

Safety Verification of Phaser Programs

Zeinab Ganjei, Ahmed Rezine, Petru Eles, Zebo Peng
dept. of computer and information science
Linköping University, Sweden
firstname.surname@liu.se

Abstract—We address the problem of statically checking control state reachability (as in possibility of assertion violations, race conditions or runtime errors) and plain reachability (as in deadlock-freedom) of *phaser programs*. Phasers are a modern non-trivial synchronization construct that supports dynamic parallelism with runtime registration and deregistration of spawned tasks. They allow for collective and point-to-point synchronizations. For instance, phasers can enforce barriers or producer-consumer synchronization schemes among all or subsets of the running tasks. Implementations are found in modern languages such as Habanero Java. Phasers essentially associate phases to individual tasks and use their runtime values to restrict possible concurrent executions. Unbounded phases may result in infinite transition systems even in the case of programs only creating finite numbers of tasks and phasers. We introduce an exact gap-order based procedure that always terminates when checking control reachability for programs generating bounded numbers of coexisting tasks and phasers. We also show verifying plain reachability is undecidable even for programs generating few tasks and phasers. We then explain how to turn our procedure into a sound analysis for checking plain reachability (including deadlock freedom). We report on preliminary experiments with our open source tool.

Index Terms—phasers, safety verification, dynamic synchronization, collective synchronization, Point-to-point synchronization, model checking

I. INTRODUCTION

We focus on safety verification of programs using *phasers* for task synchronization [1]–[3]. This sophisticated construct dynamically unifies collective and point-to-point synchronizations. For instance, it allows for dynamic registration and deregistration of tasks allowing for a more balanced usage of the computing resources when compared to static producer-consumer or barrier constructs [4]. The construct can be added to any parallel programming language with a shared address space. For instance, it can be found in Habanero Java [3], an extension of the Java programming language. Phasers build on the clock construct from the X10 programming language [1]. They can be created dynamically and spawned tasks may get registered or deregistered at runtime.

Intuitively, each phaser associates two phases (hereafter *wait* and *signal* phases) to each registered task. Apart from creating phasers and registering each other to them, tasks can individually issue *wait* and *signal* commands to a phaser they are registered to. Intuitively, *signal* commands are used to inform other registered tasks the issuing task is done with its *signal* phase. The command is non-blocking. It increments

the *signal* phase associated to the issuing task on the given phaser. The *wait* command is instead used to check whether all registered tasks are done with (i.e., have a *signal* phase that is strictly larger than) the issuing task’s *wait* phase. This command may get blocked by a task that did not yet finish the corresponding phase. Unlike classical barriers, phasers need not force registered tasks to wait for each other at each single phase. Instead they allow them to proceed with the following phases (by issuing *signal* commands), or even to exit the construct by deregistering from the phaser. Such dynamic behavior allows for better load balancing and performance, but comes at the price of making it easy to introduce programming mistakes such as assertion violations, race conditions, runtime errors and, in the important situation where *wait* and *signal* commands are decoupled for maximum flexibility, deadlocks. We summarize our contributions in this work:

- We propose an operational model based on [2], [3], [5].
- We show undecidability of checking deadlock-freedom for programs with fixed numbers of tasks and phasers.
- We describe an exact gap-order based symbolic verification procedure for checking control state reachability (as in assertion violations, race conditions or runtime errors) and plain reachability (as in checking deadlock freedom).
- We show termination of the procedure for control state reachability when numbers of tasks and phasers are fixed.
- We describe how to turn the procedure into a sound over-approximation for plain reachability.
- We report on our preliminary experiments with our open source tool.

Related work. We are not aware of automatic formal verification works that focus on constructs allowing for such a degree of dynamic parallelism. Unlike [6], we focus on fully automatic verification and consider the richer and more challenging phaser construct. The work of [5] considers the dynamic verification of phaser programs and can therefore only reason about particular program inputs and runs. The work in [7] uses Java Path Finder [8] to explore several runs, but still for one concrete input at a time. The works in [9], [10] target gap-order systems. Although phaser programs share some of their properties (larger gaps can do more), the results in [9], [10] do not apply since gap-order systems crucially forbid exact increments.

Outline. We describe a phaser program and recall some preliminaries in Sections II and III. This is followed in Section IV by a formal description of phaser programs and

of the properties we want to check. We also establish the undecidability of checking deadlock freedom. We introduce a gap-order based symbolic representation in Section V and describe in Section VI a simple verification procedure. We then show decidability of checking control state reachability and introduce a relaxation procedure for checking plain reachability. Finally, we report on our experiments and conclude the work. Descriptions of the proofs can be found in [11].

II. MOTIVATING EXAMPLE

The program listed in Fig. (1) uses Boolean shared variables $B = \{a, b, \text{done}\}$. A *main* task creates two phasers (lines 5 and 6). When creating a phaser, the task gets automatically registered to it. The main task also creates three other task instances (lines 9, 10 and 11). Several tasks can be registered to several phasers. When a task t is registered to a phaser p , a pair of numbers $(\text{wait}_p^t, \text{sig}_p^t)$, each in $\mathbb{N} \cup \{+\infty\}$, is associated to the couple (t, p) . The pair represents the individual *wait* and *signal* phases of task t on phaser p .

Registration of a task to a phaser can occur in one of three modes: SIG_WAIT, WAIT and SIG. In SIG_WAIT mode, a task may issue both `signal` and `wait` commands. In WAIT mode, a task may only issue `wait` commands on the phaser. Finally, when registered in SIG mode, a task may only issue `signal` commands. Issuing a `signal` command by a task on a phaser results in the task incrementing its signal phase associated to the phaser. This command is non-blocking. On the other-hand, issuing a `wait` command by a task on a phaser p will block until **all** tasks registered on p exhibit signal values on p that are strictly larger than the wait value of the issuing task on phaser p . In this case, the wait phase of the issuing task is incremented. Intuitively, a `signal` command allows the issuing task to state other tasks need not wait for it to complete its signal phase. In retrospect, a `wait` command allows a task to make sure all registered tasks have moved past its wait phase.

Upon creation of a phaser, wait and signal phases are initialized to 0 (except in WAIT mode where the signal phase is instead initialized to $+\infty$ in order to not block other waiters). The only other way a task may get registered to a phaser is if an already registred task does register it in the same mode (or in WAIT or SIG if the registrar is registered in SIG_WAIT). In this case, wait and signal phases of the newly registered task are initialized to those of the registrar. Tasks are therefore dynamically registered (e.g., lines 9-11). They can also dynamically deregister themselves (e.g., lines 25-26);

In this example, two producers and one consumer are synchronized using two phasers. The consumer requires the two producers to be ahead of it (wrt. the phaser main pointed to with `prod`) in order for it to consume their respective products. At the same time, the consumer needs to be ahead of both producers (wrt. the phaser main pointed to with `cons`) in order for these to produce their pair of products. It should be clear that phasers can be used as barriers for synchronizing dynamic subsets of concurrent tasks. Observe producers need not, in general, proceed in a lock step fashion. Producers may produce many items before consumers “catch up”.

We are interested in checking: (a) control state reachability as in assertions (e.g., line 44), race conditions (e.g., mutual exclusion of lines 20 and 49) or runtime errors (e.g., signaling a dropped phaser), and (b) plain reachability as in deadlocks (e.g., a producer at line 23 and a consumer at line 50 with equal phases). Intuitively, both problems concern themselves with the reachability of target sets of program configurations. The difference is that control state reachability defines the targets with the states of the tasks (their control locations and whether they are registered to some phasers). Plain reachability can, in addition, use values or relations between values of involved phases. Observe that control state reachability depends on the values of the actual phases, but these values are not used to define the target sets. For example, assertions are expressed as predicates over Boolean variables (e.g., line 44). Establishing such an assertion requires capturing the constraints imposed by the phasers on the program behaviors.

