Parameterized Verification of Asynchronous Round-Based Distributed Algorithms via Reduction to Finite-Counter Systems

Nathalie Bertrand
Univ Rennes, Inria, CNRS, IRISA, France

Pranav Ghorpade
The University of Sydney, Australia

Sasha Rubin
The University of Sydney, Australia

Abstract

Traditional model-checking techniques typically verify distributed algorithms only for a fixed number of finite-state processes. Parameterized model checking generalizes this to any number of processes, while still typically assuming that each process is finite-state. In this work, we consider asynchronous round-based distributed algorithms in which each process is infinite-state since it can execute for an infinite number of rounds. We show that the parameterized verification problem for asynchronous round-based distributed algorithms is undecidable, already for simple specifications. Nevertheless, as our main contribution, we provide a reduction to LTL model checking over finite-counter systems and prove that it is sound and complete. This enables the use of off-the-shelf, mature symbolic model checkers for finite-counter systems. We demonstrate the practical applicability of this reduction by verifying safety and liveness properties of several asynchronous round-based consensus and leader-election algorithms using the nuXmv model checker.

2012 ACM Subject Classification Theory of computation → Verification by model checking; Theory of computation → Modal and temporal logics; Theory of computation → Distributed algorithms

Keywords and phrases parametrized verification, asynchronous round-based distributed algorithms, finite-counter systems, LTL model checking

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2026.20

Related Version *Full Version*: `https://arxiv.org/abs/2606.27867` [4]

Funding *Nathalie Bertrand*: Projects ANR BisoUS (ANR-22-CE48-0012) and ANR PaVeDyS (ANR-23-CE48-0005).
Pranav Ghorpade: Supported by University of Sydney Faculty of Engineering Research Scholarship.

1 Introduction

Asynchronous distributed algorithms consist of a parameterized number of identical processes that asynchronously exchange messages with other processes and update their local states. A subclass of these is that of asynchronous round-based distributed algorithms (ARDAs), in which each process progresses through an unbounded sequence of rounds and each message is tagged with the sender's current round number. Typical examples include Ben-Or's randomized consensus algorithms [3], Bracha's consensus algorithm [9], and Raft's leader-election algorithm [35]. In the absence of synchronization, processes in such algorithms may drift across rounds, so each process must store its current round index as part of its local state. Since this index is unbounded, each process has infinitely many possible states. Consequently, the global system exhibits two orthogonal sources of unboundedness: the parameterized number of processes and the unbounded round index of each process.

Parameterized verification of distributed algorithms refers to checking their correctness for all numbers n of processes, provided that a resilience condition relating n and the maximum number t of faulty processes is satisfied (e.g., $3t < n$). Parameterized model checking provides techniques that cope with all admissible valuations of the parameters n and t, namely those satisfying the resilience condition. However, it typically assumes that processes are finite-state (see, for instance, the book [8] or the recent survey [26]) and is therefore not directly applicable to asynchronous round-based distributed algorithms. Moreover, even for fixed parameter values, the resulting system is infinite-state because of the round indices and cannot be directly handled by finite-state model checkers such as SPIN, NuSMV, or TLA+.

In this work, we introduce a modeling formalism for a broad class of ARDAs, together with a specification logic designed to express their standard correctness properties. For this formalism, we show that the parameterized verification problem is undecidable, whereas the fixed-instance verification problem is decidable. To address the parameterized case despite this undecidability, we give a sound and complete reduction to LTL model checking over finite-counter systems. Although this target problem is itself undecidable, it is well studied, and sound, necessarily incomplete, symbolic model-checking techniques for it have been developed and implemented, for example in nuXmv [10]. Thus, our reduction enables the use of existing symbolic model checkers for parameterized verification of ARDAs, while also potentially benefiting from future advances in such tools. To assess the practical applicability of the reduction, we apply it to four case studies: (i) Ben-Or's consensus with crash faults [3, Algo A], (ii) Ben-Or's consensus with Byzantine faults [3, Algo B], (iii) Bracha's consensus with Byzantine faults [9, Fig. 4], and (iv) Raft's leader-election [35, Sec. 5.2]. Using nuXmv as the back end, we verify safety and liveness properties, namely agreement, validity, termination (under some condition), and leader uniqueness. These case studies suggest that existing verification engines can efficiently handle the resulting finite-counter systems.

Brief Overview. § 2.1 introduces templates for ARDAs. A template specifies the symbolic parameters of the algorithm together with the behavior of a single process. For each admissible parameter valuation, it induces a round-based distributed transition system (RDTS) capturing the behavior of the resulting interacting processes. The parameterized semantics (parameterized RDTS) is obtained as the disjoint union of these fixed-instance RDTS over all admissible parameter valuations. § 2.2 identifies a class of *action-based* temporal properties, called *history state-count properties*, and introduces a tailored action-based logic, *history state-count logic* (HSCL). This logic is deliberately restricted so that it can be translated into state-based LTL, while remaining expressive enough to capture standard correctness conditions for consensus and leader election. § 2.3 gives a reduction from the non-halting problem for two-counter machines to the verification problem of parameterized RDTS against HSCL formulas, thereby establishing undecidability. § 3.1 gives four sound and complete transformations that reduce the parameterized RDTS to a finite-counter system while preserving history state-count properties. § 3.2 then gives two further sound and complete transformations that reduce verification of parameterized RDTS against HSCL formulas to LTL model checking over finite-counter systems. When specialized to fixed-instance semantics RDTS, our reduction yields a *one*-counter system for history state-count properties and reduces verification of RDTS against HSCL formulas to LTL model checking over *finite-state* systems. This proves that the fixed-instance verification problem is decidable. § 4 presents the case studies, and § 5 concludes the paper.

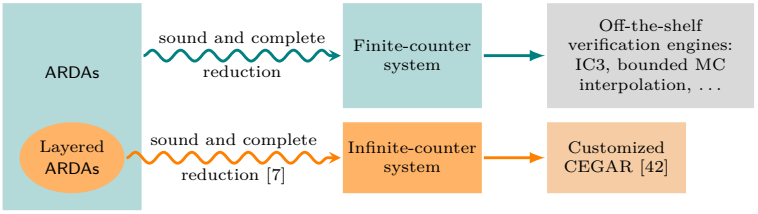

Figure 1 Our reduction improves over the verification framework for layered ARDAs by [7, 42] in both scope and the ability to reuse mature verification tools.

