hat{\mathcal{A}}) \subseteq \downarrow \mathrm{L}(\mathcal{A})$.
Proof. Since the inclusion $\mathrm{L}(\mathcal{A}) \subseteq \mathrm{L}(\hat{\mathcal{A}})$ is obvious, we prove $\mathrm{L}(\hat{\mathcal{A}}) \subseteq \downarrow \mathrm{L}(\mathcal{A})$. We proceed by induction on the number $m$ of executions of new edges (i.e. those from Eq. (6.1)). More specifically, we show that if $u \in \mathrm{~L}(\hat{\mathcal{A}})$ is accepted using $m$ executions of new edges, then there is a $u^{\prime} \in \mathrm{L}(\hat{\mathcal{A}})$ such that $u \sqsubseteq u^{\prime}$ and $u^{\prime}$ is accepted using $<m$ executions of new edges.

Suppose $u$ is accepted using a run $\rho$ with $m>0$ executions of new edges. Then $\rho$ must apply one edge $p \xrightarrow{\Delta_{p, q}(\mathcal{A})[p, q]} p$ and thus also the edge $q \xrightarrow{\Delta_{q, p}(\overline{\mathcal{A}})[\overline{[p, q]}} q$ to remove the letter $[p, q]$ from the stack. Thus, $\rho$ decomposes as $\rho=\rho_{1} \rho_{2} \rho_{3} \rho_{4} \rho_{5}$, where $\rho_{2}$ and $\rho_{4}$ are the executions of the new edges. Let $u=u_{1} u_{2} u_{3} u_{4} u_{5}$ be the corresponding decomposition of $u$.

The run $\rho_{1}$ must end in state $p$ and with some stack content $w \in \Gamma^{*}$. Then, $\rho_{3}$ is a run from $(p, w[p, q])$ to $(q, w[p, q])$ and $\rho_{5}$ is a run from $(q, w)$ to $\left(q_{f}, \varepsilon\right)$ with $q_{f} \in F$.

Since $u_{2}$ and $u_{4}$ are read while executing the new edges, we have $u_{2} \in \Delta_{p, q}(\mathcal{A})^{*}$ and $u_{4} \in \Delta_{q, p}(\overline{\mathcal{A}})^{*}$. We can therefore write $u_{2}=r_{1} \cdots r_{k}$ and $u_{4}=s_{1} \cdots s_{\ell}$ with $r_{1}, \ldots, r_{k} \in$ $\Delta_{p, q}(\mathcal{A})$ and $s_{1}, \ldots, s_{\ell} \in \Delta_{q, p}(\overline{\mathcal{A}})$. By definition, this means for each $1 \leq i \leq k$, there is a word $\tilde{r}_{i} \in M_{p, q}(\mathcal{A})$ that contains the letter $r_{i}$. Likewise, for every $1 \leq i \leq \ell$, there is a $\tilde{s}_{i} \in M_{q, p}(\overline{\mathcal{A}})$ that contains $s_{i}$.

Since $\tilde{r}_{i} \in M_{p, q}(\mathcal{A})$ and $\tilde{s}_{j} \in M_{q, p}(\overline{\mathcal{A}})$ for $1 \leq i \leq k$ and $1 \leq j \leq \ell$, there are words $x_{i}$ and $y_{j}$ in $\Gamma^{*}$ so that

$$
\begin{array}{lll}
(p, \varepsilon) \stackrel{\tilde{r}_{i}}{\Rightarrow}\left(p, x_{i}\right) & \text { and } & \left(q, x_{i}\right) \Rightarrow(q, \varepsilon) \\
(p, \varepsilon) \Rightarrow\left(p, y_{j}\right) & \text { and } & \left(q, y_{i}\right) \stackrel{\tilde{s}_{j}}{\Longrightarrow}(q, \varepsilon)
\end{array}
$$

We can therefore construct a new run $\rho^{\prime}=\rho_{1} \rho_{2}^{\prime} \rho_{3}^{\prime} \rho_{4}^{\prime} \rho_{5}$, where

$$
\begin{aligned}
& \rho_{2}^{\prime}:(p, w) \stackrel{\tilde{r}_{1}}{\Longrightarrow} \cdots \stackrel{\tilde{r}_{k}}{\Longrightarrow}\left(p, w x_{1} \cdots x_{k}\right) \Rightarrow \cdots \Rightarrow\left(p, w x_{1} \cdots x_{k} y_{\ell} \cdots y_{1}\right) \\
& \rho_{4}^{\prime}:\left(q, w x_{1} \cdots x_{k} y_{\ell} \cdots y_{1}\right) \stackrel{\tilde{s}_{1}}{\Longrightarrow} \cdots \stackrel{\tilde{s}_{\ell}}{\Rightarrow}\left(q, w x_{1} \cdots x_{k}\right) \Rightarrow \cdots \Rightarrow(q, w) .
\end{aligned}
$$

Moreover, since $\rho_{3}$ is a positive run from $(p, w)$ to $(q, w)$, we obtain $\rho_{3}^{\prime}$ from $\rho_{3}$ by replacing the prefix $w$ of every stack by $w x_{1} \cdots x_{k} y_{1} \cdots y_{\ell}$.

Then $\rho^{\prime}$ reads some word $u_{1} \tilde{r}_{1} \cdots \tilde{r}_{k} f u_{3} \tilde{s}_{1} \cdots \tilde{s}_{\ell} g u_{5}$ for $f, g \in \Sigma^{*}$. Note that since $r_{i}$ occurs in $\tilde{r}_{i}$ and $s_{i}$ occurs in $\tilde{s}_{j}$, we have $u=u_{1} u_{2} u_{3} u_{4} u_{5} \sqsubseteq u_{1} \tilde{r}_{1} \cdots \tilde{r}_{k} f u_{3} \tilde{s}_{1} \cdots \tilde{s}_{\ell} g u_{5}$.

We now show that every word in $\hat{\mathcal{A}}$ is accepted by a run with bounded stack height.
Lemma 6.5. $\mathrm{L}(\hat{\mathcal{A}})=\mathrm{L}_{h}(\hat{\mathcal{A}})$, where $h=2|Q|^{2}$.
Before we prove Lemma 6.5, we need another observation. Just like Lemma 6.4, one can show that for any $p, q \in Q$, we have $M_{p, q}(\mathcal{A}) \subseteq M_{p, q}(\hat{\mathcal{A}}) \subseteq \downarrow M_{p, q}(\mathcal{A})$ and in particular

$$
\begin{equation*}
\Delta_{p, q}(\mathcal{A})=\Delta_{p, q}(\hat{\mathcal{A}}) \quad \Delta_{q, p}(\overline{\mathcal{A}})=\Delta_{q, p}(\overline{\mathcal{A}}) \tag{6.2}
\end{equation*}
$$

where the second identity follows from the first: Duality yields $\Delta_{q, p}(\overline{\mathcal{A}})=\Delta_{q, p}(\hat{\overline{\mathcal{A}}})$ and since $\hat{\mathcal{A}}$ and $\overline{\hat{\mathcal{A}}}$ are isomorphic (i.e. they are the same up to a renaming of stack symbols), we have $\Delta_{q, p}(\hat{\mathcal{A}})=\Delta_{q, p}(\overline{\hat{\mathcal{A}}})$.

