

Bounded Phase Analysis of Message-Passing Programs^{*,**}

Ahmed Bouajjani and Michael Emmi^{***}

LIAFA, Université Paris Diderot, France
{abou,mje}@liafa.jussieu.fr

Abstract. We describe a novel technique for bounded analysis of asynchronous message-passing programs with ordered message queues. Our bounding parameter does not limit the number of pending messages, nor the number of “contexts-switches” between processes. Instead, we limit the number of process communication cycles, in which an unbounded number of messages are sent to an unbounded number of processes across an unbounded number of contexts. We show that remarkably, despite the potential for such vast exploration, our bounding scheme gives rise to a simple and efficient program analysis by reduction to sequential programs. As our reduction avoids explicitly representing message queues, our analysis scales irrespectively of queue content and variation.

1 Introduction

Software is becoming increasingly concurrent: reactivity (e.g., in user interfaces, web servers), parallelization (e.g., in scientific computations), and decentralization (e.g., in web applications) necessitate asynchronous computation. Although shared-memory implementations are often possible, the burden of preventing unwanted thread interleavings without crippling performance is onerous. Many have instead adopted asynchronous programming models in which processes communicate by posting messages/tasks to others’ message/task queues— [19] discuss why such models provide good programming abstractions. Single-process systems such as the JavaScript page-loading engine of modern web browsers [1], and the highly-scalable Node.js asynchronous web server [11], execute a series of short-lived tasks one-by-one, each task potentially queueing additional tasks to be executed later. This programming style ensures that the overall system responds quickly to incoming events (e.g., user input, connection requests). In the multi-process setting, languages such as Erlang and Scala have adopted message-passing as a fundamental construct with which highly-scalable and highly-reliable distributed systems are built.

Despite the increasing popularity of such programming models, little is known about precise algorithmic reasoning. This is perhaps not without good reason:

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

** An extended version of this paper is online [6]

*** Supported by a Fondation Sciences Mathématiques de Paris post-doctoral fellowship.

decision problems such as state-reachability for programs communicating with unbounded reliable queues are undecidable [10], even when there is only a single finite-state process (posting messages to itself). Furthermore, the known decidable under-approximations (e.g., bounding the size of queues) represent queues explicitly, are thus doomed to combinatorial explosion as the size and variability of queue content increases.

Some have proposed analyses which abstract message arrival order [23, 14, 13], or assume messages can be arbitrarily lost [2, 3]. Such analyses do not suffice when correctness arguments rely on reliable messaging—several systems specifically do ensure the ordered delivery of messages, including Scala, and recent web-browser specifications [1]. Others have proposed analyses which compute finite symbolic representations of queue contents [5, 8]. Known bounded analyses which model queues precisely either bound the maximum capacity of message-queues, ignoring executions which exceed the bound, or bound the total number of process “contexts” [21, 16], where each context involves a single process sending and receiving messages. For each of these bounding schemes there are trivial systems which cannot be adequately explored, e.g., by sending more messages than the allowed queue-capacity, having more processes than contexts, or by alternating message-sends to two processes—we discuss such examples in Section 3. All of the above techniques represent queues explicitly, though perhaps symbolically, and face combinatorial explosion as queue content and variation increase.

In this work we propose a novel technique for bounded analysis of asynchronous message-passing programs with reliable, ordered message queues. Our bounding parameter, introduced in Section 3, is not sensitive to the capacity nor content of message queues, nor the number of process contexts. Instead, we bound the number of process communication cycles by labeling each message with a monotonically-increasing phase number. Each time a message chain visits the same process, the phase number must increase. For a given parameter k , we only explore behaviors of up to k phases—though k phases can go a long way. In the leader election distributed protocol [24] for example, each election round occurs in 2 phases: in the first phase each process sends *capture* messages to the others; in the second phase some processes receive *accept* messages, and those that find themselves majority-winners broadcast *elected* messages. In these two phases an unbounded number of messages are sent to an unbounded number of processes across an unbounded number of process contexts!

We demonstrate the true strength of phase-bounding by showing in Sections 4 and 5 that the bounded phase executions of a message-passing program can be concisely encoded as a non-deterministic sequential program, in which message-queues are not explicitly represented. Our so-called “sequentialization” sheds hope for scalable analyses of message-passing programs. In a small set of simple experiments (Section 4), we demonstrate that our phase-bounded encoding scales far beyond known explicit-queue encodings as queue-content increases, and even remains competitive as queue-content is fixed while the number of phases grows. By reducing to sequential programs, we leverage highly-developed sequential program analysis tools for message-passing programs.

2 Asynchronous Message-Passing Programs

We consider a simple multi-processor programming model in which each processor is equipped with a procedure stack and a queue of pending tasks. Initially all processors are idle. When an idle processor's queue is non-empty, the oldest task in its queue is removed and executed to completion. Each task executes essentially a recursive sequential program, which besides accessing its own processor's global storage, can *post* tasks to the queues of any processor, including its own. When a task does complete, its processor again becomes idle, chooses the next pending task to execute to completion, and so on. The distinction between queues containing messages and queues containing tasks is mostly aesthetic, but in our task-based treatment queues are only read by idle processors; reading additional messages during a task's execution is prohibited. While in principle many message-passing systems, e.g., in Erlang and Scala, allow reading additional messages at any program point, we have observed that common practice is to read messages only upon completing a task [25].

Though similar to [23]'s model of asynchronous programs, the model we consider has two important distinctions. First, tasks execute across potentially several processors, rather than only one, each processor having its own global state and pending tasks. Second, the tasks of each processor are executed in exactly the order they are posted. For the case of single-processor programs, [23]'s model can be seen as an abstraction of the model we consider, since there the task chosen to execute next when a processor is idle is chosen non-deterministically among all pending tasks.