Our work proposes a sound and complete algorithm for checking control state reachability in case a bounded number of tasks and phasers are generated. The algorithm can handle arbitrarily large phases, e.g., generated using nested signaling loops. The algorithm starts from a symbolic representation of all bad configurations and successively computes sets of predecessor configurations. We show termination based on a well-quasi-ordering argument that imposes restrictions on what can be expressed with our symbolic representation. For instance putting upper bounds on differences between phases is forbidden. Deadlock configurations cannot be faithfully captured with such restricted representations. Intuitively, a deadlocked configuration will have a cycle where each involved task is waiting for the task to its right but where the wait phase of each task equals the signal phase of the task it is waiting for. We show the problem of checking deadlock freedom to be undecidable even for programs only generating a bounded number of tasks and phasers. We explain how to turn our verification algorithm into a sound but incomplete procedure for checking deadlock-freedom. Precision can then be augmented on demand to eliminate false positives.

III. PRELIMINARIES

We use \mathbb{N} and \mathbb{Z} for natural and integer numbers respectively. We write $A \uplus B$ to mean the union of disjoint sets A and B . We let $\text{Pfn}(A, B)$ be the set of partial functions from A to B and use \emptyset_A for the empty function over A , i.e., $\emptyset_A(a)$ is undefined (written $\emptyset_A(a) \uparrow$) for all $a \in A$. Given function $g \in \text{Pfn}(A, B)$ we write $g(a) \downarrow$ to mean that $g(a)$ is defined and write $g \setminus \{a\}$ to mean the restriction of g to the domain $A \setminus \{a\}$. We write $g[a \leftarrow b]$ for the function that coincides with g on A except for a that is sent to b . We abuse notation and let, for pairwise different $\{a_i \mid i \in I\}$, $g[\{a_i \leftarrow b_i \mid i \in I\}]$ mean the function that coincides with g on A except for each a_i that is sent to the corresponding b_i . We sometimes write a function g as a set $\{a \mapsto g(a) \mid a \in A\}$. It is then implicitly undefined outside of A .

```

1 bool a, b, done;
2 main()
3 {
4   done = false;
5   prod = newPhaser(SIG_WAIT);
6   cons = newPhaser(SIG_WAIT);
7   cons.signal();
8
9   async(aProducer, prod(SIG), cons(WAIT));
10  async(bProducer, prod(SIG), cons(WAIT));
11  async(abConsumer, prod(WAIT), cons(SIG));
12
13  prod.drop();
14  cons.drop();
15 }
16
17 aProducer(p(SIG), c(WAIT))
18 {
19   c.wait();
20   while(!done){
21     a = true;
22     p.signal();
23     c.wait();
24   };
25   p.drop();
26   c.drop();
27 }

```

```

28 bProducer(p(SIG), c(WAIT))
29 {
30   c.wait();
31   while(!done){
32     b = true;
33     p.signal();
34     c.wait();
35   };
36   p.drop();
37   c.drop();
38 }
39
40 abConsumer(p(WAIT), c(SIG))
41 {
42   while(!done){
43     p.wait();
44     assert(a & b);
45     a = false;
46     b = false;
47
48     if(ndet())
49       done = true;
50     c.signal();
51   };
52   c.drop();
53   p.drop();
54 }

```

Fig. 1. Two producers and one consumer are synchronized using two phasers. In this construction, the consumer requires both producers to be ahead of it (wrt. the prod phaser) in order for it to consume their respective products. At the same time, the consumer needs to be ahead of both producers (wrt. the cons phaser) in order for these to be able to produce their pair of products.

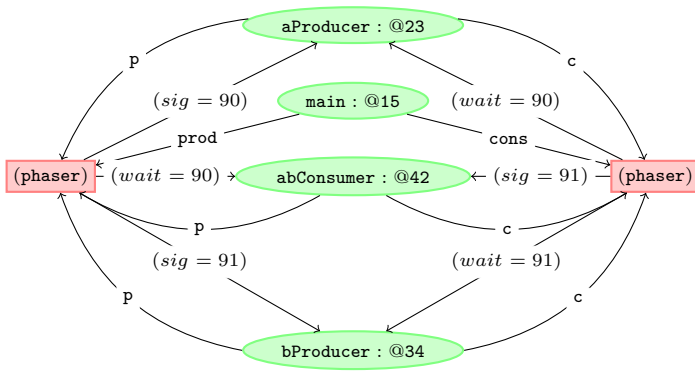


Fig. 2. Possible wait and signal phase values for Fig. (1). Observe that there is no a priori bound on the values of the different wait and signal phases. In this example, the difference between signal and wait phases is bounded. This is not always the case in general.

IV. LANGUAGE

A program may use a set B of shared Boolean variables and a set V of local phaser variables:

```

prg ::= bool b1, ..., b|B|;
      task1(v1, ..., vk1) {stmt1}
      ...
      taskn(v1, ..., vkn) {stmtn}

stmt ::= v = newPhaser() | async(task, v1, ..., vk)
       | v.drop() | v.signal() | v.wait() | exit
       | stmt; stmt | b = cond | assert(cond)
       | while(cond) {stmt} | if(cond) {stmt}

cond ::= ndet() | true | false | b | cond ∨ cond
       | cond ∧ cond | ¬cond

```

A program consists in a set of tasks T . A task is declared with $\text{task}(v_1, \dots, v_k) \{ \text{stmt} \}$ where v_1, \dots, v_k are phaser variables

that are local to the declared task. A task can also create a new phaser with $v = \text{newPhaser}()$ and store the identifier of the phaser in a local variable v . We let V be the union of all local phaser variables. When creating a phaser, a task gets registered to it. To simplify our description, we will assume all registrations to be in SIG_WAIT mode. Including the other modes is a matter of changing the initial phase values at registration and of statically ensuring the issued commands respect the registration mode. A task can deregister itself from a phaser referenced by a variable v with $v.\text{drop}()$. It can also issue signal or wait commands on a phaser on which it is registered and that is referenced by v . A task can spawn another task with $\text{async}(\text{task}, v_1, \dots, v_n)$. The issuing task registers the spawned task to the phasers it points to with v_1, \dots, v_n . The issuing task need not wait for the spawned task and may directly continue its execution.

Assume a phaser program $\text{prg} = (B, V, T)$. We inductively define the finite set S of control sequences as follows. S is the smallest set containing: (i) suffixes of each “ stmt_i ” appearing in some “ $\text{task}_i(v_{1_i}, \dots, v_{k_i}) \{ \text{stmt}_i \}$ ”; and (ii) suffixes of “ $\text{stmt}_i; \text{while}(\text{cond}) \{ \text{stmt}_i \}; \text{stmt}_j$ ” (respectively “ $\text{stmt}_i; \text{while}(\text{cond}) \{ \text{stmt}_i \}$ ”) for each “ $\text{while}(\text{cond}) \{ \text{stmt}_i \}; \text{stmt}_j$ ” (respectively “ $\text{while}(\text{cond}) \{ \text{stmt}_i \}$ ”) in S ; and (iii) suffixes of “ $\text{stmt}_i; \text{stmt}_j$ ” (respectively “ stmt_i ”) for each “ $\text{if}(\text{cond}) \{ \text{stmt}_i \}; \text{stmt}_j$ ” (respectively “ $\text{if}(\text{cond}) \{ \text{stmt}_i \}$ ”) appearing in S . We write s to mean some control sequence in S , and $\text{hd}(s)$ and $\text{tl}(s)$ to respectively mean the head and the tail of the sequence s .