We now prove Lemma 6.5. Let $u \in \mathrm{~L}(\hat{\mathcal{A}})$. We show that any minimal length accepting run $\rho$ reading $u$ must have stack height $\leq h$ and hence $u \in \mathrm{~L}_{h}(\hat{\mathcal{A}})$.

Suppose the stack height of $\rho$ is larger than $h=2|Q|^{2}$.
Claim: $\rho$ decomposes into runs $\rho_{1}, \rho_{2}, \rho_{3}, \rho_{4}, \rho_{5}$ reading $u_{1}, u_{2}, u_{3}, u_{4}, u_{5}$, respectively, so that there are $p, q \in Q$ and $w \in \Gamma^{*}$ with

- $\rho_{2}$ is a positive run from $(p, w)$ to $(p, w v)$ of length $\geq 2$
- $\rho_{3}$ is a positive run from $(p, w v)$ to $(q, w v)$
- $\rho_{4}$ is a dually positive run from $(q, w v)$ to $(q, w)$

Let $c_{h}$ be a configuration along $\rho$ with stack height at least $2|Q|^{2}+1$. Then there exist $|Q|^{2}$ configurations $c_{2} \Rightarrow c_{4} \Rightarrow \cdots \Rightarrow c_{2|Q|^{2}} \Rightarrow c_{h}$ along $\rho$ such that $c_{2 i}$ is the last time that the stack height is $2 i$ before visiting $c_{h}$. Symmetrically, we have $c_{h} \Rightarrow c_{2|Q|^{2}}^{\prime} \Rightarrow \cdots c_{4}^{\prime} \Rightarrow c_{2}^{\prime}$ where $c_{2 i}^{\prime}$ is the first occurrence of stack height $2 i$ after visiting $c_{h}$. Clearly by definition the run between consecutive $c_{2 i}$ (resp. $c_{2 i}^{\prime}$ ) configurations is positive (resp. dually positive). Additionally, the length of the run between them must be at least two, because the stack heights differ by two. Considering the pair of states at each $c_{2 i}, c_{2 i}^{\prime}$, there are $|Q|^{2}$ possibilities. Hence there must exist indices $2 i<2 j$ such that the $c_{2 i}$ and $c_{2 j}$ have the same state $p$ and $c_{2 i}^{\prime}$ and $c_{2 j}^{\prime}$ have the same state $q$. It is now clear that $\rho_{2}: c_{2 i} \Rightarrow c_{2 j}, \rho_{3}: c_{2 j} \Rightarrow c_{2 j}^{\prime}$ and $\rho_{4}: c_{2 j}^{\prime} \Rightarrow c_{2 i}^{\prime}$ satisfy the conditions of the claim.

These conditions imply that $u_{2} \in M_{\underline{p}, q}(\hat{\mathcal{A}})$ and $u_{4} \in M_{q, p}(\overline{\hat{\mathcal{A}}})$. Therefore, we have $u_{2} \in \Delta_{p, q}(\hat{\mathcal{A}})^{*}=\Delta_{p, q}(\mathcal{A})^{*}$ and $u_{4} \in \Delta_{q, p}(\overline{\hat{\mathcal{A}}})^{*}=\Delta_{q, p}(\overline{\mathcal{A}})^{*}$ by Eq. (6.2).

We obtain the run $\rho^{\prime}$ from $\rho$ as follows. We replace $\rho_{2}$ by a single execution of the edge $p \xrightarrow{\Delta_{p, q}(\mathcal{A})^{*} \mid[p, q]} p$ reading $u_{2}$. Moreover, we replace $\rho_{4}$ by a single execution of the edge $q \xrightarrow{\Delta_{q, p}(\overline{\mathcal{A}})^{*} \mid[\overline{p, q]}} q$. Then $\rho^{\prime}$ is clearly a run reading $u=u_{1} u_{2} u_{3} u_{4} u_{5}$. Furthermore, since $\rho_{2}$
has length $\geq 2$, but the single edge used instead in $\rho^{\prime}$ only incurs length $1, \rho^{\prime}$ is strictly shorter than $\rho$. This is in contradiction to the minimal length of $\rho$. This completes the proof of Lemma 6.5 and thus also Theorem 6.2.
Remark. The augmented automaton $\hat{\mathcal{A}}$ yields a very simple construction of a finite automaton (of exponential size) for $\downarrow \mathrm{L}(\mathcal{A})$. First, it is easy to construct a finite automaton for $\mathrm{L}_{h}(\hat{\mathcal{A}})$. Then, by introducing $\varepsilon$-edges, we get a finite automaton for $\downarrow \mathrm{L}_{h}(\hat{\mathcal{A}})$, which, by Lemmas 6.4 and 6.5 , equals $\downarrow \mathrm{L}(\mathcal{A})$.

It is now straightforward to construct a polynomial size 1-bounded Petri net $N=$ $\left(\Sigma, S, T, \partial_{0}, \partial_{0}, \mathbf{m}_{0}, \mathbf{m}_{f}\right)$ with $\mathrm{L}(N)=\mathrm{L}_{h}(\hat{\mathcal{A}})$, thus proving Corollary 6.3. First, by adding states, we turn $\hat{\mathcal{A}}$ into a pushdown automaton $\mathcal{A}^{\prime}=\left(Q^{\prime}, \Sigma, \hat{\Gamma}, E^{\prime}, q_{0}, q_{f}\right)$, where every edge reads at most one letter, i.e. every edge $p \xrightarrow{R \mid v} q$ in $\mathcal{A}^{\prime}$ has $R=\{x\}$ for some $x \in \Sigma \cup\{\varepsilon\}$ (this is done by 'pasting' the automaton for $R$ in place of the edge). Moreover, we add $\varepsilon$-edges, so that for every edge $p \xrightarrow{\{x\} \mid v} q$, we also have an edge $p \xrightarrow{\{\varepsilon\} \mid v} q$. Then clearly $\mathrm{L}_{h}\left(\mathcal{A}^{\prime}\right)=\downarrow \mathrm{L}_{h}(\hat{\mathcal{A}})=\downarrow \mathrm{L}(\mathcal{A})$.