2.1 Program Syntax

Let *Procs* be a set of procedure names, *Vals* a set of values, *Exprs* a set of expressions, *Pids* a set of processor identifiers, and let *T* be a type. Figure 1 gives the grammar of *asynchronous message-passing programs*. We intentionally leave the syntax of expressions *e* unspecified, though we do insist *Vals* contains **true** and **false**, and *Exprs* contains *Vals* and the (*nullary*) choice operator \star .

Each program *P* declares a single global variable *g* and a procedure sequence, each $p \in \text{Procs}$ having a single parameter *l* and top-level statement denoted s_p ; as statements are built inductively by composition with control-flow statements, s_p describes the entire body of *p*. The set of program statements *s* is denoted *Stmts*. Intuitively, a **post** ρ *p* *e* statement is an asynchronous call to a procedure *p* with argument *e* to be executed on the processor identified by ρ ; a *self-post* to one's own processor is made by setting ρ to $_$. A program in which all **post** statements are self-posts is called a *single-processor program*, and a program without **post** statements is called a *sequential program*.

The programming language we consider is simple, yet very expressive, since the syntax of types and expressions is left free, and we lose no generality by considering only single global and local variables. Our extended report [6] lists several syntactic extensions which we use in the source-to-source translations of the subsequent sections, and which easily reduce to the syntax of our grammar.

$$\begin{aligned}
P &::= \text{var } g:T (\text{proc } p (\text{var } l:T) s)^* \\
s &::= s; s \mid \text{skip} \mid x := e \\
&\mid \text{assume } e \\
&\mid \text{if } e \text{ then } s \text{ else } s \\
&\mid \text{while } e \text{ do } s \\
&\mid \text{call } x := p \ e \\
&\mid \text{return } e \\
&\mid \text{post } \rho \ p \ e \\
x &::= g \mid l
\end{aligned}$$

$$\begin{array}{c}
\text{DISPATCH} \\
\hline
\langle g, \varepsilon, f \cdot q \rangle \xrightarrow{S} \langle g, f, q \rangle \\
\\
\text{COMPLETE} \\
\hline
f = \langle \ell, \text{return } e; s \rangle \\
\langle g, f, q \rangle \xrightarrow{S} \langle g, \varepsilon, q \rangle \\
\\
\text{SELF-POST} \\
\hline
s_1 = \text{post } _ \ p \ e; \ s_2 \\
\ell_2 \in e(g, \ell_1) \quad f = \langle \ell_2, s_p \rangle \\
\hline
\langle g, \langle \ell_1, s_1 \rangle w, q \rangle \xrightarrow{S} \langle g, \langle \ell_1, s_2 \rangle w, q \cdot f \rangle
\end{array}$$

Fig. 1. The grammar of asynchronous message-passing programs P . Here T is an unspecified type, and e , p , and ρ range, resp., over expressions, procedure names, and processor identifiers.

Fig. 2. The single-processor transition rules \rightarrow^S ; see our extended report [6] for the standard sequential statements.

2.2 Single-Processor Semantics

A (*procedure*) *frame* $f = \langle \ell, s \rangle$ is a current valuation $\ell \in \text{Vals}$ to the procedure-local variable l , along with a statement $s \in \text{Stmts}$ to be executed. (Here s describes the entire body of a procedure p that remains to be executed, and is initially set to p 's top-level statement s_p ; we refer to initial procedure frames $t = \langle \ell, s_p \rangle$ as *tasks*, to distinguish the frames that populate processor queues.) The set of all frames is denoted Frames .

A *processor configuration* $\kappa = \langle g, w, q \rangle$ is a current valuation $g \in \text{Vals}$ to the processor-global variable g , along with a procedure-frame stack $w \in \text{Frames}^*$ and a pending-tasks queue $q \in \text{Frames}^*$. A processor is idle when $w = \varepsilon$. The set of all processor configurations is denoted Pconfigs . A processor configuration map $\xi : \text{Pids} \rightarrow \text{Pconfigs}$ maps each processor $\rho \in \text{Pids}$ to a processor configuration $\xi(\rho)$. We write $\xi(\rho \mapsto \kappa)$ to denote the configuration ξ updated with the mapping $(\rho \mapsto \kappa)$, i.e., the configuration ξ' such that $\xi'(\rho) = \kappa$, and $\xi'(\rho') = \xi(\rho')$ for all $\rho' \in \text{Pids} \setminus \{\rho\}$.

For expressions without program variables, we assume the existence of an evaluation function $\llbracket \cdot \rrbracket_e : \text{Exprs} \rightarrow \wp(\text{Vals})$ such that $\llbracket \star \rrbracket_e = \text{Vals}$. For convenience we define $e(g, \ell) \stackrel{\text{def}}{=} \llbracket e[g/g, \ell/l] \rrbracket_e$ to evaluate the expression e in a global valuation g by substituting the current values for variables g and l . As these are the only program variables, the substituted expression $e[g/g, \ell/l]$ has no free variables.