Related Work. Round-based distributed algorithms belong to the broader class of fault-tolerant distributed algorithms, which also includes algorithms in which processes are finite-state, such as broadcast algorithms. Threshold automata [28, 27] provide a successful framework for parameterized verification of such finite-state algorithms by abstracting them to finite-counter systems that can be analyzed by symbolic model checkers. This line of work was later extended to synchronous round-based distributed algorithms through synchronous threshold automata [40]. For asynchronous round-based distributed algorithms (ARDAs), a compositional approach was proposed in [6]. It reduces the verification to proof obligations for individual rounds, which are handled within the threshold-automata framework. A limitation is that the target property is not established directly: one must first identify suitable round invariants and then argue separately that these invariants imply the desired property.

Layered threshold automata [7], together with the PyLTA prototype [42], were later proposed as a more direct framework for verifying ARDAs. However, this framework is restricted to algorithms that admit a *layered* structure, thereby excluding, for instance, most leader-election algorithms. More importantly, the corresponding verification procedure does not retain the main tooling advantage of standard threshold automata. It first gives a sound and complete abstraction to counter systems with *infinitely* many counters, and then applies a further abstraction to finite-state systems that is *incomplete*. As a result, the approach relies on a bespoke CEGAR loop rather than on off-the-shelf symbolic model checking. In contrast, our work targets a broader class of ARDAs, including leader-election algorithms, and gives a sound and complete reduction to finite-counter systems. This enables the use of mature general-purpose symbolic model checkers, as illustrated in Fig. 1.

Traditionally, parameterized verification of ARDAs has largely relied on human-guided methods. Interactive theorem provers such as Rocq, Isabelle/HOL, and TLA+ have been used to formalize distributed algorithms and establish their correctness through mechanized, expert-driven proofs [43, 11, 5]. Deductive verification frameworks such as Ivy [36] encode distributed algorithms in first-order logic and can prove correctness automatically once suitable inductive invariants are provided. However, discovering invariants remains challenging, motivating a large body of work on automatic inductive-invariant inference [22, 45, 25, 37].

For fixed-instance verification, prior work [16, 34] applies finite-state model-checking tools to prove correctness for a fixed number of processes executing a bounded number of rounds. In contrast, our setting imposes no bound on the number of rounds. The work of [13] also addresses fixed-instance verification of round-based consensus algorithms in the Heard-Of model [14], using TLC to verify fixed instances of algorithms such as OneThirdRule and UniformVoting. However, their reduction is tailored to TLA+/TLC, whose specifications are written in the Temporal Logic of Actions, whereas our reduction targets standard LTL model checking. Moreover, our framework handles a broader class of algorithms, including Byzantine fault-tolerant consensus algorithms and leader-election algorithms.

Parameters: n, t , with $n > 2t$
Input: $v_0 \in \{0, 1\}$
Initialize: $v := v_0$, $d :=$ `Nil`, $r := 0$
At round r:
1. `Broadcast` (v, r)
2. `Wait` until $\geq (n - t)$ messages received with round tag r
3. `Compute` if some value v' appears in $> n/2$ messages: $d := v'$; `halt`, else $v := 0$ or $v := 1$
4. `Progress` $r := r + 1$

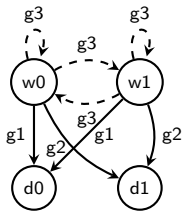

$\mathsf{P} = \{n, t\}$ $\quad \mathsf{rc} = n > 2t$
$\mathcal{L} = \{\mathsf{w0}, \mathsf{w1}, \mathsf{d0}, \mathsf{d1}\}$ $\mathcal{M} = \{\mathsf{m0}, \mathsf{m1}\}$

$\mathsf{g1} = \texttt{Quorum} \wedge \mathsf{m0} > n/2$
$\mathsf{g2} = \texttt{Quorum} \wedge \mathsf{m1} > n/2$
$\mathsf{g3} = \texttt{Quorum} \wedge \mathsf{m0} \leq n/2 \wedge \mathsf{m1} \leq n/2$
where $\texttt{Quorum} = \mathsf{m0} + \mathsf{m1} \geq n - t$

——— : type $= 0$
- - - : type $= 1$

Bcast: $\mathsf{w0} \mapsto \mathsf{m0}, \mathsf{w1} \mapsto \mathsf{m1}$,
$\mathsf{d0} \mapsto \perp, \mathsf{d1} \mapsto \perp$

Figure 2 Left: Pseudo-code of a simplified round-based voting algorithm. **Right**: Corresponding process template representation. In the guards, the variable mi denotes the number of messages of type mi received by the process.

2 Modeling Asynchronous Round-Based Distributed Algorithms

This section introduces the process templates that serve as our formal model, together with the specification language and verification problems studied in the paper.

2.1 Round-based Process Templates

The syntax of our model is motivated by the typical structure of ARDAs; see Fig. 2 (left) for an example. These algorithms are parameterized; processes communicate via broadcast, with choices governed by threshold guards; and the overall behavior of each process is composed of identical, bounded, communication-closed rounds. We now expand on these three aspects. (i) ARDAs are parameterized by integer variables (typically n for number of processes and t for an upper-bound on the number of faulty processes) and are designed to work when these variables satisfy a *resilience condition*, typically expressed as a linear arithmetic formula (e.g., $t < n/2$). (ii) In ARDAs, communication is typically by *broadcast*. Moreover, since processes may be faulty, an ARDA cannot rely on receiving messages from specific processes. Accordingly, process updates are governed by *threshold guards*, which depend only on the number of received messages of each type, rather than on the identities of the senders; see step 3 in Fig. 2 (left). (iii) In ARDAs, the behavior of a process in each round follows the same high-level structure: a finite sequence of broadcast, receive, and compute steps, potentially followed by a progress step to a future round. As a result, rounds are *structurally identical*, and the control-flow within each round is *bounded*. Messages from different rounds are distinguished using *round tags*, which enforce *communication closure* within rounds: messages tagged with round r are used exclusively during computations in round r.