The net $N$ has a place $p$ for each state $p$ of $\mathcal{A}^{\prime}$ and for each $i \in\{1, \ldots, h\}$ and $\gamma \in \hat{\Gamma}$, it has a place $(i, \gamma)$. Moreover, for each $i \in\{0, \ldots, h\}$, it has a place $s_{i}$. Here, the idea is that a configuration $c=\left(p, \gamma_{1} \cdots \gamma_{n}\right)$ of $\mathcal{A}^{\prime}$ with $\gamma_{1}, \ldots, \gamma_{n} \in \hat{\Gamma}$ is represented as a marking $\mathbf{m}_{c}=\llbracket p,\left(1, \gamma_{1}\right), \ldots,\left(n, \gamma_{n}\right), s_{n} \rrbracket$. We call a marking of this form a stack marking and will argue that every reachable marking in $N$ is a stack marking. The transitions in $N$ correspond to edges in $\mathcal{A}^{\prime}$. For each edge $p \xrightarrow{\{x\} \mid v} q$ in $\hat{\mathcal{A}}$, we add the following transitions:

- if $v=\bar{\gamma}$ for some $\gamma \in \hat{\Gamma}$, then we have for every $1 \leq n \leq h$ a transition $t$ with $\partial_{0}(t)=\llbracket p,(n, \gamma), s_{n} \rrbracket, \partial_{1}(t)=\llbracket q, s_{n-1} \rrbracket$, and $\lambda(t)=x$.
- if $v \in \hat{\Gamma}$, then for every $0 \leq n<h$, we add a transition $t$ with $\partial_{0}(t)=\llbracket p, s_{n} \rrbracket, \partial_{1}(t)=$ $\llbracket q,(n+1, v), s_{n+1} \rrbracket$, and $\lambda(t)=x$.
- if $v=\varepsilon$, then we add a transition $t$ with $\partial_{0}(t)=\llbracket p \rrbracket, \partial_{1}(t)=\llbracket q \rrbracket$, and $\lambda(t)=x$.

Then clearly every reachable marking is a stack marking and we have $c \xrightarrow{x} c^{\prime}$ for configurations $c, c^{\prime}$ of $\mathcal{A}^{\prime}$ of stack height $\leq h$ if and only if $\mathbf{m}_{c} \xrightarrow{x} \mathbf{m}_{c^{\prime}}$. Therefore, if we set $\mathbf{m}_{0}=\llbracket q_{0}, s_{0} \rrbracket$ and $\mathbf{m}_{f}=\llbracket q_{f}, s_{0} \rrbracket$ as initial and final marking, we have $\mathrm{L}(N)=\mathrm{L}_{h}\left(\mathcal{A}^{\prime}\right)=\downarrow \mathrm{L}(\mathcal{A})$. This completes the proof of Corollary 6.3.

## References

[ABJ98] Parosh Aziz Abdulla, Ahmed Bouajjani, and Bengt Jonsson. On-the-fly analysis of systems with unbounded, lossy FIFO channels. In Proceedings of the 10th International Conference on Computer Aided Verification (CAV 1998), pages 305-318, 1998. doi:10.1007/BFb0028754.
[ABQ09] Mohamed Faouzi Atig, Ahmed Bouajjani, and Shaz Qadeer. Context-bounded analysis for concurrent programs with dynamic creation of threads. In Tools and Algorithms for the Construction and Analysis of Systems, 15th International Conference, TACAS 2009, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009, York, UK, March 22-29, 2009. Proceedings, volume 5505 of Lecture Notes in Computer Science, pages 107-123. Springer, 2009. doi:10.1007/978-3-642-00768-2\_11.
[Aho68] Alfred V. Aho. Indexed grammars - an extension of context-free grammars. J. ACM, 15(4):647-671, 1968. doi:10.1145/321479.321488.
[Ber79] Jean Berstel. Transductions and context-free languages. Teubner-Verlag, 1979.
[BLS15] Georg Bachmeier, Michael Luttenberger, and Maximilian Schlund. Finite automata for the suband superword closure of cfls: Descriptional and computational complexity. In Language and

Automata Theory and Applications - 9th International Conference, LATA 2015, Nice, France, March 2-6, 2015, Proceedings, volume 8977 of Lecture Notes in Computer Science, pages 473-485. Springer, 2015. doi:10.1007/978-3-319-15579-1 \_37.
[Cou91] Bruno Courcelle. On constructing obstruction sets of words. Bulletin of the EATCS, 44:178-186, 1991.
[CPSW16] Lorenzo Clemente, Paweł Parys, Sylvain Salvati, and Igor Walukiewicz. The diagonal problem for higher-order recursion schemes is decidable. In Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '16, New York, NY, USA, July 5-8, 2016, pages 96-105. ACM, 2016. doi:10.1145/2933575.2934527.
[CV07] Rohit Chadha and Mahesh Viswanathan. Decidability results for well-structured transition systems with auxiliary storage. In CONCUR 2007 - Concurrency Theory, 18th International Conference, CONCUR 2007, Lisbon, Portugal, September 3-8, 2007, Proceedings, volume 4703 of Lecture Notes in Computer Science, pages 136-150. Springer, 2007. doi:10.1007/978-3-540-74407-8\_10.
[Dam82] Werner Damm. The IO-and OI-hierarchies. Theoretical Computer Science, 20(2):95-207, 1982.
[DFS98] Catherine Dufourd, Alain Finkel, and Philippe Schnoebelen. Reset nets between decidability and undecidability. In Automata, Languages and Programming, 25th International Colloquium, ICALP'98, Aalborg, Denmark, July 13-17, 1998, Proceedings, volume 1443 of Lecture Notes in Computer Science, pages 103-115. Springer, 1998. doi:10.1007/BFb0055044.
[DG86] Werner Damm and Andreas Goerdt. An automata-theoretical characterization of the OI-hierarchy. Information and Control, 71(1):1-32, 1986.
[GM12] Pierre Ganty and Rupak Majumdar. Algorithmic verification of asynchronous programs. ACM Transactions on Programming Languages and Systems (TOPLAS), 34(1):6, 2012. doi:10.1145/ 2160910. 2160915.
[GRB07] Gilles Geeraerts, Jean-François Raskin, and Laurent Van Begin. Well-structured languages. Acta Informatica, 44(3-4):249-288, 2007. doi:10.1007/s00236-007-0050-3.
[Gre78] Sheila A. Greibach. Remarks on blind and partially blind one-way multicounter machines. Theoretical Computer Science, 7(3):311-324, 1978. doi:10.1016/0304-3975(78)90020-8.
[Gre16] Charles Grellois. Semantics of linear logic and higher-order model-checking. PhD thesis, Univeristé Denis Diderot Paris 7, 2016.
[Hai69] Leonard H Haines. On free monoids partially ordered by embedding. Journal of Combinatorial Theory, 6(1):94-98, 1969.
[HKO16] Matthew Hague, Jonathan Kochems, and C.-H. Luke Ong. Unboundedness and downward closures of higher-order pushdown automata. In Proceedings of the 43 rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20-22, 2016, pages 151-163. ACM, 2016. doi:10.1145/2837614.2837627.
[HMOS08] Matthew Hague, Andrzej S. Murawski, C.-H. Luke Ong, and Olivier Serre. Collapsible pushdown automata and recursion schemes. In Proceedings of the Twenty-Third Annual IEEE Symposium on Logic in Computer Science, LICS 2008, 24-27 June 2008, Pittsburgh, PA, USA, pages 452-461, 2008. doi:10.1109/LICS.2008.34.
[HMU07] John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. Introduction to automata theory, languages, and computation, 3rd Edition. Pearson international edition. Addison-Wesley, 2007.
[Jan79] Matthias Jantzen. On the hierarchy of Petri net languages. RAIRO - Theoretical Informatics and Applications - Informatique Théorique et Applications, 13(1):19-30, 1979. URL: http://www. numdam.org/item?id=ITA_1979__13_1_19_0.
[JM07] Ranjit Jhala and Rupak Majumdar. Interprocedural analysis of asynchronous programs. In Proceedings of the 34th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2007, Nice, France, January 17-19, 2007, pages 339-350. ACM, 2007. doi: 10.1145/1190216.1190266.
[KO11] Naoki Kobayashi and C.-H. Luke Ong. Complexity of model checking recursion schemes for fragments of the modal mu-calculus. Logical Methods in Computer Science, 7(4), 2011.
[Kob09] Naoki Kobayashi. Types and higher-order recursion schemes for verification of higher-order programs. In Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009, Savannah, GA, USA, January 21-23, 2009, pages 416-428, 2009. doi:10.1145/1480881.1480933.
[Kob19] Naoki Kobayashi. Inclusion between the frontier language of a non-deterministic recursive program scheme and the Dyck language is undecidable. Theoretical Computer Science, 777:409-416, 2019.
[Mas74] AN Maslov. The hierarchy of indexed languages of an arbitrary level. Doklady Akademii Nauk, 217(5):1013-1016, 1974.
[May03] Richard Mayr. Undecidable problems in unreliable computations. Theoretical Computer Science, 297(1-3):337-354, 2003
[Ong15] Luke Ong. Higher-order model checking: An overview. In 30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2015, Kyoto, Japan, July 6-10, 2015, pages 1-15, 2015. doi:10.1109/LICS.2015.9.
[Par17] Paweł Parys. The complexity of the diagonal problem for recursion schemes. In 37th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2017, December 11-15, 2017, Kanpur, India, volume 93 of LIPIcs, pages 45:1-45:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPIcs.FSTTCS.2017.45.
[Sip97] Michael Sipser. Introduction to the theory of computation. PWS Publishing Company, 1997
[SV06] Koushik Sen and Mahesh Viswanathan. Model checking multithreaded programs with asynchronous atomic methods. In Computer Aided Verification, 18th International Conference, CAV 2006, Seattle, WA, USA, August 17-20, 2006, Proceedings, volume 4144 of Lecture Notes in Computer Science, pages 300-314. Springer, 2006. doi:10.1007/11817963\_29.
[TZ19] Ramanathan S. Thinniyam and Georg Zetzsche. Regular separability and intersection emptiness are independent problems. In Proceedings of FSTTCS 2019, volume 150 of Leibniz International Proceedings in Informatics (LIPIcs), pages 51:1-51:15, Dagstuhl, Germany, 2019. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.FSTTCS.2019.51.
[vL78] Jan van Leeuwen. Effective constructions in well-partially-ordered free monoids. Discrete Mathematics, 21(3):237-252, 1978. doi:10.1016/0012-365X (78) 90156-5.
[Zet15] Georg Zetzsche. An approach to computing downward closures. In ICALP 2015, volume 9135, pages $440-451$. Springer, 2015. The undecidability of $Z$ intersection is shown in the full version: http://arxiv.org/abs/1503.01068.