A. Semantics.

A configuration c of $\text{prg} = (\mathbb{B}, \mathbb{V}, \mathbb{T})$ is a tuple $(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \varphi)$ where:

- \mathcal{T} is the current finite set of task identifiers. We let t, u range over the values in \mathcal{T} .
- \mathcal{P} is the current finite set of phaser identifiers. We let p, q range over the values in \mathcal{P} .
- $\mathbf{bv} : \mathbb{B} \rightarrow \{\text{true}, \text{false}\}$ is a total mapping that associates a value to each $b \in \mathbb{B}$.
- $\mathbf{pc} : \mathcal{T} \rightarrow \mathbb{S}$ is a total mapping that associates tasks to their remaining sequences (i.e., control location).
- $\mathbf{pv} : \mathcal{T} \rightarrow \text{Pfn}(\mathbb{V}, \mathcal{P})$ is a total mapping that associates, to each task identifier in \mathcal{T} , a partial mapping from the local phaser variables \mathbb{V} to phaser identifiers \mathcal{P} . It captures the values of the phaser variables \mathbb{V} of each task.
- $\varphi : \mathcal{P} \rightarrow \text{Pfn}(\mathcal{T}, \mathbb{N}^2)$ is a total mapping that associates to each phaser $p \in \mathcal{P}$ a partial mapping $\varphi(p)$ that is defined exactly on the identifiers of the tasks registered to p . For such a task t , $\varphi(p)(t)$ is the pair $(\text{wait}_p^t, \text{sig}_p^t)$ representing wait and signal values of t on p .

The set of tasks \mathcal{T} is altered by `asynch(task, v1, ..., vn)` and `exit` statements (rules `asynch` and `exit` in Fig.(3)). The set of phasers \mathcal{P} is updated upon creation of new phasers (rule `newPhaser` in Fig.(3)). The mapping \mathbf{pv} associates values to program phaser variables. Accessing variables with undefined values, or phasers to which the task is not currently registered, leads to runtime errors (rule `runtime error`). The total mapping φ captures states of phasers. It associates to each phaser identifier p in \mathcal{P} a partial mapping $\varphi(p)$. This partial mapping is defined for a task identifier $t \in \mathcal{T}$ (i.e., $\varphi(p)(t) \downarrow$) iff the task t is registered to the phaser p . In this case, $\varphi(p)$ gives the waiting phase wait_p^t and the signaling phase sig_p^t of the task t on the phaser p . Initially, a unique “main” task t_0 starts executing its `stmtmain` with no phasers. φ is the empty function with an empty domain \emptyset_\emptyset . After a task t executes a `v := newPhaser()` statement (rule `newPhaser` in Fig.(3)), a new phaser p is associated to the variable v using \mathbf{pv} and $\varphi(p)$ becomes the partial function $\{t \mapsto (0, 0)\}$. The initial configuration is $c_{\text{init}} = (\{t_0\}, \{\}, \mathbf{bv}_{\text{false}}, \{t_0 \mapsto \text{stmt}\}, \emptyset, \emptyset)$, where a “main” task with identifier t_0 and code `stmt` is the unique initial task. No phasers are present in the initial configuration, and all Boolean variables are mapped to `false`.

Given two configurations c and c' with $c = (\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \varphi)$, we write $c \xrightarrow{t} c'$ if there is a task $t \in \mathcal{T}$ such that one of the rules in Fig.(3) holds. We use $\xrightarrow{*}$ for the reflexive transitive closure of \rightarrow and write $c \xrightarrow{*} c'$ to mean that c' is reachable from c . A configuration is said reachable if it is reachable from the initial configuration c_{init} .

1) *Control-state reachability*: Checking the possibility of assertion violations, of runtime errors and of race conditions amounts to checking reachability of configurations respectively in $\text{badConfs}_{\text{assert}}^{(n,p)}$, $\text{badConfs}_{\text{runtime}}^{(n,p)}$ and in $\text{badConfs}_{\text{race}}^{(n,p)}$ for some number of tasks n and number of phasers p . We introduce in Section V a complete procedure

for checking reachability of such sets of configurations and show it to be sound for programs with fixed upper bounds on numbers of generated phasers and tasks.

2) *Deadlocks as in plain reachability*: We are also interested in checking the possibility of deadlocks. For this we need to define the notion of a blocked task. Assume in the following a configuration $c = (\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \varphi)$.

Definition 1 (Blocked). *A task $t \in \mathcal{T}$ is blocked at phaser $p \in \mathcal{P}$ by task $u \in \mathcal{T}$ if $\text{hd}(\mathbf{pc}(t)) = \text{v.wait}()$ with $\mathbf{pv}(t)(v) = p$ and $\varphi(p)(t) = (\text{wait}_p^t, _)$ when $\varphi(p)(u) = (_, \text{sig}_p^u)$ and $\text{sig}_p^u \leq \text{wait}_p^t$.*

Intuitively, a task t is blocked by a task u if it cannot finish its `wait` command on some phaser because it is waiting for task u that did not issue enough `signal` commands on the same phaser.

Definition 2 (Deadlock). *$(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \varphi)$ is a deadlock configuration if each task of a non empty subset $\mathcal{U} \subseteq \mathcal{T}$ is blocked by some task in \mathcal{U} .*

Theorem 1 (Deadlock-Freedom). *It is undecidable in general, even for programs with only three phasers and four tasks, to check for deadlock-freedom.*

The idea of the proof is to encode the reachability problem of any given 3-counters `reset-VAS` (vector addition system with reset arcs) as the reachability problem of a configuration with a cycle involving three phasers and three tasks (in addition to the main task). Indeed, reachability of configuration $(s_F, 0, 0, 0)$ (three counters with zero values at some control location s_F) is undecidable for `reset-VASs`. The idea then is to spawn three tasks and as many phasers. The value of each counter is captured with the difference between the signal and the wait of a pair of tasks on one phaser. Resets are encoded by asking a task to drop a phaser and exit and spawning a new task. The encoding ensures that a deadlock is reached exactly when the vector addition system reaches configuration $(s_F, 0, 0, 0)$. (See [11] for more details.)

V. SYMBOLIC VERIFICATION OF PHASER PROGRAMS

We briefly introduce `gap-order` constraints and use them to define a symbolic representation (hereafter constraints) that we use in Section VI for checking reachability.

A. Gap-order constraints and graphs [9], [10], [12], [13].

`Gap-order` constraints can be regarded as a particular case of the octagons or the *unit two variables per inequality* (`utvpi`) constraints. Assume in this section that x and y are integer variables and that k is an integer constant. We use X and Y to mean finite sets of integer variables. A valuation val is a total function $X \rightarrow \mathbb{Z}$. Valuations are implicitly extended to preserve constants (i.e. $val(k) = k$ for any $k \in \mathbb{Z}$). A *gap-order clause* δ over X is an inequality of the form $a - b \geq k$ where $a, b \in X \cup \{0\}$. A *gap-order constraint* Δ over X is a finite conjunction of `gap-order` clauses over the same set X . Observe that $(x = y + 2 \wedge y \leq 5)$ is essentially a `gap-order` constraint because it can be equivalently rewritten as