Exploiting the identical structure of rounds in ARDAs, our round-based process template provides a round-independent representation capturing the behavior of a *single process* within a *single round*, together with the rules governing round progress. Similar representations for ARDAs have appeared in prior work, though in more restrictive forms [7] or in a less direct form requiring additional modeling components (e.g., auxiliary locations) [6].

Before proceeding to the formal definition of templates, we illustrate it on the algorithm in Fig. 2 (left) for which Fig. 2 (right) gives a process template. The set of parameters is $\mathsf{P} = \{n, t\}$, with admissible valuations satisfying the resilience condition $\mathsf{rc} = t < n/2$. The locations $\mathcal{L} = \{\mathsf{w0}, \mathsf{w1}, \mathsf{d0}, \mathsf{d1}\}$ represent a process's position within a round: either in a *wait location* (w0 or w1, depending on its current vote) or in a *decide location* (d0 or d1, depending on its decision). These locations correspond to the *wait* and *compute* steps in the pseudo-code. The *broadcast* and *progress* steps are treated implicitly: after updating

its vote, a process immediately progresses to the next round, broadcasts its new vote, and re-enters a wait location. The message types $\mathcal{M} = \{\mathsf{m0}, \mathsf{m1}\}$ represent the possible votes a process may broadcast. The transition rules, represented by edges, are labeled with a threshold guard on received messages. For example, guard $\mathsf{g1}$ corresponds to the conjunction "at least $n - t$ messages are received" and "more than $n/2$ are of type $\mathsf{m0}$". Transitions are also equipped with a type that specifies how the round id evolves: type 0 indicates that the process remains in its current round, while in general type k indicates that it progresses by k rounds. Our example only has type 0 and type 1 transitions represented by solid and dashed edges respectively. More generally, round-tag increases of up to $b \in \mathbb{N}$ are allowed, where b is referred to as the jump bound. A broadcast function Bcast maps each location to the message sent while entering it: here, for $i = 1, 2$, wi maps to mi, and di maps to $\perp$ (i.e., no broadcast). As exhibited by this example, we assume that the subgraph formed by the solid edges (i.e., those with $\mathsf{type} = 0$) is acyclic.

▶ **Definition 1** (Round-based process templates). *Let $b \in \mathbb{N}$ be a* jump bound. *A* round-based process template*, or simply a* template*, is a tuple $\langle \mathsf{P}, \mathsf{rc}, \mathcal{L}, \mathcal{I}, \mathcal{M}, \mathsf{Bcast}, \mathsf{Rules} \rangle$ where:*

- P *is a finite set of* symbolic parameters*, including n, the number of processes;*
- $\mathsf{rc} \in \mathsf{LinArith}(\mathsf{P})$ *is a* resilience condition*;*
- $\mathcal{L}$ *is a finite set of* locations*, and $\mathcal{I} \subseteq \mathcal{L}$ is the set of* initial locations*;*
- $\mathcal{M}$ *is a finite set of* message types*;*
- $\mathsf{Bcast} : \mathcal{L} \to \mathcal{M} \cup \{\perp\}$ *is the* broadcast function*;*
- Rules *is a finite set of* transition rules*, where transition rule ρ is a record of the form $\rho = \{\mathsf{from} \in \mathcal{L},\ \mathsf{to} \in \mathcal{L},\ \mathsf{type} \in [\![0, b]\!],\ \mathsf{guard} \in \mathsf{LinArith}(\mathcal{M} \cup \mathsf{P})\}$. The set Rules is required to satisfy the following conditions:*
  1. *the subgraph induced by the rules ρ with $\rho.\mathsf{type} = 0$ is acyclic;*
  2. *no rule ρ with $\rho.\mathsf{type} = 0$ has $\rho.\mathsf{to} \in \mathcal{I}$;*
  3. *for all rules $\rho, \rho' \in \mathsf{Rules}$ such that $\rho.\mathsf{type} = 0$ and $\rho'.\mathsf{from}$ is reachable from $\rho.\mathsf{to}$ by a finite sequence of type-0 rules, the guard $\rho'.\mathsf{guard}$ is monotone in the message variables occurring in $\rho.\mathsf{guard}$.*

In this definition, rc specifies the admissible parameter valuations; $\mathcal{L}$ and $\mathcal{M}$ denote, respectively, the locations a process may occupy and the message types it may send in a round; and each rule specifies a source location, a target location, a round increment of at most b, and the guard under which it is enabled. Condition 1 ensures that each process visits each location at most once within a round. Condition 2 ensures that processes enter initial locations only at the start of a new round. Condition 3 is a technical monotonicity condition adapted from [41]. It is needed for the completeness proof of the received-message abstraction, one of the six transformations presented in § 3. Importantly, the overall reduction remains sound even without this condition. Moreover, as argued in [41], the monotonicity condition is often implicitly assumed by distributed algorithms designers, and are thus satisfied by a wide range of distributed algorithms, such as the ones from [9, 3, 38, 12, 39].

Informally, Condition 3 ensures that once a process takes a type-0 rule, increasing the message counts that appear in that rule's guard cannot disable any rule that may be taken later in the same round. For example, if $\rho.\mathsf{guard} = \mathsf{m}_1 + \mathsf{m}_2 < \mathsf{n}/2$, then every rule ρ' whose source location is reachable from $\rho.\mathsf{to}$ by a finite sequence of type-0 rules, and hence may be taken after ρ within the same round, must have a guard that is monotone in m_1 and m_2. Here, a guard $g \in \mathsf{LinArith}(X)$ is *monotone* in a set $Y \subseteq X$ if, for all valuations $\mu, \mu' : X \to \mathbb{N}$, whenever μ satisfies g, $\mu'(y) \geq \mu(y)$ for all $y \in Y$, and $\mu'(x) = \mu(x)$ for all $x \in X \setminus Y$, it also holds that μ' satisfies g.

While a template describes the symbolic parameters and the behavior of a single process, the semantics of several interacting processes is represented by an infinite-state transition system. We first define the fixed-instance semantics, for a given parameter valuation, and then lift it to the parameterized semantics over all admissible parameter valuations.