## Appendix A. Compiling away internal actions

We have commented in the examples of Section 2 that internal actions and internal updates of the global state are useful in modeling asynchronous programs from their programming language syntax. Indeed, we note that the definition of asynchronous programs in [GM12] additionally uses a separate alphabet of internal actions, in addition to the alphabet of handler posts. We show how a model of asynchronous programs with internal actions can be reduced to our, simpler, model.

Let $\mathcal{C}$ be a language class over an alphabet $\Sigma$ of handler names. The definition of asynchronous programs with internal actions, as used by [SV06, JM07, GM12], is as follows. ${ }^{3}$ An asynchronous program over $\mathcal{C}$ with internal actions (aka AP over $\mathcal{C}$ with internal actions) is a tuple $\mathfrak{P}=\left(D, \Sigma, \Sigma_{\mathbf{i}}, \mathscr{L}, R, d_{0}, \mathbf{m}_{0}\right)$, where $D, \Sigma, d_{0}, \mathbf{m}_{0}$ are as in Definition 2.2, $\Sigma_{\mathbf{i}}$ is an alphabet of internal actions disjoint from $\Sigma$, the set $\mathscr{L}=\left(L_{\sigma}\right)_{\sigma \in \Sigma}$ consists of languages from $\mathcal{C}$ (one for each handler $\sigma \in \Sigma$ ), and $R=\left(D, \Sigma \cup \Sigma_{\mathrm{i}}, \delta\right)$ a deterministic finite state automaton where $D$ is the set of states, $\Sigma \cup \Sigma_{i}$ the alphabet and $\delta$ the transition relation specifying the effect of each internal action on the global state $D$. We will write $d \underset{R}{w}{ }^{*} d^{\prime}$ to mean that there is a sequence of transitions with labels $w_{1} w_{2} \ldots w_{n}=w$ in the automaton $R$ using which we can reach $d^{\prime}$ from $d$.

For alphabets $\Sigma, \Sigma^{\prime}$ with $\Sigma \subseteq \Sigma^{\prime}$, the Parikh map Parikh: $\Sigma^{* *} \rightarrow \mathbb{M}[\Sigma]$ maps a word $w \in \Sigma^{\prime *}$ to a multiset Parikh $(w)$ such that $\operatorname{Parikh}(w)(a)$ is the number of occurrences of $a$ in $w$

[^4]for each $a \in \Sigma$. For example, $\operatorname{Parikh}(a b b a b)(a)=2, \operatorname{Parikh}(a b b a b)(b)=3$ and $\operatorname{Parikh}(\varepsilon)=\llbracket \rrbracket$. For a language $L$, we define $\operatorname{Parikh}(L)=\{\operatorname{Parikh}(w) \mid w \in L\}$. If the alphabet $\Sigma$ is not clear from the context, we write $\mathrm{Parikh}_{\Sigma}$ (usually $\Sigma^{\prime}=\Sigma$ ).

The semantics of a $\mathfrak{P}$ is given as a labeled transition system over the set of configurations, with a transition relation $\rightarrow \subseteq(D \times \mathbb{M}[\Sigma]) \times \Sigma \times(D \times \mathbb{M}[\Sigma])$ defined as follows: let $\mathbf{m}, \mathbf{m}^{\prime} \in$ $\mathbb{M}[\Sigma], d, d^{\prime} \in D$ and $\sigma \in \Sigma$

$$
(d, \mathbf{m} \oplus \llbracket \sigma \rrbracket) \underset{\mathfrak{P}}{\underset{\mathfrak{P}}{\prime}}\left(d^{\prime}, \mathbf{m} \oplus \mathbf{m}^{\prime}\right)
$$

iff

$$
\exists w \in\left(\Sigma \cup \Sigma_{\mathrm{i}}\right)^{*}: d \underset{R}{\stackrel{w}{\Rightarrow}} d^{\prime} \wedge w \in L_{\sigma} \wedge \mathbf{m}^{\prime}=\operatorname{Parikh}_{\Sigma}(w) .
$$

We now show that internal actions can be compiled away. Thus, for the algorithms presented in this paper, we use the more abstract, language-theoretic version of Definition 2.2, while we use internal actions as syntactic sugar in examples.