Figure 2 defines the transition relation \rightarrow^S for the asynchronous behavior of each processor; the standard transitions for the sequential statements are listed in our extended report [6]. The SELF-POST rule creates a new frame to execute the given procedure, and places the new frame in the current processor's pending-tasks queue. The COMPLETE rule returns from the final frame of a task, rendering the processor idle, and the DISPATCH rule schedules the least-recently posted task on a idle processor.

$$\begin{array}{c}
 \text{SWITCH} \\
 \hline
 \rho_2 \in \text{enabled}(m, \xi) \\
 \hline
 \langle \rho_1, \xi, m \rangle \xrightarrow{M} \langle \rho_2, \xi, m \rangle \\
 \\
 \text{STEP} \\
 \hline
 \begin{array}{l}
 \xi_1(\rho) \xrightarrow{S} \kappa \quad \xi_2 = \xi_1(\rho \mapsto \kappa) \\
 \rho \in \text{enabled}(m_1, \xi_1) \quad m_2 = \text{step}(m_1, \xi_1, \xi_2) \\
 \hline
 \langle \rho, \xi_1, m_1 \rangle \xrightarrow{M} \langle \rho, \xi_2, m_2 \rangle
 \end{array} \\
 \\
 \text{POST} \\
 \hline
 \begin{array}{l}
 \xi_1(\rho_1) = \langle g_1, \langle \ell_1, \mathbf{post} \ \rho_2 \ p \ e; \ s \rangle w_1, q_1 \rangle \\
 \xi_1(\rho_2) = \langle g_2, w_2, q_2 \rangle \\
 \rho_1 \neq \rho_2 \quad \ell_2 \in e(g_1, \ell_1) \quad f = \langle \ell_2, s_p \rangle \\
 \rho_1 \in \text{enabled}(m_1, \xi_1) \quad m_2 = \text{step}(m_1, \xi_1, \xi_3) \\
 \xi_2 = \xi_1(\rho_1 \mapsto \langle g_1, \langle \ell_1, s \rangle w_1, q_1 \rangle) \\
 \xi_3 = \xi_2(\rho_2 \mapsto \langle g_2, w_2, q_2 \cdot f \rangle) \\
 \hline
 \langle \rho_1, \xi_1, m_1 \rangle \xrightarrow{M} \langle \rho_1, \xi_3, m_2 \rangle
 \end{array}
 \end{array}$$

Fig. 3. The multi-processor transition relation \rightarrow_M parameterized by a scheduler $M = \langle D, \text{empty}, \text{enabled}, \text{step} \rangle$

```

// translation of var g: T
var G[k]: T

// translation of
// proc p (var l: T) s
proc p (var l: T, phase: k) s

// translation of g
G[phase]

// translation of call x := p e
call x := p (e, phase)

// translation of post _ p e
if phase+1 < k then
  call p (e, phase+1)
    
```

Fig. 4. The k -phase sequential translation $((P))_k$ of a single-processor asynchronous message-passing program P

2.3 Multi-processor Semantics

In reality the processors of multi-processor systems execute independently in parallel. However, as long as they either do not share memory, or access a sequentially consistent shared memory, it is equivalent, w.r.t. the observations of any single processor, to consider an *interleaving semantics*: at any moment only one processor executes. In order to later restrict processor interleaving, we make explicit the *scheduler* which arbitrates the possible interleavings. Formally, a *scheduler* $M = \langle D, \text{empty}, \text{enabled}, \text{step} \rangle$ consists of a data type D of scheduler objects $m \in D$, a scheduler constructor $\text{empty} \in D$, a scheduler decision function $\text{enabled} : (D \times (\text{Pids} \rightarrow \text{Pconfigs})) \rightarrow \wp(\text{Pids})$, and a scheduler update function $\text{step} : (D \times (\text{Pids} \rightarrow \text{Pconfigs}) \times (\text{Pids} \rightarrow \text{Pconfigs})) \rightarrow D$. The arguments to enabled allow a scheduler to decide which processors are enabled depending on the execution history. A scheduler is *deterministic* when $|\text{enabled}(m, \xi)| \leq 1$ for all $m \in D$ and $\xi : \text{Pids} \rightarrow \text{Pconfigs}$, and is *non-blocking* when for all m and ξ , if there is some $\rho \in \text{Pids}$ such that $\xi(\rho)$ is either non-idle or has pending tasks, then there exists $\rho' \in \text{Pids}$ such that $\rho' \in \text{enabled}(m, \xi)$ and $\xi(\rho')$ is either non-idle or has pending tasks. A *configuration* $c = \langle \rho, \xi, m \rangle$ is a currently executing processor $\rho \in \text{Pids}$, along with a processor configuration map ξ , and a scheduler object m .

Figure 3 defines the multi-processor transition relation \rightarrow_M , parameterized by a scheduler M . The SWITCH rule non-deterministically schedules any enabled processor, while the STEP rule executes one single-processor program step on the currently scheduled processor, and updates the scheduler object. Finally, the POST rule creates a new frame to execute the given procedure, and places the the new frame on the target processor's pending-tasks queue.

Until further notice, we assume M is a completely non-deterministic scheduler; i.e., all processors are always enabled. In Section 5 we discuss alternatives.

An M -execution of a program P (from c_0 to c_j) is a configuration sequence $c_0c_1 \dots c_j$ such that $c_i \rightarrow_M c_{i+1}$ for $0 \leq i < j$. An initial condition $\iota = \langle \rho_0, g_0, \ell_0, p_0 \rangle$ is a processor identifier ρ_0 , along with a global-variable valuation $g_0 \in \text{Vals}$, a local-variable valuation $\ell_0 \in \text{Vals}$, and a procedure $p_0 \in \text{Procs}$. A configuration $c = \langle \rho_0, \xi, \text{empty} \rangle$ of a program P is $\langle \rho_0, g_0, \ell_0, p_0 \rangle$ -initial when $\xi(\rho_0) = \langle g_0, \varepsilon, \langle \ell_0, s_{p_0} \rangle \rangle$ and $\xi(\rho) = \langle g_0, \varepsilon, \varepsilon \rangle$ for all $\rho \neq \rho_0$. A configuration $\langle \rho, \xi, m \rangle$ is g_f -final when $\xi(\rho') = \langle g_f, w, q \rangle$ for some $\rho' \in \text{Pids}$, and $w, q \in \text{Frames}^*$. We say a global valuation g is M -reachable in P from ι when there exists an M -execution of P from some c_0 to some c such that c_0 is ι -initial and c is g -final¹.