Given a template $\mathcal{T}$ and a parameter valuation $\nu : \mathsf{P} \to \mathbb{N}$ such that $\nu \models \mathsf{rc}$, the system of $\nu(n)$ identical processes executing $\mathcal{T}$ is represented by the action-labeled transition system $\mathsf{RDTS}(\mathcal{T}, \nu) = \langle S_\nu, I_\nu, Act_\nu, Tr_\nu \rangle$. Its semantics is sketched here and detailed in App. A.1. Elements of S_ν are *global states* (or configurations), each composed of a process state and a network state. The *process state* $\mathsf{PState} : [\![1, \nu(n)]\!] \to \langle \mathsf{loc} : \mathcal{L}, \mathsf{rd} : \mathbb{N}, \mathsf{rcvMsg} : \mathcal{M} \times \mathbb{N} \to \mathbb{N} \rangle$ stores for each process: its current round, location, and multiset of received message types. The latter records how many messages of a given type were received with a given round tag. The *network state* $\mathsf{NState} : \mathcal{M} \times \mathbb{N} \to \mathbb{N}$ collects the number of messages of each type and round tag that were broadcast. In every configuration, the number of received messages (in process state) is pointwise bounded by the number of broadcast messages (in network state). Initial global states I_ν are those where all processes are in initial locations at round 0 with empty set of received messages, and the network state contains no message. Actions in Act_ν are of two types: (i) a reception $\mathsf{Receive}(i, \langle m, r \rangle)$, where process i receives a message of type m tagged with round r, or (ii) an update $\mathsf{Update}(i, \rho, r)$, where process i updates its state according to rule ρ of the template at round r. A message reception is enabled for a process if its count of received messages (of that type and round tag) is strictly less than the number of similar broadcast messages in the network state; the process then increments its message count accordingly. An update $\mathsf{Update}(i, \rho, r)$ is enabled if process i is at location $\rho.\mathsf{frm}$ in round r, and its received messages with round tag r satisfy the guard $\rho.\mathsf{guard}$; the process location and its round are then updated to $\rho.\mathsf{to}$ and $r + \rho.\mathsf{type}$, and the process broadcasts a message $\mathsf{Bcast}(\rho.\mathsf{to})$ (unless $\rho.\mathsf{to} = \bot$) tagged with round $r + \rho.\mathsf{type}$ on the network state. An *execution* of $\mathsf{RDTS}(\mathcal{T}, \nu)$ is a finite or infinite alternating sequence of configurations and transitions starting from an initial configuration.

To reason about all admissible parameter valuations simultaneously, we lift the fixed-instance transition system RDTS to its parameterized counterpart $\mathsf{PRDTS}(\mathcal{T})$, which takes the disjoint union of all $\mathsf{RDTS}(\mathcal{T}, \nu)$ for $\nu \models \mathsf{rc}$ (see App. A.2).

2.2 History State-Count Properties and Logic

The choice of a property class is central to any verification framework, as it determines both the expressiveness of the specifications and the feasibility of the reductions. We therefore need a specification formalism that balances three requirements: (i) expressiveness, to capture standard correctness conditions for ARDAs, as enumerated in Fig. 3 (left); (ii) reduction-compatibility, to establish a sound and complete reduction from parameterized RDTS to finite-counter systems; (iii) LTL-translatability, so that the final verification problem can be expressed as standard LTL model checking.

Previously introduced specification formalisms for threshold automata fall short in expressing standard correctness properties of ARDAs. Specifically, $\mathsf{ELTL_{FT}}$ [27] lacks quantification over rounds and thus cannot express standard properties such as agreement. Finally, [7] does not define a formal specification logic and relies on assumptions about its successive abstractions that guarantee preservation of correctness.

These limitations motivate our choice of specification formalism. We propose a semantic class of *history state-count* (HSC) properties, chosen to satisfy expressiveness and reduction-compatibility. We also introduce HSC logic (HSCL), a syntactic fragment of temporal logic whose formulas denote a subset of HSC properties and are designed to be LTL-translatable

while retaining expressiveness. In contrast to prior state-based approaches [27, 7], the HSC properties are action-based. This distinction is important for our reduction: although relevant action histories could, in principle, be encoded by auxiliary state predicates, the specifications and reduction are most naturally formulated in terms of the occurrence of transitions. Thus, instead of taking a restricted fragment of state-based LTL as the specification language, we use action-based HSCL and later translate it into LTL over the reduced finite-counter system.

2.2.1 History State-Count Properties

A History State-Count (HSC) property applies to history state-count traces of executions. For an execution π, its *history state-count trace* (hsc-trace) is a function $\mathsf{hsc}(\pi) : \mathcal{L} \times \mathbb{N} \to \mathbb{N}$ that records, along π, the aggregate number of process visits to each location in every round. Formally, for a location ℓ and a round tag r, $\mathsf{hsc}(\pi)(\ell, r)$ denotes the number of update actions, with some rule ρ and round tag r' taken along π such that $\rho.\mathsf{to} = \ell$ and $r' + \rho.\mathsf{type} = r$. Hence, the hsc-trace is entirely determined by the multiset of pairs (ρ, r) corresponding to the update actions occurring along the execution. Since each process can visit any location at most once per round (Condition 1, Def. 1), the hsc-trace equivalently captures, for every pair of a location and a round tag, the number of processes that have visited it.

For a template $\mathcal{T}$ with locations $\mathcal{L}$, a *history state-count* (HSC) property is a family $\mathcal{P} = \{\mathcal{P}_\nu \subseteq \mathsf{Dom} \mid \nu : \mathsf{P} \to \mathbb{N}\}$, where $\mathsf{Dom} = \{f : \mathcal{L} \times \mathbb{N} \to \mathbb{N}\}$. We write $\mathsf{Traces}(\mathsf{RDTS}(\mathcal{T}, \nu)) = \{\mathsf{hsc}(\pi) \mid \pi \text{ is an execution of } \mathsf{RDTS}(\mathcal{T}, \nu)\}$. We write $\mathsf{RDTS}(\mathcal{T}, \nu) \models \mathcal{P}$ iff $\mathsf{Traces}(\mathsf{RDTS}(\mathcal{T}, \nu)) \subseteq \mathcal{P}_\nu$, and define $\mathsf{PRDTS}(\mathcal{T}) \models \mathcal{P}$ iff, for all valid parameter valuations ν, $\mathsf{RDTS}(\mathcal{T}, \nu) \models \mathcal{P}$.