Lemma A.1. Let $\mathcal{C}$ be a full trio. Given an $\mathrm{AP} \mathfrak{P}_{\mathfrak{i}}$ with internal actions over $\mathcal{C}$, one can construct an $\mathrm{AP} \mathfrak{P}$ over $\mathcal{C}$ such that both have identical sets of runs.
Proof. The proof is along similar lines to that of Lemmas 4.3, 4.5 in [GM12]. Given $\mathfrak{P}_{\mathbf{i}}=\left(D, \Sigma, \Sigma_{\mathbf{i}}, \mathscr{L}, R, d_{0}, \mathbf{m}_{0}\right)$ we construct $\mathfrak{P}=\left(D, \Sigma,\left(L_{c}\right)_{c \in \mathfrak{C}}, d_{0}, \mathbf{m}_{0}\right)$ such that

$$
(d, \mathbf{m} \oplus \llbracket \sigma \rrbracket) \underset{\mathfrak{P}}{\stackrel{\sigma}{\vec{~}}}\left(d^{\prime}, \mathbf{m} \oplus \mathbf{m}^{\prime}\right) \quad \text { iff } \quad(d, \mathbf{m} \oplus \llbracket \sigma \rrbracket) \underset{\mathfrak{P}_{\mathrm{i}}}{\stackrel{\sigma}{\rightarrow}}\left(d^{\prime}, \mathbf{m} \oplus \mathbf{m}^{\prime}\right)
$$

Let $L\left(R_{d, d^{\prime}}\right)$ be the language of the automaton $R$ when $d$ is the initial state and $d^{\prime}$ is the accepting state. We define $L_{d \sigma d^{\prime}}$ as:

$$
L_{d \sigma d^{\prime}}:=\operatorname{Proj}_{\Sigma}\left(L_{\sigma} \cap L\left(R_{d, d^{\prime}}\right)\right)
$$

First observe that the projection operation is a homomorphism and $L\left(R_{d, d^{\prime}}\right)$ ) is a regular language; hence by virtue of $\mathcal{C}$ being a full trio $L_{d \sigma d^{\prime}}$ as defined above is in $\mathcal{C}$. The conditions $\exists w \in\left(\Sigma \cup \Sigma_{\mathbf{i}}\right)^{*}: d \underset{R}{\stackrel{w}{*}} d^{\prime} \wedge w \in L_{\sigma} \wedge \mathbf{m}^{\prime}=\operatorname{Parikh}_{\Sigma}(w)$ and $\exists w \in L_{d \sigma d^{\prime}}: \operatorname{Parikh}(w)=\mathbf{m}^{\prime}$ are seen to be equivalent from the definition of $L_{d \sigma d^{\prime}}$, concluding the proof of the lemma.

## Appendix B. Proof of Proposition 3.3

We prove the following proposition.
Let $\mathfrak{P}=\left(D, \Sigma,\left(L_{c}\right)_{c \in \mathfrak{C}}, d_{0}, \mathbf{m}_{0}\right)$ be an asynchronous program. Then $\downarrow \operatorname{Runs}(\downarrow \mathfrak{P})=$ $\downarrow$ Runs $(\mathfrak{P})$. In particular,
(1) For every $d \in D, \mathfrak{P}$ can reach $d$ if and only if $\downarrow \mathfrak{P}$ can reach $d$.
(2) $\mathfrak{P}$ is terminating if and only if $\downarrow \mathfrak{P}$ is terminating.
(3) $\mathfrak{P}$ is bounded if and only if $\downarrow \mathfrak{P}$ is bounded.

Proof. A run of the asynchronous program $\mathfrak{P}$ is defined as a sequence $c_{0}, \sigma_{1}, c_{1}, \sigma_{2}, \ldots$ containing alternating elements of configurations $c_{i}$ and letters $\sigma_{i}$ beginning with the configuration $c_{0}=\left(d_{0}, \mathbf{m}_{0}\right)$. First we observe that

$$
\begin{equation*}
\operatorname{Runs}(\mathfrak{P}) \subseteq \operatorname{Runs}(\downarrow \mathfrak{P}) \tag{B.1}
\end{equation*}
$$

because every transition enabled in $\mathfrak{P}$ is also enabled in $\downarrow \mathfrak{P}$. Next, we claim:

$$
\begin{equation*}
\forall \rho \in \operatorname{Runs}(\downarrow \mathfrak{P}) \exists \rho^{\prime} \in \operatorname{Runs}(\mathfrak{P}) \rho \unlhd \rho^{\prime} \tag{B.2}
\end{equation*}
$$

Let $\left.\rho\right|_{k}=\left(d_{0}, \mathbf{m}_{0}\right), \sigma_{1},\left(d_{1}, \mathbf{m}_{1}\right), \sigma_{2}, \ldots, \sigma_{k},\left(d_{k}, \mathbf{m}_{k}\right)$ be the $2 k+1$ length prefix of $\rho$. We show that for each $\left.\rho\right|_{k}$ there exists $\rho_{k}^{\prime} \in \operatorname{Runs}(\mathfrak{P})$ such that $\left.\rho\right|_{k} \unlhd \rho_{k}^{\prime}$ and in addition, $\forall k \forall i \leq k \rho_{k}^{\prime}(i)=\rho_{k+1}^{\prime}(i)$. We can then define $\rho^{\prime}(i):=\rho_{i}^{\prime}(i)$ and clearly $\rho \unlhd \rho^{\prime}$.

We prove by induction on $k$.
Base Case: $\left.\rho\right|_{0}=\left.\rho^{\prime}\right|_{0}=\left(d_{0}, \mathbf{m}_{0}\right)$.
Induction Step: Let $\left.\rho\right|_{k}=\left(d_{0}, \mathbf{m}_{0}\right), \sigma_{1},\left(d_{1}, \mathbf{m}_{1}\right), \sigma_{2}, \ldots,\left(d_{k}, \mathbf{m}_{k}\right) \in \operatorname{Runs}(\downarrow \mathfrak{P})$. By induction hypothesis there is $\rho_{k-1}^{\prime}=\left(d_{0}, \mathbf{m}_{0}\right), \sigma_{1},\left(d_{1}, \mathbf{m}_{1}^{\prime}\right), \sigma_{2}, \ldots,\left(d_{k-1}, \mathbf{m}_{k-1}^{\prime}\right) \in \operatorname{Runs}(\mathfrak{P})$ such that $\rho_{k-1} \unlhd \rho_{k-1}^{\prime}$.