Definition 1. The state-reachability problem is to determine for an initial condition ι , valuation g , and program P , whether g is reachable in P from ι .

3 Phase-Bounded Execution

Because processors execute tasks precisely in the order which they are posted to their unbounded task-queues, our state-reachability problem is undecidable, even with only a single processor accessing finite-state data [10]. Since it is not algorithmically possible to consider every execution precisely, in what follows we present an incremental under-approximation. For a given bounding parameter k , we consider a subset of execution (prefixes) precisely; as k increases, the set of considered executions increases, and in the limit as k approaches infinity, every execution of any program is considered—though for many programs, every execution is considered with a finite value of k .

In a given execution, a task-chain $t_1t_2 \dots t_i$ from t_1 to t_i is a sequence of tasks² such that the execution of each t_j posts t_{j+1} , for $0 < j < i$, and we say that t_1 is an ancestor of t_i . We characterize execution prefixes by labeling each task t posted in an execution with a phase number $\varphi(t) \in \mathbb{N}$:

$$\varphi(t) = \begin{cases} 0 & \text{if } t \text{ is initially pending.} \\ \varphi(t') & \text{if } t \text{ is posted to processor } \rho \text{ by } t', \\ & \text{and } t \text{ has no phase-}\varphi(t') \text{ ancestor on processor } \rho. \\ \varphi(t') + 1 & \text{if } t \text{ is posted by } t', \text{ otherwise.} \end{cases}$$

For instance, considering Figure 5a, supposing all on a single processor, an initial task A_1 posts A_2, A_3 , and A_4 , then A_2 posts A_5 and A_6 , and then A_3 posts A_7 , which in turn posts A_8 and A_9 . Task A_1 has phase 0. Since each post is made to the same processor, the phase number is incremented for each posted task. Thus the phase 1 tasks are $\{A_2, A_3, A_4\}$, the phase 2 tasks are $\{A_5, A_6, A_7\}$, and the phase 3 tasks are $\{A_8, A_9\}$. Notice that tasks of a given phase only execute

¹ In the presence of the **assume** statement, only the values reached in completed executions are guaranteed to be valid.

² We assume each task in a given execution has implicitly a unique task-identifier.

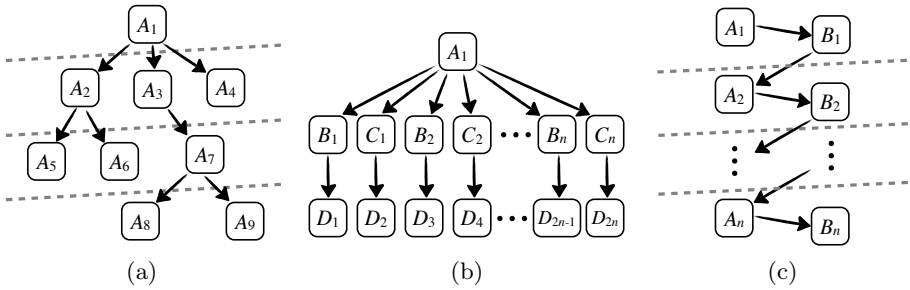


Fig. 5. Phase-bounded executions with processors A , B , C , and D ; each task’s label (e.g., A_i) indicates the processor it executes on (e.g., A). Arrows indicate the posting relation, indices indicate execution order on a given processor, and dotted lines indicate phase boundaries.

after all tasks of the previous phase have completed, i.e., execution order is in phase order; only executing tasks up to a given phase does correspond to a valid execution prefix.

Definition 2. An execution is k -phase when $\varphi(t) < k$ for each executed task t .

The execution in Figure 5a is a 4-phase execution, since all tasks have phase less than 4. Despite there being an arbitrary number $4n + 1$ of posted tasks, the execution in Figure 5b is 1-phase, since there are no task-chains between same-processor tasks. Contrarily, the execution in Figure 5c requires n phases to execute all $2n$ tasks, since every other occurrence of an A_i task creates a task-chain between A -tasks.

Note that bounding the number of execution phases does not necessarily bound the total number of tasks executed, nor the maximum size of task queues, nor the amount of switching between processors. Instead, a bound k restricts the maximum length of task chains to $k \cdot |\text{Pids}|$. In fact, phase-bounding is incomparable to bounding the maximum size of task queues. On the one hand, every execution of a program in which one root task posts an arbitrary, unbounded number of tasks to other processors (e.g., in Figure 5b) are explored with 1 phase, though no bound on the size of queues will capture all executions. On the other hand, all executions with a single arbitrarily-long chain of tasks (e.g., in Figure 5c) are explored with size 1 task queues, though no limited number of phases captures all executions. In the limit as the bounding parameter increases, either scheme does capture all executions.

Theorem 1 (Completeness). For every execution h of a program P , there exists $k \in \mathbb{N}$ such that h is a k -phase execution.