2.2.2 History State-Count Logic (HSCL)

Informally, HSCL has two types of atomic constraints: (i) *universal, round-local constraints* of the form $\forall r.\alpha_r$, which, for each round r, allow one to bound a weighted sum of visits to locations in that round by a threshold; and (ii) *cumulative constraints*, written β, which allow one to bound the corresponding weighted sum aggregated over all rounds. Formulas in HSCL are Boolean combinations (using $\neg, \wedge$) of these two types. Note that existential quantification over rounds is expressible as $\neg\forall r.\alpha_r \equiv \exists r.\neg\alpha_r$.

▶ **Definition 2** (HSC). *Let $\mathcal{X}$ be a finite set of locations and $\mathcal{Y}$ a finite set of parameters. Formulas of* HSCL *follow the grammar:*

$$\varphi \;::=\; \forall r.\,\alpha_r \;\mid\; \beta \;\mid\; \neg\varphi \;\mid\; (\varphi \wedge \varphi)$$

$$\text{with } \alpha_r \;::=\; \sum_{\ell \in \mathcal{X}} c_\ell \cdot \kappa(\ell, r) \;\leq\; t, \quad \beta \;::=\; \sum_{\ell \in \mathcal{X}} c_\ell \cdot \sum_{r \in \mathbb{N}} \kappa(\ell, r) \;\leq\; t$$

where $c_\ell \in \mathbb{N}$ are non-negative weights*; $\kappa(\ell, r)$ is a* variable *(for $\ell \in \mathcal{X}$, $r \in \mathbb{N}$) interpreted as the number of process-visits to ℓ in round r; and the* threshold *t is a linear term over $\mathcal{Y}$ of the form $t = b_0 + \sum_{y \in \mathcal{Y}} b_y \cdot y$, with $b_0, b_y \in \mathbb{Z}$. We call instances of α_r* round-local atoms *and instances of β* cumulative atoms[1].

For a process template $\mathcal{T}$ with a location set $\mathcal{L}$ and a parameter set P, we interpret $\mathcal{X}$ as a subset of $\mathcal{L}$ and $\mathcal{Y}$ as a subset of P. A HSCL formula φ represents the HSC property $\mathcal{P}^\varphi = \{\mathcal{P}^\varphi_\nu \mid \nu : \mathsf{P} \to \mathbb{N}\}$ where for a parameter valuation ν, set $\mathcal{P}^\varphi_\nu = \{f : \mathcal{L} \times \mathbb{N} \to \mathbb{N} \mid$

[1] Although the sum in β may diverge, since $c_\ell \geq 0$ and t is finite, satisfaction of β is determined by a finite partial sum.

Agreement: Either no correct process decides on 0, or no correct process decides on 1.
$\mathtt{A} := \forall r.\kappa(\mathsf{d0}, r) \leq 0 \ \vee \forall r.\kappa(\mathsf{d1}, r) \leq 0$ $\quad \mathsf{LTL}(\mathtt{A}) \ = \mathbf{G}\big(\mathsf{local}(\mathsf{d0}) \leq 0\big) \ \vee \ \mathbf{G}\big(\mathsf{local}(\mathsf{d1}) \leq 0\big)$

Validity: If no correct process starts with 0, then no correct process decides on 0.
$\mathtt{V} := \forall r.\kappa(\mathsf{d0}, r) \leq 0$ $\quad \mathsf{LTL}(\mathtt{V}) \ = \mathbf{G}\big(\mathsf{local}(\mathsf{d0}) \leq 0\big)$

Termination: Eventually, all correct processes decide.
$\mathtt{T} := \neg(\sum_r \kappa(\mathsf{d0}, r) + \kappa(\mathsf{d1}, r) \leq \mathtt{N_c} - 1)$ $\quad \mathsf{LTL}(\mathtt{T}) \ = \neg\mathbf{G}\big(\mathsf{cumul}(\mathsf{d0}) + \mathsf{cumul}(\mathsf{d1}) \leq \mathtt{N_c}\big)$

Restricted Termination: If one correct process decides, then all correct processes decide.
$\mathtt{RT} := \neg(\sum_r \kappa(\mathsf{d0}, r) + \kappa(\mathsf{d1}, r) \leq 0) \implies \mathtt{T}$ $\quad \mathsf{LTL}(\mathtt{RT}) \ = \neg\mathbf{G}\big(\mathsf{cumul}(\mathsf{d0}) + \mathsf{cumul}(\mathsf{d1}) \leq 0\big) \implies \mathsf{LTL}(\mathtt{T})$

Leader Uniqueness: At most one leader process per round.
$\mathtt{LU} := \forall r.\kappa(\mathsf{ldr}, r) \leq 1$ $\quad \mathsf{LTL}(\mathtt{LU}) \ = \mathbf{G}\big(\mathsf{local}(\mathsf{ldr}) \leq 1\big)$

Figure 3 Expressing correctness properties of ARDAs in HSCL and their LTL translations (see § 3.2). $\mathcal{X} = \{\mathsf{d0}, \mathsf{d1}, \mathsf{ldr}\}$ denote decision-0, decision-1, and leader locations, respectively, and $\mathcal{Y} = \{\mathtt{N}_c\}$ denotes the number of correct processes. Formulas for (restricted) termination assume that processes halt after making a decision.

$f, \nu \models \varphi\}$. Here, $f, \nu \models \varphi$ iff φ evaluates to true under the substitution of each threshold term t in φ with its value under ν (i.e., $t[\mathcal{Y} \leftarrow \nu]$), and of each variable $\kappa(x, r)$ with $f(x, r)$. We write $\mathsf{RDTS}(\mathcal{T}, \nu) \models \varphi$ if $\mathsf{RDTS}(\mathcal{T}, \nu) \models \mathcal{P}^{\varphi}$,and $\mathsf{PRDTS}(\mathcal{T}) \models \varphi$ if $\mathsf{PRDTS}(\mathcal{T}) \models \mathcal{P}^{\varphi}$.