$$
\begin{gathered}
\left(d_{k-1}, \mathbf{m}_{k-1}\right) \stackrel{\sigma_{k}}{\downarrow \mathfrak{P}}\left(d_{k}, \mathbf{m}_{k}\right) \\
\Rightarrow \exists \mathbf{m}_{k-1}^{\prime \prime}: \mathbf{m}_{k-1}=\mathbf{m}_{k-1}^{\prime \prime} \oplus \llbracket \sigma_{k} \rrbracket \wedge\left(d_{k-1}, \mathbf{m}_{k-1}^{\prime \prime} \oplus \llbracket \sigma_{k} \rrbracket\right) \stackrel{\sigma_{k}}{\downarrow \mathfrak{P}}\left(d_{k}, \mathbf{m}_{k}\right) \\
\Rightarrow \exists w \in \Sigma^{*}: w \in \downarrow L_{d_{k-1} \sigma_{k} d_{k}} \wedge \operatorname{Parikh}(w) \oplus \mathbf{m}_{k-1}^{\prime \prime}=\mathbf{m}_{k} \\
\Rightarrow \exists w^{\prime} \in \Sigma^{*}: w \sqsubseteq w^{\prime} \wedge w^{\prime} \in L_{d_{k-1} \sigma_{k} d_{k}} \\
\Rightarrow\left(d_{k-1}, \mathbf{m}_{k-1}\right) \frac{\sigma_{\mathfrak{F}}}{\underset{\mathfrak{P}}{ }}\left(d_{k}, \mathbf{m}_{k} \oplus \mathbf{m}_{\Delta}\right) \text { where } \mathbf{m}_{\Delta} \oplus \operatorname{Parikh}(w)=\operatorname{Parikh}\left(w^{\prime}\right) \\
\Rightarrow\left(d_{k-1}, \mathbf{m}_{k-1}^{\prime}\right) \frac{\sigma_{k}}{\mathfrak{P}}\left(d_{k}, \mathbf{m}_{k}^{\prime}\right) \text { where } \mathbf{m}_{k}^{\prime}=\mathbf{m}_{k} \oplus \mathbf{m}_{\Delta} \oplus\left(\mathbf{m}_{k-1}^{\prime} \ominus \mathbf{m}_{k-1}\right)
\end{gathered}
$$

Defining $\rho_{k}^{\prime}:=\rho_{k-1}^{\prime}, \sigma_{k},\left(d_{k}, \mathbf{m}_{k}^{\prime}\right)$ we see that $\left.\rho\right|_{k} \unlhd \rho_{k}^{\prime}$, completing the proof of Equation B.2. We are now ready to show that $\downarrow \operatorname{Runs}(\downarrow \mathfrak{P})=\downarrow \operatorname{Runs}(\mathfrak{P})$. The direction $\downarrow \operatorname{Runs}(\mathfrak{P}) \subseteq$ $\downarrow$ Runs( $\downarrow \mathfrak{P}$ ) follows immediately from Equation (B.1). Conversely, let

$$
\begin{aligned}
& \rho=\left(s_{0}, \mathbf{n}_{0}\right), \sigma_{1},\left(s_{1}, \mathbf{n}_{1}\right), \sigma_{2}, \ldots \in \downarrow \operatorname{Runs}(\downarrow \mathfrak{P}) \\
\Rightarrow & \exists \rho^{\prime} \in \operatorname{Runs}(\downarrow \mathfrak{P}) \rho \unlhd \rho^{\prime} \\
\Rightarrow & \exists \rho^{\prime \prime} \in \operatorname{Runs}(\mathfrak{P}) \rho^{\prime} \unlhd \rho^{\prime \prime} \\
\Rightarrow & \rho \in \downarrow \operatorname{Runs}(\mathfrak{P})
\end{aligned}
$$

$$
\Rightarrow \exists \rho^{\prime \prime} \in \operatorname{Runs}(\mathfrak{P}) \rho^{\prime} \unlhd \rho^{\prime \prime} \quad \text { by Equation (B.2) }
$$

We have proved that $\downarrow \operatorname{Runs}(\downarrow \mathfrak{P})=\downarrow \operatorname{Runs}(\mathfrak{P})$. We now show that each of the three properties i.e. safety, termination and boundedness only depend on the downclosure of the runs.
$\underline{\text { Safety: }}$

$$
d \text { is reachable in } \mathfrak{P}
$$

$$
\text { iff } \exists \rho=\left(d_{0}, \mathbf{m}_{0}\right), \sigma_{1},\left(d_{1}, \mathbf{m}_{1}\right), \sigma_{2}, \ldots, \sigma_{k},\left(d_{k}, \mathbf{m}_{k}\right) \in \operatorname{Runs}(\mathfrak{P}): d_{k}=d
$$

By Equation B.1, we know $\rho \in \operatorname{Runs}(\mathfrak{P})$ implies $\rho \in \downarrow \operatorname{Runs}(\mathfrak{P})$. Conversely, if there is $\rho^{\prime}=\left(s_{0}, \mathbf{n}_{0}\right), \sigma_{1},\left(s_{1}, \mathbf{n}_{1}\right), \sigma_{2}, \ldots, \sigma_{k},\left(s_{k}, \mathbf{n}_{k}\right) \in \downarrow \operatorname{Runs}(\mathfrak{P})$ with $s_{k}=d$, then by Equation B.2, there is $\rho=\left(d_{0}, \mathbf{m}_{0}\right), \sigma_{1},\left(d_{1}, \mathbf{m}_{1}\right), \sigma_{1}, \ldots, \sigma_{k},\left(d_{k}, \mathbf{m}_{k}\right) \in \operatorname{Runs}(\mathfrak{P})$ with $\rho^{\prime} \unlhd \rho$ which implies $d_{k}=d$. Hence we have
$d$ is reachable in $\mathfrak{P}$
iff $\exists \rho=\left(s_{0}, \mathbf{n}_{0}\right), \sigma_{1},\left(s_{1}, \mathbf{n}_{1}\right), \sigma_{2}, \ldots, \sigma_{k},\left(s_{k}, \mathbf{n}_{k}\right) \in \downarrow \operatorname{Runs}(\mathfrak{P}): s_{k}=d$

By a similar argument as above we also have:
Termination:
$\mathfrak{P}$ does not terminate
iff $\exists \rho \in \operatorname{Runs}(\mathfrak{P}): \rho$ is infinite
iff $\exists \rho \in \downarrow \operatorname{Runs}(\mathfrak{P}): \rho$ is infinite

Boundedness:

$$
\begin{aligned}
& \mathfrak{P} \text { is bounded } \\
& \text { iff } \exists N \in \mathbb{N} \forall \rho \in \operatorname{Runs}(\mathfrak{P}) \forall i|\rho(2 i) . m|<N \\
& \text { iff } \exists N \in \mathbb{N} \forall \rho \in \downarrow \operatorname{Runs}(\mathfrak{P}) \forall i|\rho(2 i) . m|<N
\end{aligned}
$$

In each of the three cases, the property only depends on the downclosure and hence one may equivalently replace $\mathfrak{P}$ by $\downarrow \mathfrak{P}$ since $\downarrow \operatorname{Runs}(\downarrow \mathfrak{P})=\downarrow \operatorname{Runs}(\mathfrak{P})$.