4 Phase-Bounding for Single-Processor Programs

Characterizing executions by their phase-bound reveals a simple and efficient technique for bounded exploration. This seems remarkable, given that phase-bounding explores executions in which arbitrarily many tasks execute, making

the task queue arbitrarily large. The first key ingredient is that once the number of phases is bounded, each phase can be executed in isolation. For instance, consider again the execution of Figure 5a. In phase 1, the tasks A_2 , A_3 , and A_4 pick up execution from the global valuation g_1 which A_1 left off at, and leave behind a global valuation g_2 for the phase 2 tasks. In fact, given the sequence of tasks in each phase, the only other “communication” between phases is a single passed global valuation; executing that sequence of tasks on that global valuation is a faithful simulation of that phase.

The second key ingredient is that the ordered sequence of tasks executed in a given phase is exactly the ordered sequence of tasks posted in the previous phase. This is obvious, since tasks are executed in the order they are posted. However, combined with the first ingredient we have quite a powerful recipe. Supposing the global state g_i at the beginning of each phase i is known initially, we can simulate a k -phase execution by executing each task posted to phase i as soon as it is posted, with an independent virtual copy of the global state, initially set to g_i . That is, our simulation will store a vector of k global valuations, one for each phase. Initially, the i^{th} global valuation is set to the state g_i in which phase i begins; tasks of phase i then read from and write to the i^{th} global valuation. It then only remains to ensure that the global valuations g_i used at the beginning of each phase $0 < i < k$ match the valuations reached at the end of phase $i - 1$.

This simulation is easily encoded into a non-deterministic sequential program with k copies of global storage. The program begins by non-deterministically setting each copy to an arbitrary value. Each task maintains their current phase number i , and accesses the i^{th} copy of global storage. Each posted task is simply called instead of posted, its phase number set to one greater than its parent—posts to tasks with phase number k are ignored. At the end of execution, the program ensures that the i^{th} global valuation matches the initially-used valuation for phase $i + 1$, for $0 \leq i < k - 1$. When this condition holds, any global valuation observed along the execution is reachable within k phases in the original program. Figure 4 lists a code-to-code translation which implements this simulation.

Theorem 2. *A global-valuation g is reachable in a k -phase execution of a single-processor program P if and only if g is reachable in $((P))_k$ —the k -phase sequential translation of P .*

When the underlying sequential program model has a decidable state-reachability problem, Theorem 2 gives a decision procedure for the phase-bounded state-reachability problem, by applying the decision procedure for the underlying model to the translated program. This allows us for instance to derive a decidability result for programs with finite data domains.

Corollary 1. *The k -phase state-reachability problem is decidable for single-processor programs with finite data domains.*

More generally, given any underlying sequential program model, our translation makes applicable any analysis tool for said model to message-passing programs, since the values of the additional variables are either from the finite domain $\{0, \dots, k - 1\}$, or in the domain of the original program variables.

Note that our simulation of a k -phase execution does not explicitly store the unbounded task queue. Instead of storing a multitude of possible unbounded task sequences, our simulation stores exactly k global state valuations. Accordingly, our simulation is not doomed to the unavoidable combinatorial explosion encountered by storing (even bounded-size) task queues explicitly. To demonstrate the capability of our advantage, we measure the time to verify two fabricated yet illustrative examples (listed in full in our extended report [6], comparing our bounded-phase encoding with a bounded task-queue encoding. In the bounded task-queue encoding, we represent the task-queue explicitly by an array of integers, which stores the identifiers of posted procedures³. When control of the initial task completes, the program enters a loop which takes a procedure identifier from the head of the queue, and calls the associated procedure. When the queue reaches a given bound, any further posted tasks are ignored.

The first program $P_1(i)$, parameterized by $i \in \mathbb{N}$, has a single Boolean global variable \mathbf{b} , i procedures named p_1, \dots, p_i , which assert \mathbf{b} to be **false** and set \mathbf{b} to **true**, and i procedures named q_1, \dots, q_i which set \mathbf{b} to **false**. Initially, $P_1(i)$ sets \mathbf{b} to **false**, and enters a loop in which each iteration posts some p_j followed by some q_j . Since a q_j task must be executed between each p_j task, each of the assertions are guaranteed to hold. Figure 6a compares the time required to verify $P_1(i)$ (using the BOOGIE verification engine [4]) for various values of i , and various bounds n on loop unrolling. Note that although every execution of $P_1(i)$ has only 2 phases, to explore all n loop iterations in any given execution, the size of queues must be at least $2n$, since two tasks are posted per iteration. Even for this very simple program, representing (even bounded) task-queues explicitly does not scale, since the number of possible task-queues grows astronomically as the size of task-queues grow. This ultimately prohibits the bounded tasks-queue encodings from exploring executions in which more than a mere few simple tasks execute. On the contrary, our bounded-phase simulation easily explores every execution up to the loop-unrolling bound in a few seconds.

To be fair, our second program P_2 is biased to support the bounded task-queue encoding. Following the example of Figure 5c, P_2 again has a single Boolean global variable \mathbf{b} , and two procedures: p_1 asserts \mathbf{b} to be **false**, sets \mathbf{b} to **true**, and posts p_2 , while p_2 sets \mathbf{b} to **false** and posts p_1 . Initially, the program P_2 sets \mathbf{b} to **false** and posts a single p_1 task. Again here, since a p_2 task must execute between each p_1 task, each of the assertions are guaranteed to hold. Figure 6b compares the time required to verify P_2 for various bounds n on the number of tasks explored⁴. Note that although every execution of P_2 uses only size 1 task-queues, to explore all n tasks in any given execution, the number of phases must be at least n , since each task must execute in its own phase. Although

³ For simplicity our examples do not pass arguments to tasks; in general, one should also store in the task-queue array the values of arguments passed to each posted procedure.

⁴ The number n of explored tasks is controlled by limiting the number of loop unrollings in the bounded task-queue encoding, and limiting the recursion depth, and phase-bound, in the bounded-phase encoding.