$$\begin{array}{c}
\frac{\text{hd}(\mathbf{pc}(t)) = v := \text{newPhaser}() \wedge p \notin \mathcal{P} \wedge \\
p' = \mathcal{P} \cup \{p\} \wedge \mathbf{pv}' = \mathbf{pv}[t \leftarrow \mathbf{pv}(t)[v \leftarrow p]] \\
\wedge \varphi' = \varphi[p \leftarrow \{t \mapsto (0, 0)\}]}{(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \varphi) \xrightarrow{t} (\mathcal{T}, \mathcal{P}', \mathbf{bv}, \mathbf{pc}[t \leftarrow \text{tl}(\mathbf{pc}(t))], \mathbf{pv}', \varphi')} \quad (\text{newPhaser})} \\
\frac{\text{hd}(\mathbf{pc}(t)) = \text{assert}(\text{cond}) \wedge \\
\mathbf{bv}(\text{cond}) = \text{true}}{(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \varphi) \xrightarrow{t} (\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}[t \leftarrow \text{tl}(\mathbf{pc}(t))], \mathbf{pv}, \varphi)} \quad \left(\begin{array}{c} \text{assert.} \\ \text{ok} \end{array} \right) \\
\frac{\text{hd}(\mathbf{pc}(t)) = \text{v.signal}() \wedge \\
\mathbf{pv}(t)(v) = p \wedge \varphi(p)(t) = (\text{wait}_p^t, \text{sig}_p^t) \wedge \\
\varphi' = \varphi[p \leftarrow \varphi(p)[t \leftarrow (\text{wait}_p^t, 1 + \text{sig}_p^t)]]}{(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \varphi) \xrightarrow{t} (\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}[t \leftarrow \text{tl}(\mathbf{pc}(t))], \varphi')} \quad (\text{signal}) \\
\frac{\text{hd}(\mathbf{pc}(t)) = \text{v.drop}() \wedge \mathbf{pv}(t)(v) = p \wedge \\
\varphi(p)(t) \downarrow \wedge \varphi' = \varphi[p \leftarrow \varphi(p)[t \leftarrow \uparrow]]}{(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \varphi) \xrightarrow{t} (\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}[t \leftarrow \text{tl}(\mathbf{pc}(t))], \mathbf{pv}, \varphi')} \quad (\text{drop}) \\
\frac{\text{hd}(\mathbf{pc}(t)) = \text{asynch}(\text{task}, v_1, \dots, v_k)\{s_1\} \wedge \text{paramOf}(\text{task}) = (w_1, \dots, w_k) \wedge \\
\text{for each } i : 1 \leq i \leq k. \mathbf{pv}(t)(v_i) = p_i \wedge \varphi(p_i) \downarrow \wedge \\
u \notin \mathcal{T} \wedge \mathbf{pv}' = \mathbf{pv}[u \leftarrow \{w_i \mapsto \mathbf{pv}(t)(v_i) \mid 1 \leq i \leq k\}] \wedge \mathbf{pc}' = \mathbf{pc}[u \leftarrow s_1] \wedge \\
\varphi' = \varphi[\{p_i \leftarrow \varphi(p_i)[u \leftarrow \varphi(p_i)(t)] \mid \mathbf{pv}(t)(v_i) = p_i \text{ for } p_i \in \mathcal{P} \text{ and } 1 \leq i \leq k\}]}{(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \varphi) \xrightarrow{t} (\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}'[t \leftarrow \text{tl}(\mathbf{pc}(t))], \mathbf{pv}, \varphi')} \quad (\text{asynch}) \\
\frac{\text{hd}(\mathbf{pc}(t)) = \text{v.wait}() \wedge \mathbf{pv}(t)(v) = p \wedge \varphi(p)(t) = (\text{wait}_p^t, \text{sig}_p^t) \wedge \\
\forall u \in \mathcal{T}. (\varphi(p)(u) = (\text{wait}_p^u, \text{sig}_p^u) \Rightarrow \text{wait}_p^t < \text{sig}_p^u) \wedge \\
\varphi' = \varphi[p \leftarrow \varphi(p)[t \leftarrow (1 + \text{wait}_p^t, \text{sig}_p^t)]]}{(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \varphi) \xrightarrow{t} (\mathcal{T}, \mathcal{P}, \mathbf{pv}, \mathbf{pc}[t \leftarrow \text{tl}(\mathbf{pc}(t))], \varphi')} \quad (\text{wait}) \\
\frac{\text{hd}(\mathbf{pc}(t)) = \text{exit} \wedge \mathbf{pv}' = \mathbf{pv} \setminus \{t\} \wedge \mathbf{pc}' = \mathbf{pc} \setminus \{t\} \wedge \\
\varphi' = \varphi[\{p \leftarrow (\varphi(p) \setminus \{t\}) \mid p \in \mathcal{P}\}]}{(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \varphi) \xrightarrow{t} (\mathcal{T} \setminus \{t\}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}', \mathbf{pv}', \varphi')} \quad (\text{exit}) \\
\frac{\text{hd}(\mathbf{pc}(t)) = \text{assert}(\text{cond}) \wedge \\
\mathbf{bv}(\text{cond}) = \text{false}}{(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \varphi) \in \text{badConf}_{\text{assert}}^{(|\mathcal{T}|, |\mathcal{P}|)} \quad \left(\begin{array}{c} \text{assert.} \\ \text{fault} \end{array} \right) \\
\frac{\{t_0, \dots, t_n\} \subseteq \mathcal{T} \wedge \{p_0, \dots, p_n\} \subseteq \mathcal{P} \wedge \\
\forall i : 0 \leq i \leq n. \text{hd}(\mathbf{pc}(t_i)) = v_i.\text{wait}() \wedge \\
\mathbf{pv}(t_i)(w_i) = p_{(i+1)\%n} \wedge \mathbf{pv}(t_i)(v_i) = p_i \wedge \\
\text{wait}_{p_{(i+1)\%n}}^{t_i} \geq \text{sig}_{p_{(i+1)\%n}}^{t_{(i+1)\%n}}} \\
(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \varphi) \in \text{badConf}_{\text{deadlock}}^{(|\mathcal{T}|, |\mathcal{P}|)} \quad (\text{deadlock}) \\
\frac{\text{hd}(\mathbf{pc}(t)) = s \wedge (s = \text{v.drop}() \vee s = \text{v.signal}() \\
\vee s = \text{v.wait}() \vee s = \text{asynch}(\text{task}, \dots, v, \dots)) \\
\wedge (\mathbf{pv}(t)(v) \uparrow \vee \varphi(\mathbf{pv}(t)(v))(t) \uparrow)}{(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \varphi) \in \text{badConf}_{\text{runtime}}^{(|\mathcal{T}|, |\mathcal{P}|)} \quad \left(\begin{array}{c} \text{runtime} \\ \text{error} \end{array} \right) \\
\frac{\text{hd}(\mathbf{pc}(t)) = b := \text{cond} \wedge \text{hd}(\mathbf{pc}(u)) = s' \wedge t \neq u \wedge \\
\left(\begin{array}{l} s' = b := \text{cond}' \vee b \text{ appears in cond' and} \\ \left(\begin{array}{l} s' = \text{if}(\text{cond}') \{ \text{stmt}' \} \vee \\ s' = \text{while}(\text{cond}') \{ \text{stmt}' \} \vee \\ s' = \text{assert}(\text{cond}') \vee s' = b' = \text{cond}' \end{array} \right) \end{array} \right)}{(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \varphi) \in \text{badConf}_{\text{race}}^{(|\mathcal{T}|, |\mathcal{P}|)} \quad (\text{race})}
\end{array}$$

Fig. 3. Operational semantics of phaser statements.

the conjunction $(x - y \geq 2 \wedge y - x \geq -2 \wedge 0 - y \geq -5)$. Given a gap-order constraint Δ over X and a valuation $val : X \rightarrow \mathbb{Z}$, we write $val \models \Delta$ to mean that $val(a) - val(b) \geq k$ holds for each gap-order clause $\delta : a - b \geq k$ appearing in Δ . We let $Sat(\Delta)$ be the set $\{val : X \rightarrow \mathbb{Z} \mid val \models \Delta\}$.