HSCL can express the five standard correctness properties listed in Fig. 3. The same figure gives the corresponding HSCL formulas and their LTL translations. We discuss the translation from HSCL to LTL in § 3.2. As we will see, HSCL formulas infact translate to restricted LTL properties, namely obligation properties [33]. These five properties guided the design of HSCL, whose syntax is tailored to encode them effectively. Although HSCL is intentionally minimal, our reduction to a finite-counter system applies more broadly to arbitrary HSC properties. Thus, the logic is not intended as a general-purpose specification language, but rather as a concise intermediate representation that facilitates translation to LTL.

2.3 Verification Problems

We can now define the four verification problems studied in this work. They are obtained by combining a choice of semantics, fixed-instance or parameterized, with a choice of specification formalism: either an arbitrary HSC property or a property expressed in HSCL.

P1 Fixed-instance HSC-Verification Given a template $\mathcal{T}$, an HSC property $\mathcal{P}$, and a parameter valuation ν with $\nu \models \mathsf{rc}$, does $\mathsf{RDTS}(\mathcal{T}, \nu) \models \mathcal{P}$?

P2 Parameterized HSC-Verification Given a template $\mathcal{T}$ and an HSC property $\mathcal{P}$, does $\mathsf{PRDTS}(\mathcal{T}) \models \mathcal{P}$?

P3 Fixed-instance HSCL-Verification Given a template $\mathcal{T}$, an HSCL formula φ, and a parameter valuation ν with $\nu \models \mathsf{rc}$, does $\mathsf{RDTS}(\mathcal{T}, \nu) \models \varphi$?

P4 Parameterized HSCL-Verification Given a template $\mathcal{T}$ and an HSCL formula φ, does $\mathsf{PRDTS}(\mathcal{T}) \models \varphi$?

Note that P1 and P2 are mathematical problems, whereas P3 and P4 are computational decision problems. The reductions developed in the next section relate these four problems to model checking problems over finite-state, one-counter, and finite-counter systems; see Table 1. In particular, the reduction for P3 yields LTL model checking over a finite-state system, and is therefore decidable. By contrast, the reduction for P4 yields LTL model checking over

a finite-counter system. Here, a finite-counter system [2], is a symbolic representation of a transition system whose states are valuations of finitely many variables, called *counters*, ranging over the non-negative integers, and whose set of initial states and transition relation are definable by linear-arithmetic formulas over these counters. For model checking over such systems, we allow LTL atoms to be linear-arithmetic formulas over the counters; see App. A.3 for formal definitions.

LTL model checking over finite-counter systems is undecidable in general, thus the reduction above does not by itself give a decision procedure for P4. We next show that P4 is indeed undecidable, by extending the undecidability result of [40] from synchronous round-based distributed algorithms to our asynchronous setting.

Table 1 Reductions to model-checking (MC) problems over one-counter, finite-counter, and finite-state systems. All reductions are sound and complete.

Problem	Reduced problem	Ref.
P1 Fixed-instance HSC-Verification	MC of one-counter systems w.r.t. HSC	Thm. 9
P2 Parameterized HSC-Verification	MC of finite-counter systems w.r.t. HSC	Thm. 10
P3 Fixed-instance HSCL-Verification	MC of finite-state systems w.r.t. LTL	Thm. 13
P4 Parameterized HSCL-Verification	MC of finite-counter systems w.r.t. LTL	Thm. 14

▶ **Theorem 3.** *Parameterized* HSCL*-Verification is undecidable.*

Proof. We reduce from the non-halting problem for two-counter machines (known to be undecidable [8]). Let $\mathcal{C}$ be a two-counter machine with a finite set of control states S, an initial state s_{init}, a halting state s_{halt}, and a finite set of commands Δ. Each command in Δ has the form (s, op, s'), where op is an operation of the form $\mathsf{inc}(c_i)$, $\mathsf{dec}(c_i)$, or $\mathsf{iszero}(c_i)$ for $i \in \{1, 2\}$. Commands with operations $\mathsf{dec}(c_i)$ and $\mathsf{iszero}(c_i)$ are enabled only when $c_i > 0$ and $c_i = 0$, respectively. From $\mathcal{C}$ we construct a round-based process template $\mathcal{T}_{\mathcal{C}}$ and an HSCL formula $\varphi_{\mathcal{C}}$ such that $\mathcal{C}$ has no halting execution iff $\mathsf{PRDTS}(\mathcal{T}_{\mathcal{C}}) \models \varphi_{\mathcal{C}}$.

The construction maintains the following invariant. The r-th configuration (s, v_1, v_2) of an execution of $\mathcal{C}$ is represented by a global configuration in which all processes have just jumped to round r, with one process, referred to as the controller, in location l_s, and v_i many processes in location l_{c_i} (for $i = 1, 2$), and the rest of the processes are in l_{res} ("reserve"). Thus, initially, one process is in $l_{s_{\mathsf{init}}}$ and all other processes are in l_{res}.

A transition in $\mathcal{C}$ from configuration (s, v_1, v_2) on a command, say, $(s, \mathsf{inc}(c_1), s')$, is simulated as follows. First, the controller moves from l_s to an intermediate location $l_{(s,\mathsf{inc}(c_1),s')}$ in the same round and broadcasts message m_{inc,c_1}; this message enables a process to move from l_{res} to another intermediate location l_{inc,c_1} of the same round and broadcast the message m_{ack}; this acknowledgment then allows all processes to move to their corresponding locations in the next round, i.e., the controller to $l_{s'}$, the process in l_{inc,c_1} to l_{c_1}, and the remaining processes to stay in their current locations. Upon entering the new round, each non-controller process broadcasts a ready message indicating whether it is in l_{c_1}, l_{c_2}, or l_{res}. Once the controller enters the next round and receives all $n-1$ ready messages, the simulated transition is complete. The other commands are simulated similarly; in particular, the ready-messages counts are used to test the enabledness of decrement and zero-test commands.