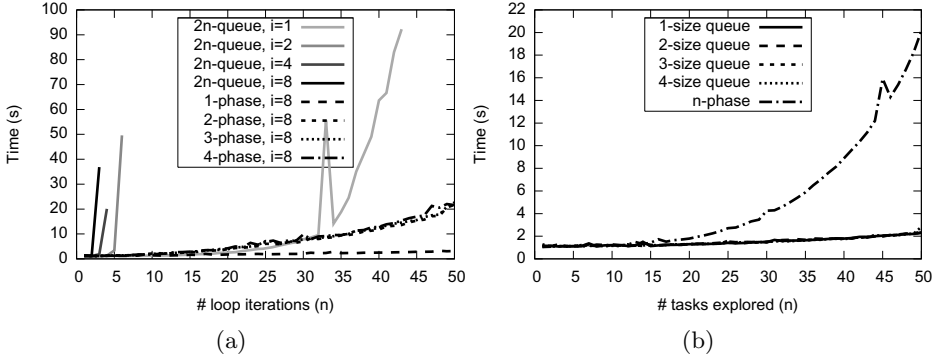


Fig. 6. Time required to verify (a) the program $P_1(i)$, and (b) the program P_2 with the BOOGIE verification engine using various encodings (bounded queues, bounded phase), and various loop unrolling bounds. Time-out is set to 100s.

verification time for the bounded-phase encoding does increase with n faster than the bounded task-queue encoding—as expected—due to additional copies of the global valuation, and more deeply in-lined procedures, the verification time remains manageable. In particular, the time does not explode uncontrollably: even 50 tasks are explored in under 20s.

5 Phase-Bounding for Multi-processor Programs

Though state-reachability under a phase bound is immediately and succinctly reducible to sequential program analysis for single-processor programs, the multi-processor case is more complicated. The added complexity arises due to the many orders in which tasks on separate processors can contribute to others’ task-queues. As a simple example, consider the possible bounded-phase executions of Figure 5b with four processors, A , B , C , and D . Though B ’s tasks B_1, \dots, B_n must be executed in order, and C ’s tasks C_1, \dots, C_n must also be executed in order, the order of D ’s tasks are not pre-determined: the arrival order of D ’s tasks depends on how B ’s and C ’s tasks *interleave*. Suppose for instance B_1 executes to completion before C_1 , which executes to completion before B_2 , and so on. In this case D ’s tasks arrive to D ’s queue, and ultimately execute, in the index order D_1, D_2, \dots as depicted. However, there exist executions for every possible order of D ’s tasks respecting $D_1 < D_3 < \dots$ and $D_2 < D_4 < \dots$ (where $<$ denotes an ordering constraint)—many possible orders indeed! In fact, due to the capability of such unbounded interleaving, the problem of state-reachability under a phase-bound is undecidable for multi-processor programs, even for programs with finite data domains.

Theorem 3. *The k -phase bounded state-reachability problem is undecidable for multi-processor programs with finite data domains.*

Note that Theorem 3 holds independently of whether memory is shared between processors: the fact that a task-queue can store any possible (unbounded) shuffling of tasks posted by two processors lends the power to simulate Post's correspondence problem [20].

Theorem 3 insists that phase-bounding alone will not lead to the elegant encoding to sequential programs which was possible for single-processor programs. If that were possible, then the translation from a finite-data program would lead to a finite-data sequential program, and thus a decidable state-reachability problem. Since a precise algorithmic solution to bounded-phase state-reachability is impossible for multi-processor programs, we resort to a further incremental yet orthogonal under-approximation, which limits the number of considered processor interleavings. The following development is based on delay-bounded scheduling [12].

We define a *delaying scheduler* $M = \langle D, \text{empty}, \text{enabled}, \text{step}, \text{delay} \rangle$, as a scheduler $\langle D, \text{empty}, \text{enabled}, \text{step} \rangle$, along with a function $\text{delay} : (D \times \text{Pids} \times (\text{Pids} \rightarrow \text{Pconfigs})) \rightarrow D$. Furthermore, we extend the transition relation of Figure 3 with a postponing rule of Figure 7 which we henceforth refer to as a *delay (operation)*, saying that processor ρ is delayed. Note that a delay operation may or may not change the set of enabled processors in any given step, depending on the scheduler. A delaying scheduler is *delay-accessible* when for every configuration c_1 and non-idle or task-pending processor ρ , there exists a sequence $c_1 \rightarrow_M \dots \rightarrow_M c_j$ of DELAY-steps such that ρ is enabled in c_j . Given executions h_1 and h_2 of (delaying) schedulers M_1 and M_2 resp., we write $h_1 \sim h_2$ when h_1 and h_2 are identical after projecting away delay operations.

Definition 3. *An execution with at most k delay operators is called k -delay.*

Consider again the possible executions of Figure 5b, but suppose we fix a deterministic scheduler M which without delaying would execute D 's tasks in index order: D_1, D_2, \dots ; furthermore suppose that delaying a processor ρ in phase i causes M to execute the remaining phase i tasks of ρ in phase $i + 1$, while keeping the tasks of other processors in their current phase. Without using any delays, the execution of Figure 5b is unique, since M is deterministic. However, as Figure 8 illustrates, using a single delay, it is possible to also derive the order $D_1, D_3, \dots, D_{2n-1}, D_2, D_4, \dots, D_{2n}$ (among others): simply delay processor C once before C_1 posts D_2 . Since this forces the D_{2i} tasks posted by each C_i to occur in the second phase, it follows they must all happen after the D_{2i-1} tasks posted by each B_i .