A gap-order graph (or graph for short) \wp over X is a graph (V, E) with vertices $V = X \cup \{0\}$ where edges in E are of the form $a \xrightarrow{k} b$ with $a, b \in V$ and weight k in $\mathbb{Z} \cup \{-\infty, +\infty\}$. We let $\text{varsOf}(\wp) = X$. Given a gap-constraint Δ over X , we can build the graph $\text{graphOf}(\Delta)$ with vertices $X \cup \{0\}$ and where E only contains a representative $a \xrightarrow{k} b$ edge for each clause $a - b \geq k$ appearing in Δ . A valuation $val : X \rightarrow \mathbb{Z}$ satisfies a graph $\wp = (V, E)$ (written $val \models \wp$) iff $val(a) - val(b) \geq k$ for each $a \xrightarrow{k} b \in E$. We let $Sat(\wp)$ be the set $\{val : X \rightarrow \mathbb{Z} \mid val \models \wp\}$. Clearly, $Sat(\text{graphOf}(\Delta)) = Sat(\Delta)$. The closure $\text{clo}(\wp)$ of a graph $\wp = (V, E)$ is the unique complete graph with the same vertices V and where $a \xrightarrow{k'} b$ is an edge of $\text{clo}(\wp)$ iff $k' \in \mathbb{Z} \cup \{-\infty, +\infty\}$ is the least upper bound of all weight-sums for any path in \wp from a to b . Closure allows us to deduce $(0 - x \geq -7)$ from $(y - x \geq -2 \wedge 0 - y \geq -5)$. The result of the closure procedure is a special graph \wp_{false} denoting the graph without any satisfying

valuation each time a weight $k = +\infty$ is generated. The closure of a graph can be computed in polynomial time and we get $Sat(\text{clo}(\wp)) = Sat(\wp)$. We define the *degree* of a graph \wp (written $\text{degreeOf}(\wp)$) to be 0 if no edge in $\text{clo}(\wp)$ has a negative weight apart from $-\infty$. Otherwise, $\text{degreeOf}(\wp)$ is the largest natural $k \in \mathbb{N}$ such that there is an edge in $\text{clo}(\wp)$ with weight $-k$. For instance, the degree of the graph resulting from $(x - y \geq 2 \wedge y - x \geq -4)$ is 4. We systematically close all manipulated graphs and write $\mathcal{G}(X)$ for the set of closed graphs over X . Given a graph \wp , we write $\wp[x/y]$ to mean the graph obtained by replacing the vertex x by the vertex y . We abuse notation and write $\wp[\{x_i/y_i \mid i \in I\}]$, for pairwise different x_i elements to mean the simultaneous application of the individual substitutions. For a set of variables Y , we write $\wp \ominus Y$ to mean the graph obtained by removing the variables in Y from the vertices of \wp . Given two closed graphs \wp and \wp' over the same X , we write $\wp \sqsubseteq_{\mathcal{G}} \wp'$ to mean that each directed edge in \wp is labeled with a larger weight in \wp' . As a result, $Sat(\wp') \subseteq Sat(\wp)$. Finally, we write $\wp \otimes \wp'$ to mean the closure of the graph obtained with merging the two sets of vertices and edges. As a result, $Sat(\wp \otimes \wp') = Sat(\wp) \cap Sat(\wp')$.

B. Constraints as a symbolic representation.

A constraint ϕ is a tuple $(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \gamma)$ where the only difference with the definition of a configuration $(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \varphi)$ is the adoption of a gap-order constraint γ instead of φ . More specifically, $\gamma : \mathcal{P} \rightarrow \cup_{u \subseteq \mathcal{T}} \mathcal{G}(\cup_{t \in u} \{\omega_p^t, \sigma_p^t\})$ is a total mapping that associates a gap-order graph to each phaser $p \in \mathcal{P}$. Intuitively, we use variables ω_p^t and σ_p^t to constrain in graph $\gamma(p)$ possible values of both wait ($wait_p^t$) and signal (sig_p^t) phases of each task t registered to phaser p . As a result, we can check if task t is registered to phaser p according to graph $\varphi = \gamma(p)$ by checking if $\{\omega_p^t, \sigma_p^t\} \subseteq \text{varsOf}(\varphi)$. We will write $\text{Reg}(p, \varphi)$ to mean the set of tasks $\{t \mid \{\omega_p^t, \sigma_p^t\} \subseteq \text{varsOf}(\varphi)\}$. We also write $\text{isReg}(t, p, \varphi)$ for the predicate $t \in \text{Reg}(p, \varphi)$. Observe that the language semantics impose that, for each phaser p and for any pair t, u of tasks in $\text{Reg}(p, \varphi)$, the predicate $0 \leq wait_p^t \leq sig_p^u$ is an invariant. For this reason, we always safely strengthen, in any obtained $\gamma(p) = \varphi$, weights k in $\sigma_p^t \xrightarrow{k} \omega_p^u$, $\sigma_p^t \xrightarrow{k} 0$ and $\omega_p^t \xrightarrow{k} 0$ with $\max(k, 0)$. The following definition helps us characterize configurations for which our procedure terminates.

Definition 3 (degree and freeness of constraints). *A constraint $(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \gamma)$ has as degree the largest degree among all its graphs $\gamma(p)$ for $p \in \mathcal{P}$ if \mathcal{P} is not empty and 0 otherwise. Furthermore, a constraint is said to be “free” if, for any $p \in \mathcal{P}$, the only edges in $\gamma(p)$ with weights different from $-\infty$ are edges of the forms (i) $\sigma_p^t \xrightarrow{k(\sigma_p^t, \omega_p^u)} \omega_p^u$, (ii) $\sigma_p^t \xrightarrow{k(\sigma_p^t)} 0$, or (iii) $\omega_p^t \xrightarrow{k(\omega_p^t)} 0$ for some $t, u \in \text{Reg}(p, \gamma(p))$ and $k(\sigma_p^t, \omega_p^u), k(\sigma_p^t), k(\omega_p^t) \in \mathbb{N}$*

Free constraints are only allowed to impose, for the same phaser, non-negative lower bounds on differences between signals and waits, between signals and 0, and between waits and 0. Like degree-0-constraints, free constraints are not allowed to put a positive upper bound on how much a signal is larger than a wait. Unlike degree-0-constraints, they are not allowed to put bounds on the differences among signal values, or among wait values. For instance a free constraint cannot impose $\sigma_p^t - \sigma_p^u = 0$ while a degree-0-constraint can. Intuitively, freeness does not oblige our verification procedure to maintain exact differences when firing “signal” or “wait” instructions, jeopardizing termination. This will be stated in Section VI.

C. Denotations of constraints.

Given a configuration $c = (\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \varphi)$ and a constraint $\phi = (\mathcal{T}', \mathcal{P}', \mathbf{bv}', \mathbf{pc}', \mathbf{pv}', \gamma')$, we say that c satisfies ϕ , and write $c \models \phi$, if c satisfies (up to a renaming of the tasks and the phasers) conditions imposed by ϕ . More concretely, $c \models \phi$ if $\mathbf{bv} = \mathbf{bv}'$ and there are bijections $\tau : \mathcal{T} \rightarrow \mathcal{T}'$ and $\pi : \mathcal{P} \rightarrow \mathcal{P}'$ such that: (i) $\mathbf{pc}(t) = \mathbf{pc}'(\tau(t))$ for each $t \in \mathcal{T}$; and (ii) $\pi(\mathbf{pv}(t)(v)) = \mathbf{pv}'(\tau(t))(v)$ for each $t \in \mathcal{T}$ and $v \in V$; and (iii) the renaming of tasks and phasers in φ wrt. τ and π satisfies γ , i.e., (iii.a) for each $t \in \mathcal{T}$

and each $p \in \mathcal{P}$, $\varphi(p)(t) \downarrow$ iff $\text{isReg}(\tau(t), \pi(p), \gamma(\pi(p)))$, and (iii.b) for each $p' \in \mathcal{P}'$, $\wp(\bigwedge_{t' \in \text{Reg}(p', \gamma(p'))} (\omega_{p'}^{t'}, \sigma_{p'}^{t'}) = \varphi(\pi^{-1}(p'))(\tau^{-1}(t')))) \models \gamma(p')$. We let $\llbracket \phi \rrbracket$ denote $\{c \mid c \models \phi\}$. Intuitively, $\llbracket (\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \gamma) \rrbracket$ contains all configurations c with the same number of tasks and phasers and such that there are renamings of tasks and phasers that preserve in c the correspondence between \mathbf{pc} , \mathbf{pv} and γ . We write $\llbracket \Phi \rrbracket$, for a set Φ of constraints, to mean the union $\cup_{\phi \in \Phi} \llbracket \phi \rrbracket$. Given a program (B, V, T) , we can exactly characterize with a finite set of constraints all configurations involving n tasks and p phasers and satisfying the premises of rules (runtime error), (assert. fault), (race) and (deadlock) from Fig.(3).