## Appendix C. Proof of Proposition 5.1

We begin with a simple observation. For every $\operatorname{HORS} \mathscr{S}=(\Sigma, \mathcal{N}, S, \mathcal{R})$, there exists another HORS $\mathscr{S}^{\prime}=\left(\Sigma^{\prime}, \mathcal{N}^{\prime}, S, \mathcal{R}^{\prime}\right)$ where $\Sigma^{\prime}=\{\mathrm{br}, \mathrm{e}\} \cup \tilde{\Sigma}$ with br of arity 2 , e of arity 0 and all symbols in $\tilde{\Sigma}$ of arity 1 ; such that $\mathcal{L}_{\mathrm{p}}(\mathscr{S})=\mathcal{L}_{\mathrm{p}}\left(\mathscr{S}^{\prime}\right)$.

By rewriting every terminal symbol $A$ of arity $n \geq 2$ using the rule

$$
A x_{1} x_{2} \cdots x_{n} \rightarrow \operatorname{br}\left(x _ { 1 } \operatorname { b r } \left(x_{2} \operatorname{br}\left(\cdots \operatorname{br}\left(x_{n-1} x_{n}\right)\right)\right.\right.
$$

and by replacing the occurrence of every nullary symbol in $\mathscr{S}$ by e, we get $\mathscr{S}^{\prime}$, which has the same path language.
We assume due to the above observation that $\Sigma=\{\mathrm{br}, \mathrm{e}\} \cup \tilde{\Sigma}$ where br is binary, e is nullary and all letters in $\tilde{\Sigma}$ are unary. Define $\mathscr{S}^{\prime}=\left(\Sigma^{\prime}, \mathcal{N}^{\prime}, S^{\prime}, \mathcal{R}^{\prime}\right)$ as $\Sigma^{\prime}=\tilde{\Sigma} \cup\{\mathrm{e}\}, \mathcal{N}^{\prime}=\mathcal{N} \cup\{B$ : $\circ \rightarrow(\circ \rightarrow 0)\}, \mathcal{R}^{\prime}=\{r[\mathrm{br} / B] \mid r \in \mathcal{R}\} \cup\{B x y \rightarrow x, B x y \rightarrow y\}$, where by $r[\mathrm{br} / B]$ we mean the rule $r$ with br uniformly replaced by $B$. Note that the only new non-terminal symbol introduced is $B$, which is of order 1. Hence the obtained word scheme $\mathscr{S}^{\prime}$ is of the same order as $\mathscr{S}$.
$\mathcal{L}_{\mathrm{p}}(\mathscr{S}) \subseteq \mathcal{L}_{\mathrm{w}}\left(\mathscr{S}^{\prime}\right)$ : Let $w \in \mathcal{L}_{\mathrm{p}}(\mathscr{S})$. Therefore there exists a finite path $\Pi$ in a sentential form $t$ such that $\forall \bar{n} \in \operatorname{dom}(\Pi) w_{\Pi}(\bar{n}) \in \Sigma$. We derive $t^{\prime}:=t[\mathrm{br} / B]$ using the corresponding rules in $\mathcal{R}^{\prime}$. Note that the corresponding path $\Pi^{\prime}$ in $t^{\prime}$ satisfies $\forall \bar{n} \in \operatorname{dom}\left(\Pi^{\prime}\right) w_{\Pi^{\prime}}(\bar{n}) \in \Sigma \cup\{B\}$. We then apply either $B x y \rightarrow x$ or $B x y \rightarrow y$ to each $B$ in $t^{\prime}$ according to the path $\Pi^{\prime}$ to obtain $w$ in the word scheme.
$\mathcal{L}_{\mathrm{w}}\left(\mathscr{S}^{\prime}\right) \subseteq \mathcal{L}_{\mathrm{p}}(\mathscr{S})$ : We define the order $\leq_{\text {pre }}$ on sequences of natural numbers $\bar{n}, \bar{m}$ as $\bar{n} \leq_{\text {pre }} \bar{m}$ if $\bar{m}=\bar{n} \bar{k}$ for some sequence $\bar{k}$.

Suppose that for given a sentential form $t^{\prime}$ of $\mathscr{S}^{\prime}$ there exists a sentential form $t$ of $\mathscr{S}$ and a map $\alpha: \operatorname{dom}\left(t^{\prime}\right) \rightarrow \operatorname{dom}(t)$ (simply called embedding henceforth) satisfying the following conditions:

- $\forall \bar{n} \in \operatorname{dom}\left(t^{\prime}\right)$

$$
t^{\prime}(\bar{n})=\left\{\begin{array}{l}
B \text { if } t(\alpha(\bar{n}))=\mathrm{br}  \tag{C.1}\\
t(\alpha(\bar{n})) \text { otherwise }
\end{array}\right.
$$

- $\forall \bar{n}, \bar{m} \in \operatorname{dom}\left(t^{\prime}\right), \bar{n} \leq_{\text {pre }} \bar{m}$ implies

$$
\begin{equation*}
\alpha(\bar{n}) \leq_{\text {pre }} \alpha(\bar{m}), \text { and } \tag{C.2}
\end{equation*}
$$

$$
\begin{align*}
& \left(\forall \bar{l}\left(\bar{n}<_{\text {pre }} \bar{l}<_{\text {pre }} \bar{m}\right) \text { implies } t^{\prime}(\bar{l}) \notin \tilde{\Sigma}\right) \text { implies } \\
& \left(\forall \bar{k} \alpha(\bar{n})<_{\text {pre }} \bar{k}<_{\text {pre }} \alpha(\bar{m}) \text { implies } t(\bar{k}) \notin \tilde{\Sigma}\right) \tag{C.3}
\end{align*}
$$

Informally, Equations C.1, C. 2 and C. 3 state the following: $\alpha$ preserves labels, except for the case of br where it maps to $B, \alpha$ preserves the order $\leq_{\text {pre }}$ and the images of any two nodes with node labels from $\tilde{\Sigma}$ with no node label from $\tilde{\Sigma}$ in between are mapped to two such nodes with the same property.

We will show by induction on the length of the derivation that such a pair $(t, \alpha)$ always exists given some $t^{\prime}$. Let us see how the existence of such $t$ and $\alpha$ gives us the proposition. Consider a word $w=w_{1} w_{2} \ldots w_{n} \in \mathcal{L}_{\mathrm{w}}\left(\mathscr{S}^{\prime}\right)$. In other words, there is a term $t^{\prime}=w_{1}\left(w_{2} \ldots\left(w_{n}(\mathrm{e})\right) \ldots\right)$ such that ${S^{\prime} \overrightarrow{\mathscr{S}}^{\prime}}^{*} t^{\prime}$. Corresponding to this, we have a sentential form $S \overrightarrow{\mathscr{S}}^{*} t$ and $\alpha$ satisfying the given conditions. In particular, there exists a path $\Pi$ with $\operatorname{dom}(\Pi) \subseteq \operatorname{dom}(t)$ which is the path in $t$ connecting the $\alpha(\varepsilon)$ to $\alpha\left(0^{n}\right)$. It is immediate that $\operatorname{Proj}_{\tilde{\Sigma}}(\Pi)=w$. It remains to show the existence of $t$ and $\alpha$ by induction on the length of derivations.