Theorem 4 (Completeness). *Let M be any delay-accessible scheduler. For every execution h of a program P , there exists an M -execution h' and $k \in \mathbb{N}$ such that h' is a k -delay execution and $h' \sim h$.*

Note that Theorem 4 holds for *any* delay-accessible scheduler M —even deterministic schedulers. As it turns out there is one particular scheduler M_{bfs} for which we know a convenient sequential encoding, and this scheduler is described in our extended report [6]. For the moment, the important points to note are

$$\text{DELAY} \frac{m_2 = \text{delay}(m_1, \rho, \xi)}{\langle \rho, \xi, m_1 \rangle \xrightarrow{M} \langle \rho, \xi, m_2 \rangle}$$

Fig. 7. The delay operation




Fig. 8. A 2-phase delaying execution varying the 1-phase execution of Figure 5b

```
// translation of var g: T
var G[k+d]: T
var shift[Pids][k], delay: d
var ancestors[Pids][k+d]: B

// translation of proc p (var l: T) s
proc p (var l: T, pid: Pids, phase: k)

// translation of g
G[ phase + shift[pid][phase] ]

// code to be sprinkled throughout
while * and delay < d do
    shift[pid][phase]++; delay++

// translation of call x := p e
call x := p (e, pid, phase)

// translation of post rho p e
let p = phase + shift[pid][phase] in
let p' = p + (if ancestors[rho][p] then 1 else 0) in
if p' < k then
    ancestors[rho][p' + shift[rho][p']] := true;
    call p (e, rho, p')
    ancestors[rho][p' + shift[rho][p']] := false
```

Fig. 9. The k -phase d -delay sequential translation $((P))_{k,d}^{\text{bfs}}$ of a multi-processor message-passing asynchronous program P

that M_{bfs} is deterministic, non-blocking, and delay-accessible. Essentially, determinism allows us to encode the scheduler succinctly in a sequential program; the non-blocking property ensures this scheduler does explore some execution, rather than needlessly ceasing to continue; delay-accessibility combined with Theorem 4 ensure the scheduler is complete in the limit. Figure 9 lists a code-to-code translation which encodes bounded-phase and bounded-delay exploration of a given program according to the M_{bfs} scheduler as a sequential program.

Our translation closely follows the single-processor translation of Section 4, the key differences being:

- the **phase** of a posted task is not necessarily incremented, since posted tasks may not have same-processor ancestors in the current phase, and
- at any point, the currently executing task may increment a **delay** counter, causing all following tasks on the same processor to shift forward one additional phase.

As the global values reached by each processor at the end of each phase $i - 1$ must be ensured to match the initial values of phase i , for $0 < i < k + d$, so must the values for the **shift** counter: an execution is only valid when for each processor $\rho \in \text{Pids}$ and each phase $0 < i < k$, $\text{shift}[\rho][i - 1]$ matches the initial value of $\text{shift}[\rho][i]$.

Theorem 5. *A global valuation g is reachable in a k -phase d -delay M_{bfs} -execution of a multi-processor program P if and only if g is reachable in $((P))_{k,d}^{\text{bfs}}$.*

As is the case for our single-processor translation, our simulation does not explicitly store the unbounded tasks queue, and is not doomed to combinatorial explosion faced by storing tasks-queues explicitly.

6 Related Work

Our work follows the line of research on compositional reductions from concurrent to sequential programs. The initial so-called “sequentialization” [22] explored multi-threaded programs up to one context-switch between threads, and was later expanded to handle a parameterized amount of context-switches between a statically-determined set of threads executing in round-robin order [21, 18]. [17] later extended the approach to handle programs parameterized by an unbounded number of statically-determined threads, and shortly after, [12] further extended these results to handle an unbounded amount of dynamically-created tasks, which besides applying to multi-threaded programs, naturally handles asynchronous event-driven programs [23]. [9] pushed these results even further to a sequentialization which attempts to explore as many behaviors as possible within a given analysis budget. Each of these sequentializations necessarily do provide a bounding parameter which limits the amount of interleaving between threads or tasks, but none are capable of precisely exploring tasks in creation order, which is abstracted away from their program models [23]. [15]’s sequentialization is sensitive to task priorities, their reduction assumes a finite number of statically-determined tasks.

In a closely-related work, [16] propose a “context-bounded” analysis of shared-memory multi-pushdown systems communicating with message-queues. According to this approach, one “context” involves a single process reading from its queue, and posting to the queues of other processes, and the number of contexts per execution is bounded. Our work can be seen as an extension in a few ways. First, and most trivially, in their setting a process cannot post to its own message queue; this implies that at least $2k$ contexts must be used to simulate k phases of a single-processor program. Second, there are families of 1-phase executions which require an unbounded number of task-contexts to capture; the execution order $D_1D_2D_3 \dots D_{2n}$ of Figure 5b is such an example. We conjecture that bounded phase and delay captures context-bounding—i.e., there exists a polynomial function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that every k -context bounded execution of any program P is also a $f(k)$ -phase and delay bounded execution. Finally, though phase-bounding leads to a convenient sequential encoding, we are unaware whether a similar encoding is possible for context-bounding.

[5] and [7] have proposed analyses of message-passing programs by computing explicit finite symbolic representations of message-queues. As our sequentialization does not represent queues explicitly, we do not restrict the content

of queues to conveniently-representable descriptions. Furthermore, reduction to sequential program analyses is easily implementable, and allows us to leverage highly-developed and optimized program analysis tools.