Lemma 1 (Characterizing badness). *Given a program (B, V, T) and natural numbers (n, p) , we can exhibit finite sets of constraints $\text{badCstrs}_{\text{race}}^{(n, p)}$, $\text{badCstrs}_{\text{assert}}^{(n, p)}$, $\text{badCstrs}_{\text{runtime}}^{(n, p)}$ and $\text{badCstrs}_{\text{deadlock}}^{(n, p)}$ such that:*

$$\begin{aligned} \text{badConfs}_{\text{race}}^{(n, p)} &= \llbracket \text{badCstrs}_{\text{race}}^{(n, p)} \rrbracket \\ \text{badConfs}_{\text{assert}}^{(n, p)} &= \llbracket \text{badCstrs}_{\text{assert}}^{(n, p)} \rrbracket \\ \text{badConfs}_{\text{runtime}}^{(n, p)} &= \llbracket \text{badCstrs}_{\text{runtime}}^{(n, p)} \rrbracket \\ \text{badConfs}_{\text{deadlock}}^{(n, p)} &= \llbracket \text{badCstrs}_{\text{deadlock}}^{(n, p)} \rrbracket \end{aligned}$$

In addition, we can choose the constraints in $\text{badCstrs}_{\text{deadlock}}^{(n, p)}$ to be of degree 0 while those in $\text{badCstrs}_{\text{race}}^{(n, p)}$, $\text{badCstrs}_{\text{assert}}^{(n, p)}$ or in $\text{badCstrs}_{\text{runtime}}^{(n, p)}$ to be free.

D. Entailment.

We say that a constraint $\phi = (\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \gamma)$ is weaker than a constraint $\phi' = (\mathcal{T}', \mathcal{P}', \mathbf{bv}', \mathbf{pc}', \mathbf{pv}', \gamma')$, written $\phi \sqsubseteq \phi'$, to mean the following. First, the two constraints have the same number of phasers and tasks, agree on the values of the Boolean variables and, up to renamings, on the values of the phaser variables and on which tasks are registered to which phasers. Second, the constraints on the wait and signal values are stronger in ϕ' than in ϕ . More formally, $\phi \sqsubseteq \phi'$ if $\mathbf{bv} = \mathbf{bv}'$ and there are bijections $\tau : \mathcal{T} \rightarrow \mathcal{T}'$ and $\pi : \mathcal{P} \rightarrow \mathcal{P}'$ s.t. for each $t \in \mathcal{T}$ and $p \in \mathcal{P}$ the following four conditions hold: (i) $\mathbf{pc}(t) = \mathbf{pc}'(\tau(t))$; and (ii) $\pi(\mathbf{pv}(t)(v)) = \mathbf{pv}'(\tau(t))(v)$; and (iii) $\pi(\text{Reg}(p, \gamma(p))) = \text{Reg}(\pi(p), \gamma'(\pi(p)))$; and (iv) $\gamma(p) \sqsubseteq_{\mathcal{G}} \gamma'(\pi(p)) \left[\left\{ \omega_{\pi(p)}^{\tau(t)} / \omega_p^t, \sigma_{\pi(p)}^{\tau(t)} / \sigma_p^t \mid t \in \text{Reg}(p, \gamma(p)) \right\} \right]$. Clearly, $\phi \sqsubseteq \phi'$ implies $\llbracket \phi' \rrbracket \subseteq \llbracket \phi \rrbracket$. We say that \sqsubseteq is sound.

We can show that \sqsubseteq is a well-quasi-order¹ over constraints of bounded degrees and involving fixed numbers of tasks and phasers since $\sqsubseteq_{\mathcal{G}}$ is itself a well-quasi-ordering over graphs of bounded degrees over a finite set of variables ([9], [12]).

Lemma 2 (WQO). *Given $k, n, p \in \mathbb{N}$, the entailment relation \sqsubseteq over the set of constraints of degree k involving at most n tasks and p phasers is a well-quasi-order.*

¹A reflexive and transitive binary relation \preceq is a well-quasi-order over a set A if there is no infinite sequence a_0, a_1, \dots of A elements s.t. $a_i \not\preceq a_j$ for all $i < j$.