Base case: This is trivial since $t=t^{\prime}=S$.
Induction step: Suppose $t_{0}^{\prime} \underset{r \in \mathcal{R}^{\prime}}{\vec{\prime}} t^{\prime}$ where $t_{0}^{\prime}$ is a sentential form. By induction hypothesis, there is a sentential form $t_{0}$ of $\mathscr{S}$ and $t_{0}^{\prime}$ embeds into $t_{0}$ via map $\alpha_{0}$. Assume that the rule $r$ is applied at position $\bar{n}$ on $t_{0}^{\prime}$. We now have two cases to consider:
 hypothesis, we have $t_{0}$ and and an embedding $\alpha_{0}$ of $t_{0}^{\prime}$ into $t_{0}$. Refering to Figure 2 , we see that $t$ can be taken to be $t_{0}$ and $\alpha$ maps all nodes in the subtree $u_{0}^{\prime}$ into $u_{0}$ as before, while the subtree $l_{0}^{\prime}$ rooted at $\bar{n}$ is mapped into $l_{0}$. It is immediate that $\alpha$ preserves the order and since by induction hypothesis $\alpha_{0}(\bar{n})$ has label br, Equation C. 3 is also satisfied by $\alpha$ since no new $\tilde{\Sigma}$ labelled nodes have been added.
Case 2: We demonstrate for the case when a non-terminal of arity two is replaced for an easier reading of the proof: the rule $r$ is $N x y \underset{\boldsymbol{\mathscr { S }}^{\prime}}{ } t_{1}$ for some nonterminal $N \neq B$. Refering to Figure 3 , the rule replaces the subtree rooted at $\bar{n}$ in $t_{0}^{\prime}$ with $t_{1}\left[x / l_{0}^{\prime}, y / r_{0}^{\prime}\right]$ where $l_{0}^{\prime}, r_{0}^{\prime}$ are respectively the left and right subtrees of the subtree rooted at $\bar{n}$ in $u_{0}^{\prime}$. For nodes in the subtree $u_{0}^{\prime}$ of $t^{\prime}, \alpha$ mimics $\alpha_{0}$. By induction we know that the subtrees $l_{0}^{\prime}, r_{0}^{\prime}$ of $t_{0}^{\prime}$ embed respectively into $l_{0}, r_{0}$ of $t_{0}$. Thus $\alpha$ maps every subtree $l_{0}^{\prime}$ (resp. $r_{0}^{\prime}$ ) rooted at $\bar{m}$ (resp. $\bar{k}$ ) in $t_{1}\left[x / l_{0}^{\prime}, y / r_{0}^{\prime}\right]$ into the corresponding subtree $l_{0}\left(\right.$ resp. $\left.r_{0}\right)$ rooted at $\alpha(\bar{m})$ (resp. $\alpha(\bar{k})$ ). Note that $\alpha(\bar{m})$ can be defined as the root of the $i$-th occurrence of $l_{0}$ from left to right if $\bar{m}$ is the $i$-th occurrence of $l_{0}^{\prime}$ from left to right. Nodes in $t_{1}\left[x / l_{0}^{\prime}, y / r_{0}^{\prime}\right]$ which are not in any of the $l_{0}^{\prime}$ (or $r_{0}^{\prime}$ ) subtrees (i.e.) between $\bar{n}$ and $\bar{m}($ or $\bar{k})$ ) have corresponding nodes in $t_{1}\left[x / l_{0}, y / r_{0}\right]$ to which they can be mapped. Label preservation and order preservation immediately follow by appeal to the induction hypothesis. In order to see that Equation C. 3 holds, consider consecutive $\tilde{\Sigma}$ nodes $\bar{n}^{\prime}, \bar{m}^{\prime}$ in $t^{\prime}$ (i.e. $\bar{n}^{\prime}<_{p r e} \bar{m}^{\prime}$ and there are no $\tilde{\Sigma}$ labels in between). If both $\bar{n}^{\prime}$ and $\bar{m}^{\prime}$ are in $u_{0}^{\prime}$ or in one of the $l_{0}^{\prime}$ (or $r_{0}^{\prime}$ ) then the induction hypothesis applies. In the case where $\bar{n}^{\prime} \in u_{0}^{\prime}$ and $\bar{m}^{\prime} \in l_{0}^{\prime}$ (resp. $r_{0}^{\prime}$,) this means that no $\tilde{\Sigma}$ labels are present


Figure 2: Construction of embedding $\alpha$ from $\alpha_{0}$ in the case of the rule $B x y \rightarrow x$
in the path from $\bar{n}$ to $\bar{m}$ (resp. $\bar{k}$ ), $\bar{n}^{\prime}$ to $\bar{n}$ nor $\bar{m}$ (resp. $\bar{k}$ ) to $\bar{m}^{\prime}$. The path from $\alpha(\bar{n})$ to $\alpha(\bar{m})$ is identical to that from $\bar{n}$ to $\bar{m}$ (resp. $\bar{k}$ ). The induction hypothesis also implies that the first a label for some $\mathrm{s} \in \tilde{\Sigma}$ in $l_{0}^{\prime}$ must be mapped to first a label in $l_{0}$ (or $u_{0}^{\prime}$ does not contain any $\tilde{\Sigma}$ labels). Hence there are no $\tilde{\Sigma}$ labels between $\alpha(\bar{m})($ resp. $\alpha(\bar{k}))$ and $\alpha\left(\bar{m}^{\prime}\right)$ and similarly between $\alpha\left(\bar{n}^{\prime}\right)$ and $\alpha(\bar{n})$. The final case is when either $\bar{n}^{\prime}$ or $\bar{m}^{\prime}$ lies between $\alpha(\bar{n})$ and $\alpha(\bar{m})$ (resp. $\alpha(\bar{k})$ ). There are subcases here to consider when the other point lies in $u_{0}^{\prime}, l_{0}$ or also between $\bar{n}$ and $\bar{m}$. In all of these subcases, it easily follows that there are no extra $\tilde{\Sigma}$ labels introduced in between two consecutive nodes which are in the image of $\alpha$.


Figure 3: Construction of embedding $\alpha$ from $\alpha_{0}$ when the rule is $N x y \rightarrow t_{1}$


[^0]:    Key words and phrases: Higher-order asynchronous programs, Higher-order recursion schemes and Decidability.

    An abridged version of this paper appeared in TACAS 2021; missing proofs have been added to create this version.

[^1]:    © R. Majumdar, R. S. Thinniyam, and G. Zetzsche
    (c) Creative Commons

[^2]:    ${ }^{1}$ The "mild language-theoretic assumption" is that the class of languages forms an effective full trio: it is closed under intersections with regular languages, homomorphisms, and inverse homomorphisms. Many language classes studied in formal language theory and verification satisfy these conditions.

[^3]:    ${ }^{2}$ The models of HORS (used in model checking higher-order programs [Kob09]) and word schemes (used in language-theoretic exploration of downclosures [HKO16, CPSW16]) are somewhat different. Thus, we show an explicit reduction between the two formalisms.

[^4]:    ${ }^{3}$ The language class $\mathcal{C}$ in [GM12] is fixed to be the class of context free languages. Their definition generalizes to any language class.