The template semantics is an over-approximation of the intended simulation. For example, several processes in l_{res} may receive m_{inc,c_1} and move to l_{inc,c_1} in the same round, thereby increasing the value of c_1 by more than one. The formula $\varphi_{\mathcal{C}} := \neg\psi_{\mathcal{C}}$ where $\psi_{\mathcal{C}}$ conjuncts

[2] finite-counter systems are instances of Presburger counter systems [17] with quantifier-free predicates

conditions ensuring well-formedness of the simulation with a halting condition: in each round, at most one process may visit some location l_s with $s \in S$, at most one process may visit an intermediate location l_{inc,c_i} or l_{dec,c_i}, and some process eventually visits $l_{s_{\mathsf{halt}}}$. Thus, $\varphi_{\mathcal{C}}$ holds on all executions of $\mathsf{PRDTS}(\mathcal{T}_{\mathcal{C}})$ iff no well-formed simulated execution of $\mathcal{C}$ reaches s_{halt}. ◀

3 Reductions

This section presents the reductions summarized in Table 1. They are obtained by a sequence of sound and complete transformation steps. Due to space constraints, we present only the intuition behind each step. For the formal definitions and proofs refer to the long version of the paper [4].

3.1 Reductions for HSC Properties

In this subsection, we present a reduction pipeline for HSC properties, shown in Fig. 4. The pipeline is sound and complete for HSC properties, and reduces fixed-instance semantics to one-counter systems and parameterized semantics to finite-counter systems. At a high level, the counters of the reduced systems record, for each relevant round, (i) the number of processes at each location and (ii) the number of messages of each type that have been broadcast. For the fixed-instance semantics, given a template $\mathcal{T}$ and valuation ν, the system $\mathsf{RDTS}(\mathcal{T}, \nu)$ reduces to a finite-counter system in which all counters except one are bounded by the number of processes $\nu(n)$; hence, the target is a *one-counter system*. For the parameterized semantics, we exploit the fact that our reductions for fixed instances guarantee that the set of counters is identical across parameter valuations. Taking the disjoint union over all admissible ν thus yields a *finite-counter system*. The reduction pipeline consists of four steps. At each step X, we denote by $\mathsf{RDTS}^X(\mathcal{T}, \nu) = \langle S^X_\nu, I^X_\nu, Act^X_\nu, Tr^X_\nu \rangle$ the fixed-instance semantics after the step X.

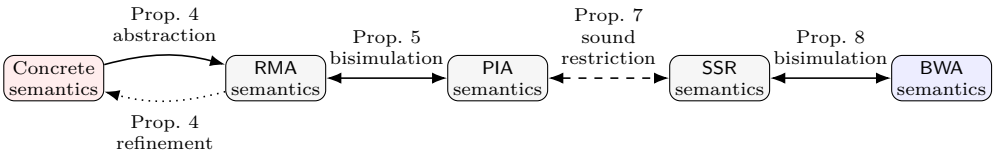

Figure 4 Reduction pipeline for fixed-instance and parameterized verification of HSC properties. All steps are sound and complete. RMA stands for received-message abstraction, PIA for process-identity abstraction, SSR for semi-synchronous restriction, and BWA for bounded-window abstraction. In the fixed-instance (resp. parameterized) setting, the target semantics BWA is a one-counter (resp. finite-counter) system.

Step 1. Received-Message Abstraction (RMA)

The first step exploits the fact that the hsc-trace of an execution is insensitive to Receive transitions. Consequently, receive transitions can be treated as stutter steps. In the abstract system $\mathsf{RDTS}^{\mathsf{RMA}}(\mathcal{T}, \nu)$, the receive transitions are dropped and the elements of S^{RMA}_ν remain pairs of process and network states, except that the process state is simplified to $\mathsf{PState} : [\![1, \nu(n)]\!] \to \langle \mathsf{loc} : \mathcal{L}, \mathsf{rd} : \mathbb{N} \rangle$, omitting rcvMsg. Since Update transitions in $\mathsf{RDTS}(\mathcal{T}, \nu)$ are enabled only when their guards are justified by received messages, enabledness must now be redefined: in the abstract system, guards are evaluated *existentially* with respect to subsets of messages recorded in the network state NState. While abstraction of received

messages has appeared in prior work [41, 7] in a state-based form, we give here an action-based formulation that treats receive transitions as stuttering steps. This yields refinement and forward simulation in the sense of Lynch et al. [32].

With Receive transitions as stutter steps, we establish a forward simulation [32, § 3.2] between the concrete semantics and the RMA semantics, proving that RMA is sound for HSC properties. The completeness of RMA is more subtle. The abstraction discards the concrete record of which messages a process has already *received and used* to satisfy previous guards. As a result, in the abstract enabledness relation, the witness for one update may be $R_1 \leq \mathsf{NState}$, while the witness for the very next update may be a different $R_2 \leq \mathsf{NState}$, without respecting the concrete monotonicity constraint that the second guard must hold on a super-set of messages already used (i.e., $R_2 \supseteq R_1$). Consequently, the abstraction may suggest that two updates can be carried out consecutively, even though no message-delivery schedule in the concrete system can realize them. Condition 3 in Def. 1 is imposed to exclude this phenomenon: once a process takes a type-0 rule, increasing the message counts used to justify its guard cannot disable any rule that may be taken later in the same round. Under this condition, the abstract witnesses for future updates can always be arranged consistently with a concrete receive schedule. Thus, with Receive transitions as stutter steps, we establish a refinement [32, § 3.1] between the RMA semantics to the concrete semantics, proving that RMA is complete for HSC properties.

▶ **Proposition 4** (Soundness and Completeness of RMA). *For every template $\mathcal{T}$ and parameter valuation ν:* $\mathsf{Traces}(\mathsf{RDTS}(\mathcal{T}, \nu)) = \mathsf{Traces}(\mathsf{RDTS}^{\mathsf{RMA}}(\mathcal{T}, \nu))$.