$$\begin{array}{c}
\frac{\begin{array}{l} \mathbf{pc}' = \mathbf{pc}[t \leftarrow v := \text{newPhaser}(); \mathbf{pc}(t)] \wedge \\ \mathbf{pv}(t)(v) = p \wedge \mathbf{pv}' = \mathbf{pv}[t \leftarrow \mathbf{pv}(t)[v \leftarrow q]] \wedge \\ \{\omega_p^t \mapsto 0, \sigma_p^t \mapsto 0\} \models \gamma(p) \wedge \gamma' = \gamma \setminus \{p\} \wedge \\ (\text{isReg}(u, p, \gamma(p)) \implies u = t) \end{array}}{(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \gamma) \xrightarrow{t} (\mathcal{T}, \mathcal{P} \setminus \{p\}, \mathbf{bv}, \mathbf{pc}', \mathbf{pv}', \gamma')} \quad (\text{newPhaser I})} \\
\frac{\begin{array}{l} \mathbf{pc}' = \mathbf{pc}[t \leftarrow v := \text{newPhaser}(); \mathbf{pc}(t)] \wedge \\ \mathbf{pv}(t)(v) = p \wedge \mathbf{pv}' = \mathbf{pv}[t \leftarrow \mathbf{pv}(t)[v \leftarrow q]] \wedge \\ \{\omega_p^t \mapsto 0, \sigma_p^t \mapsto 0\} \models \gamma(p) \wedge \gamma' = \gamma \setminus \{p\} \wedge \\ (\text{isReg}(u, p, \gamma(p)) \implies u = t) \end{array}}{(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \gamma) \xrightarrow{t} (\mathcal{T}, \mathcal{P} \setminus \{p\}, \mathbf{bv}, \mathbf{pc}', \mathbf{pv}', \gamma')} \quad (\text{newPhaser II})} \\
\frac{\begin{array}{l} \mathbf{pc}' = \mathbf{pc}[t \leftarrow v.\text{signal}(); \mathbf{pc}(t)] \wedge \mathbf{pv}(t)(v) = p \wedge \text{isReg}(t, p, \gamma(p)) \wedge \wp = (\gamma(p) \otimes \text{graphOf}(\wedge_{u \in \text{Reg}(p, \gamma(p))} (\sigma_p^u > \omega_p^u \geq 0))) \\ \wedge \text{isSat}(\wp) \wedge \gamma' = \gamma[p \leftarrow ((\wp[\sigma_p^t/\sigma] \otimes \text{graphOf}(\sigma_p^t = \sigma - 1)) \ominus \{\sigma\})] \end{array}}{(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \gamma) \xrightarrow{t} (\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}', \mathbf{pv}, \gamma')} \quad (\text{signal})} \\
\frac{\begin{array}{l} \mathbf{pc}' = \mathbf{pc}[t \leftarrow v.\text{wait}(); \mathbf{pc}(t)] \wedge \mathbf{pv}(t)(v) = p \wedge \text{isReg}(t, p, \gamma(p)) \wedge \wp = (\gamma(p) \otimes \text{graphOf}(\wedge_{\{u \in \text{Reg}(p, \gamma(p))\}} (\sigma_p^u \geq \omega_p^t > 0))) \\ \wedge \text{isSat}(\wp) \wedge \gamma' = \gamma[p \leftarrow ((\wp[\omega_p^t/\omega] \otimes \text{graphOf}(\omega_p^t = \omega - 1)) \ominus \{\omega\})] \end{array}}{(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \gamma) \xrightarrow{t} (\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}', \mathbf{pv}, \gamma')} \quad (\text{wait})} \\
\frac{\begin{array}{l} \mathbf{pc}' = \mathbf{pc}[t \leftarrow v.\text{drop}(); \mathbf{pc}(t)] \wedge \mathbf{pv}(t)(v) = p \wedge \neg \text{isReg}(t, p, \gamma(p)) \wedge \\ \gamma' = \gamma[p \leftarrow (\gamma(p) \otimes \text{graphOf}((\sigma_p^t \geq \omega_p^t \geq 0) \wedge_{u \in \text{Reg}(p, \gamma(p))} (\sigma_p^u \geq \omega_p^t \geq 0) \wedge (\sigma_p^t \geq \omega_p^u \geq 0)))] \end{array}}{(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \gamma) \xrightarrow{t} (\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}', \mathbf{pv}, \gamma')} \quad (\text{drop})} \\
\frac{\begin{array}{l} \mathbf{pc}' = \mathbf{pc}[t \leftarrow \text{asynch}(\text{task}, v_1, \dots, v_k)\{s_1\}; \mathbf{pc}(t)] \wedge \text{paramOf}(\text{task}) = (w_1, \dots, w_k) \wedge u \in \mathcal{T} \setminus \{t\} \wedge \\ \mathbf{pv}' = \mathbf{pv} \setminus \{u\} \wedge \mathbf{pc}(u) = s_1 \wedge \forall i : 1 \leq i \leq k. \mathbf{pv}(t)(v_i) = \mathbf{pv}(u)(w_i) = p_i \wedge (\text{isReg}(t, p_i, \gamma(p_i)) \Leftrightarrow \text{isReg}(u, p_i, \gamma(p_i))) \wedge \\ \wp_i = (\gamma(p_i) \otimes \text{graphOf}(\omega_{p_i}^t = \omega_{p_i}^u \wedge \sigma_{p_i}^t = \sigma_{p_i}^u)) \wedge \text{isSat}(\wp_i) \wedge \gamma_0 = \gamma \wedge \gamma_i = \gamma_{i-1}[p_i \leftarrow (\wp_i \ominus \{\omega_{p_i}^u, \sigma_{p_i}^u\})] \end{array}}{(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \gamma) \xrightarrow{t} (\mathcal{T} \setminus \{u\}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}' \setminus \{u\}, \mathbf{pv}', \gamma_n)} \quad (\text{asynch})} \\
\frac{\begin{array}{l} t \notin \mathcal{T} \wedge \mathbf{pc}' = \mathbf{pc}[t \leftarrow \text{exit}] \wedge f \in \text{Pfn}(\mathcal{V}, \mathcal{P}) \wedge \mathbf{pv}' = \mathbf{pv} \uplus \{t \mapsto f\} \wedge \\ \mathcal{Q} \subseteq \mathcal{P} \wedge \gamma' = \gamma[\{p \leftarrow \wp(p) \otimes \text{graphOf}((\sigma_p^t \geq \omega_p^t \geq 0) \wedge_{u \in \text{Reg}(p, \gamma(p))} (\sigma_p^u \geq \omega_p^t \geq 0) \wedge (\sigma_p^t \geq \omega_p^u \geq 0)) \mid p \in \mathcal{Q}\}] \end{array}}{(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \gamma) \xrightarrow{t} (\mathcal{T} \cup \{t\}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}', \mathbf{pv}', \gamma')} \quad (\text{exit})}
\end{array}$$

Fig. 4. Derivation rules for computing $\text{pre}(t, \phi)$ for phaser statements as union of all $\{\phi' \mid \phi \xrightarrow{t} \phi'\}$ with $\phi = (\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \gamma)$ and $t \in \mathcal{T}$. Derivations for other program statements are straightforward.

VI. VERIFICATION PROCEDURE

Input: A program $\text{prg} = (\mathcal{B}, \mathcal{V}, \mathcal{T})$, a set Φ_{bad} of pairwise \sqsubseteq -incomparable constraints, maximum upper bounds t^\bullet and p^\bullet (in $\mathbb{N} \cup \{+\infty\}$) on coexisting tasks and phasers.
Output: A symbolic run to Φ_{bad} or the value *unreachable*

- 1 Initialize both *Working* and *Visited* to $\{(\phi, \phi) \mid \phi \in \Phi_{bad}\}$;
- 2 **while** there exists $(\phi, \tau) \in \text{Working}$ **do**
- 3 remove (ϕ, τ) from *Working*;
- 4 let $(\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \gamma) = \phi$;
- 5 **if** $|\mathcal{T}| > t^\bullet$ **or** $|\mathcal{P}| > p^\bullet$ **then** continue;
- 6 **if** $c_{init} \models \phi$ **then** return τ ;
- 7 **foreach** $t \in \mathcal{T}$ **do**
- 8 **foreach** $\phi' \in \text{pre}(t, \phi)$ **do**
- 9 **if** $\psi \not\sqsubseteq \phi'$ for all $(\psi, _)$ **in** *Visited* **then**
- 10 Remove from *Working* and *Visited* each $(\psi, _)$ for which $\phi' \sqsubseteq \psi$;
- 11 Add $(\phi', \phi' \cdot t \cdot \tau)$ to both *Working* and *Visited*;
- 12 **return** *unreachable* ;

Procedure $\text{check}(\text{prg}, \Phi_{bad}, t^\bullet, p^\bullet)$, a simple working list procedure for checking constraints reachability.

We discuss in the following the procedure *check* depicted above and assume a program prg and a set Φ_{bad} of constraints the reachability of which we want to check. Φ_{bad} can for example be any subset of $\text{badCstrs}_{\text{deadlock}}^{(n,p)}$ (degree 0) or of $\text{badCstrs}_{\text{assert}}^{(n,p)}$ (free) in case we want to check the possibility of a deadlock or of an assertion violation.

It is not difficult to show that $\llbracket \text{pre}(t, \phi) \rrbracket$ (obtained as described in Fig.(4)) coincides with $\{c' \mid c' \xrightarrow{t} c \text{ and } c \in \llbracket \phi \rrbracket\}$.

Using the soundness of \sqsubseteq , we can show by induction the partial correctness of the procedure $\text{check}(\text{prg}, \Phi_{bad}, +\infty, +\infty)$.

Lemma 3 (Partial correctness). *If $\text{check}(\text{prg}, \Phi_{bad}, +\infty, +\infty)$ returns *unreachable*, then $c_{init} \not\rightarrow^* \llbracket \Phi_{bad} \rrbracket$. If it returns a trace $\phi_n \cdot t_n \cdots t_1 \cdot \phi_1$ then there are c_n, \dots, c_1 with $c_n = c_{init}$, $c_1 \in \llbracket \Phi_{bad} \rrbracket$ and $c_i \xrightarrow{t_i} c_{i-1}$ for $i : 1 < i \leq n$.*

Theorem 2 (Free termination). *$\text{check}(\text{prg}, \Phi_{bad}, t^\bullet, p^\bullet)$ terminates for $t^\bullet, p^\bullet \in \mathbb{N}$ and free Φ_{bad} .*

Proof. Sketch. Freeness is preserved by the pre computation (Fig.(4)). Suppose the procedure does not terminate. The infinite sequence of constraints passing the test at line 9 of the procedure violates well-quasi-orderedness of \sqsubseteq over free constraints with fixed numbers of tasks and phasers. \square

In order to check reachability of arbitrary constraints, we may need to force termination. We do this by soundly bounding the degree of generated constraints using a relaxation ρ_k . The relaxation $\rho_k((\mathcal{T}, \mathcal{P}, \mathbf{bv}, \mathbf{pc}, \mathbf{pv}, \gamma))$ replaces, in each graph $\gamma(p)$, each weight k'' s.t. $k'' < -k$ with $-\infty$.