7 Conclusion

By introducing a novel phase-based characterization of message-passing program executions, we enable bounded program exploration which is not limited by message-queue capacity nor the number of processors. We show that the resulting phase-bounded analysis problems can be solved by concise reduction to sequential program analysis. Preliminary evidence suggests our approach is at worst competitive with known task-order respecting bounded analysis techniques, and can easily scale where those techniques quickly explode.

Acknowledgments. We thank Constantin Enea, Cezara Dragoi, Pierre Ganty, and the anonymous reviewers for helpful feedback.

References

- [1] HTML5: A vocabulary and associated APIs for HTML and XHTML, <http://dev.w3.org/html5/spec/Overview.html>
- [2] Abdulla, P.A., Jonsson, B.: Verifying programs with unreliable channels. In: LICS 1993: Proc. 8th Annual IEEE Symposium on Logic in Computer Science, pp. 160–170. IEEE Computer Society (1993)
- [3] Abdulla, P.A., Bouajjani, A., Jonsson, B.: On-the-Fly Analysis of Systems with Unbounded, Lossy FIFO Channels. In: Vardi, M.Y. (ed.) CAV 1998. LNCS, vol. 1427, pp. 305–318. Springer, Heidelberg (1998)
- [4] Barnett, M., Leino, K.R.M., Moskal, M., Schulte, W.: Boogie: An intermediate verification language, <http://research.microsoft.com/en-us/projects/boogie/>
- [5] Boigelot, B., Godefroid, P.: Symbolic verification of communication protocols with infinite state spaces using QDDs. *Formal Methods in System Design* 14(3), 237–255 (1999)
- [6] Bouajjani, A., Emmi, M.: Bounded phase analysis of message-passing programs (2011), <http://hal.archives-ouvertes.fr/hal-00653085/en>
- [7] Bouajjani, A., Habermehl, P.: Symbolic reachability analysis of fifo-channel systems with nonregular sets of configurations. *Theor. Comput. Sci.* 221(1-2), 211–250 (1999)
- [8] Bouajjani, A., Habermehl, P., Vojnar, T.: Verification of parametric concurrent systems with prioritised FIFO resource management. *Formal Methods in System Design* 32(2), 129–172 (2008)
- [9] Bouajjani, A., Emmi, M., Parlato, G.: On Sequentializing Concurrent Programs. In: Yahav, E. (ed.) SAS 2011. LNCS, vol. 6887, pp. 129–145. Springer, Heidelberg (2011)
- [10] Brand, D., Zafriropulo, P.: On communicating finite-state machines. *J. ACM* 30(2), 323–342 (1983)
- [11] Dahl, R.: Node.js: Evented I/O for V8 JavaScript, <http://nodejs.org/>

- [12] Emmi, M., Qadeer, S., Rakamarić, Z.: Delay-bounded scheduling. In: POPL 2011: Proc. 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, pp. 411–422. ACM (2011)
- [13] Ganty, P., Majumdar, R.: Algorithmic verification of asynchronous programs. CoRR, abs/1011.0551 (2010), <http://arxiv.org/abs/1011.0551>
- [14] Jhala, R., Majumdar, R.: Interprocedural analysis of asynchronous programs. In: POPL 2007: Proc. 34th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, pp. 339–350. ACM (2007)
- [15] Kidd, N., Jagannathan, S., Vitek, J.: One Stack to Run Them All: Reducing Concurrent Analysis to Sequential Analysis Under Priority Scheduling. In: van de Pol, J., Weber, M. (eds.) SPIN 2010. LNCS, vol. 6349, pp. 245–261. Springer, Heidelberg (2010)
- [16] La Torre, S., Madhusudan, P., Parlato, G.: Context-Bounded Analysis of Concurrent Queue Systems. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 299–314. Springer, Heidelberg (2008)
- [17] La Torre, S., Madhusudan, P., Parlato, G.: Model-Checking Parameterized Concurrent Programs Using Linear Interfaces. In: Touili, T., Cook, B., Jackson, P. (eds.) CAV 2010. LNCS, vol. 6174, pp. 629–644. Springer, Heidelberg (2010)
- [18] Lal, A., Reps, T.W.: Reducing concurrent analysis under a context bound to sequential analysis. *Formal Methods in System Design* 35(1), 73–97 (2009)
- [19] Miller, M.S., Tribble, E.D., Shapiro, J.S.: Concurrency Among Strangers. In: De Nicola, R., Sangiorgi, D. (eds.) TGC 2005. LNCS, vol. 3705, pp. 195–229. Springer, Heidelberg (2005)
- [20] Post, E.L.: A variant of a recursively unsolvable problem. *Bull. Amer. Math. Soc.* 52(4), 264–268 (1946)
- [21] Qadeer, S., Rehof, J.: Context-Bounded Model Checking of Concurrent Software. In: Halbwachs, N., Zuck, L.D. (eds.) TACAS 2005. LNCS, vol. 3440, pp. 93–107. Springer, Heidelberg (2005)
- [22] Qadeer, S., Wu, D.: KISS: Keep it simple and sequential. In: PLDI 2004: Proc. ACM SIGPLAN Conference on Programming Language Design and Implementation, pp. 14–24. ACM (2004)
- [23] Sen, K., Viswanathan, M.: Model Checking Multithreaded Programs with Asynchronous Atomic Methods. In: Ball, T., Jones, R.B. (eds.) CAV 2006. LNCS, vol. 4144, pp. 300–314. Springer, Heidelberg (2006)
- [24] Svensson, H., Arts, T.: A new leader election implementation. In: Erlang 2005: Proc. ACM SIGPLAN Workshop on Erlang, pp. 35–39. ACM (2005)
- [25] Trottier-Hebert, F.: Learn you some Erlang for great good!, <http://learnyousomeerlang.com/>