- 1 **foreach** $\phi'' \in \text{pre}(t, \phi)$ **do**
- 2 | Let $\phi' = \rho_k(\phi'')$;

Fig. 5. Systematic relaxation

Theorem 3 (Forced termination). *Procedure* $\text{check}(\text{prg}, \Phi_{\text{bad}}, t^*, p^*)$ for $t^*, p^* \in \mathbb{N}$, with line 8 replaced by the lines of Fig. (5), is sound and guaranteed to terminate.

Proof. Soundness is due to the validity of $\rho_k(\phi) \sqsubseteq \phi$ while the termination argument relies, similarly to Theorem (2), on well-quasi orderness of \sqsubseteq on the set of constraints with bounded degree and fixed numbers of tasks and phasers. \square

VII. EXPERIMENTAL RESULTS

We report on experiments with our open source prototype *hjVerify* (<https://gitlab.ida.liu.se/apv/hjVerify>) for the verification of phaser programs. We conducted experiments on 12 different programs (some of which are from [5]). We considered both deadlock and assertion reachability problems. For each property, we considered correct and buggy versions. This gave 48 different instances with 2 to 3 phasers and 2 to 4 tasks (except for the parameterized case). Our tool uses global phaser and task variables as in [5]. We have experimented with adapting the view abstraction technique [14] to verify phaser programs generating arbitrary many tasks, i.e., parameterized verification where the number of phasers is fixed. (see [11] for more details.) We report on two parameterized examples. Experiments were conducted on a 2.9GHz processor with 8GB of memory.

program	property	safe / buggy	times
01.Loopless	deadlock:	ok / trace	1s / 1s
	assertion:	ok / trace	1s / 1s
02.Iterative averaging	deadlock:	ok / trace	1s / 1s
	assertion:	ok / trace	1s / 1s
03.Ordered phasers	deadlock:	ok / trace	1s / 1s
	assertion:	ok / trace	13s / 1s
04.Conditional	deadlock:	ok / trace	2s / 1s
	assertion:	ok / trace	4s / 7s
05.Loop Synch.	deadlock:	ok / trace	178s / 145s
	assertion:	ok / trace	7s / 13s
06.Nested forks	deadlock:	ok / trace	2s / 1s
	assertion:	ok / trace	1s / 1s
07.Conditional membership	deadlock:	ok / trace	1s / 1s
	assertion:	ok / trace	12s / 3s
08.Producer-consumer	deadlock:	ok / trace	37s / 222s
	assertion:	ok / trace	79s / 34s
09.Parameterized loopless	deadlock:	ok / trace	20s / 1s
	assertion:	ok / trace	67s / 1s
10.Parameterized iterative-averaging	deadlock:	ok / trace	1s / 1s
	assertion:	ok / trace	1s / 1s
11.Running-2	deadlock:	ok / trace	5s / 1s
	assertion:	ok / trace	26s / 4s
12.Running-3	deadlock:	ok / trace	4318s / 128s
	assertion:	ok / trace	18631s / 54s

Our implemented procedure does not eagerly concretize all task states as described in the predecessor computation of Section V. Instead we collect conditions on the phases of the tasks that did not take any action yet and lazily concretize them. Reported times for checking deadlocks are the sums of the times required to check reachability for each cycle. The prototype is only a proof of concept. For instance, the example (12.Running-3) is a variant of (11.Running-2) where a task instance is spawned twice leading to two symmetrical tasks (out of four). This required up to three orders of magnitude more time to check. We believe partial order reduction techniques would help here. Other relevant heuristics would be to make

use of priority queues and to organize the minimal sets. All examples are available on the tool homepage.

VIII. CONCLUSION

We have proposed a gap-order based reachability analysis for phaser programs. We have showed our analysis to be exact and guaranteed to terminate when checking runtime, race and assertion errors. We have established the undecidability of deadlock verification and explained how to turn our analysis into a sound over-approximation. To the best of our knowledge, this is beyond the capabilities of current verification techniques which currently only target concrete inputs to phaser programs. We are currently working on tackling the parameterized case and have obtained preliminary encouraging results. Apart from improving the scalability of the tool and from using it in combination with predicate abstraction and abstract interpretation in order to analyze actual source code, we are investigating the applicability of the presented techniques for the verification of similar synchronization constructs.

REFERENCES

- [1] P. Charles, C. Grothoff, V. Saraswat, C. Donawa, A. Kielstra, K. Ebcioglu, C. von Praun, and V. Sarkar, "X10: An object-oriented approach to non-uniform cluster computing," *SIGPLAN Not.*, vol. 40, no. 10, pp. 519–538, Oct. 2005.
- [2] J. Shirako, D. M. Peixotto, V. Sarkar, and W. N. Scherer, "Phasers: a unified deadlock-free construct for collective and point-to-point synchronization," in *Proceedings of the 22nd annual international conference on Supercomputing*. ACM, 2008, pp. 277–288.
- [3] V. Cavé, J. Zhao, J. Shirako, and V. Sarkar, "Habanero-java: the new adventures of old x10," in *Proceedings of the 9th International Conference on Principles and Practice of Programming in Java*. ACM, 2011, pp. 51–61.
- [4] J. Shirako, D. M. Peixotto, V. Sarkar, and W. N. Scherer, "Phaser accumulators: A new reduction construct for dynamic parallelism," in *Parallel & Distributed Processing, 2009. IPDPS 2009. IEEE International Symposium on*. IEEE, 2009, pp. 1–12.
- [5] T. Cogumbreiro, R. Hu, F. Martins, and N. Yoshida, "Dynamic deadlock verification for general barrier synchronisation," in *20th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, ser. PPOPP 2015. New York, NY, USA: ACM, 2015, pp. 150–160.
- [6] D.-K. Le, W.-N. Chin, and Y.-M. Teo, "Verification of static and dynamic barrier synchronization using bounded permissions," in *Int. Conf. on Formal Engineering Methods*. Springer, 2013, pp. 231–248.
- [7] P. Anderson, B. Chase, and E. Mercer, "Jpf verification of habanero java programs," *SIGSOFT Softw. Eng. Notes*, vol. 39, no. 1, pp. 1–7, 2014.
- [8] K. Havelund and T. Pressburger, "Model checking java programs using java pathfinder," *International Journal on Software Tools for Technology Transfer (STTT)*, vol. 2, no. 4, pp. 366–381, 2000.
- [9] R. Mayr and P. Totzke, "Branching-time model checking gap-order constraint systems," *Fundamenta Informaticae*, vol. 143, no. 3–4, pp. 339–353, 2016.
- [10] L. Bozzelli and S. Pinchinat, "Verification of gap-order constraint abstractions of counter systems," *Theoretical Computer Science*, vol. 523, pp. 1 – 36, 2014.
- [11] G. Zeinab, R. Ahmed, E. Petru, and P. Zebo, "Safety verification of phaser programs," *CoRR*, vol. abs/1708.02801, 2017. [Online]. Available: <https://arxiv.org/abs/1708.02801>
- [12] P. Z. Revesz, "A closed-form evaluation for datalog queries with integer (gap)-order constraints," *Theoretical Computer Science*, vol. 116, no. 1, pp. 117–149, 1993.
- [13] S. Lahiri and M. Musuvathi, "An efficient decision procedure for utvpi constraints," in *Frontiers of Combining Systems (FroCos '05)*. Springer Verlag, May 2005.
- [14] P. A. Abdulla, F. Haziza, and L. Holík, "All for the price of few," in *International Workshop on Verification, Model Checking, and Abstract Interpretation*. Springer, 2013, pp. 476–495.