Step 2. Process-Identity Abstraction (PIA)

The second step exploits the fact that the hsc-trace of an execution is insensitive to the identities of the processes performing Update transitions. Consequently, process identities can be dropped, resulting in an abstraction in which the exact states of processes are no longer recorded, only their counts. In the resulting abstract system $\mathsf{RDTS}^{\mathsf{PIA}}(\mathcal{T}, \nu)$, the elements of S_ν^{PIA} remain pairs of process and network states, except that the process state is simplified to $\mathsf{PState} : \mathcal{L} \times \mathbb{N} \to \mathbb{N}$. Here $\mathsf{PState}(\ell, r)$ denote the number of processes in location ℓ of round r. Elements of Act_ν^{PIA} are of the form $\mathsf{Update}(\rho, r)$, again omitting process identities. Since the seminal paper [20], this counting abstraction is classical in the verification of systems composed of identical anonymous processes. The process-identity abstraction induces an action-based bisimulation [1, Def. 7.15] between $\mathsf{RDTS}^{\mathsf{RMA}}(\mathcal{T}, \nu)$ and $\mathsf{RDTS}^{\mathsf{PIA}}(\mathcal{T}, \nu)$, which entails the following result:

▶ **Proposition 5** (Soundness and Completeness of PIA). *For every template $\mathcal{T}$ and parameter valuation ν,* $\mathsf{Traces}(\mathsf{RDTS}^{\mathsf{PIA}}(\mathcal{T}, \nu)) = \mathsf{Traces}(\mathsf{RDTS}^{\mathsf{RMA}}(\mathcal{T}, \nu))$.

The system $\mathsf{RDTS}^{\mathsf{PIA}}(\mathcal{T}, \nu)$ is a counter system with *infinitely many* counters. Indeed, each round $r \in \mathbb{N}$ introduces $|\mathcal{L}| + |\mathcal{M}|$ counters: a location counter $\mathsf{PState}(\ell, r)$ for every $\ell \in \mathcal{L}$ and a message counter $\mathsf{NState}(m, r)$ for every $m \in \mathcal{M}$. To obtain a finite-counter reduction, a natural question is whether all counters must be tracked at all times. Intuitively, counters become redundant once they can no longer influence future updates. To capture this, we define the *relevant window* of a configuration as the interval of round tags whose counters may be nonzero and can affect future updates. Under the PIA semantics, for a configuration cfg, this window spans from the smallest ($r_{\mathsf{min}}(\mathsf{cfg})$) to the largest ($r_{\mathsf{max}}(\mathsf{cfg})$) round occupied by a process in cfg. A uniformly bounded relevant window across all configurations enables a direct reduction to a finite number of counters, as discussed in Step 4. However, due

to asynchrony, the *round drift* $r_{\mathsf{max}}(\mathsf{cfg}) - r_{\mathsf{min}}(\mathsf{cfg})$ is unbounded across configurations of $\mathsf{RDTS}^{\mathsf{PIA}}(\mathcal{T}, \nu)$, preventing such a reduction. To this end, our next step constrains the enabledness of updates to enforce a uniformly bounded relevant window while preserving soundness for HSC properties.

Step 3. Semi-Synchronous Restriction (SSR)

The third step exploits the fact that the hsc-trace of an execution is insensitive to the order of Update transitions. Consequently, verification can be restricted to a subset of executions of $\mathsf{RDTS}^{\mathsf{PIA}}(\mathcal{T}, \nu)$ such that every execution has a counterpart within this subset that executes the same multiset of Update actions. To achieve this, we impose a *semi-synchronous restriction* (SSR), which permits an action $a = \mathsf{Update}(\rho, r)$ from a configuration cfg only if its target round $\mathsf{tgt}(a) = r + \rho.\mathsf{type}$ is at least $r_{\mathsf{max}}(\mathsf{cfg})$, i.e., the highest round currently occupied by a process in cfg, also called the *frontier round* of cfg. The SSR therefore restricts executions by forbidding transitions that would move a process below the frontier round. As a result, since round jumps are bounded by b, the set of relevant counters in a configuration cfg is confined to the interval $[\![r_{\mathsf{max}}(\mathsf{cfg}) - b,\, r_{\mathsf{max}}(\mathsf{cfg})]\!]$, which has a uniform size of $b{+}1$ across all configurations. To establish the soundness of SSR for HSC properties, we use commutativity arguments similar to the one used in [13].

An execution $\pi = \mathsf{cfg}_1\, a_1\, \mathsf{cfg}_2\, a_2 \cdots$ of $\mathsf{RDTS}^{\mathsf{PIA}}(\mathcal{T}, \nu)$ is called *semi-synchronous* iff its target sequence $\mathsf{tgt}(a_1)\mathsf{tgt}(a_2)\ldots$ is nondecreasing. This notion precisely characterizes the executions of $\mathsf{RDTS}^{\mathsf{SSR}}(\mathcal{T}, \nu)$. The equivalence follows from the observation that, for every transition $(\mathsf{cfg}, a, \mathsf{cfg}') \in Tr_{\nu}^{\mathsf{PIA}}$, the frontier round evolves as $r_{\mathsf{max}}(\mathsf{cfg}') = \max(r_{\mathsf{max}}(\mathsf{cfg}), \mathsf{tgt}(a))$.

▶ **Lemma 6.** *In* $\mathsf{RDTS}^{\mathsf{PIA}}(\mathcal{T}, \nu)$*, for every execution* π *there exists a semi-synchronous execution* π' *that starts from the same initial configuration as* π *and contains the same multiset of update actions.*

The proof proceeds by transforming any execution π into a semi-synchronous execution through adjacent swaps of out-of-order update actions, thereby eliminating inversions in the target-round sequence. Each swap is justified by the *commutativity property*: two consecutive updates commute whenever the source round of the later differs from the target round of the earlier. To ensure convergence, swaps are performed in a disciplined manner that gives a monotone sequence of prefixes converging to a well-defined limiting execution, which is semi-synchronous.

The same commutativity property gives a stronger result that further reduces nondeterminism in semi-synchronous executions, which we exploit in the experiments reported in § 4. For every execution, there exists a semi-synchronous execution (with the same initial configuration and the same multiset of update actions) with the additional property that we call *strong*: for each round r, all *jump updates* targeting r (whose source round is smaller than r) happen before all *local updates* targeting r (whose source round equals r). Moreover, all jump updates are executed as a multiset in a fixed order which can alternatively be viewed as occurring synchronously in a single step.

▶ **Proposition 7** (Soundness and Completeness of SSR)**.** *For every template* $\mathcal{T}$ *and parameter valuation* ν*,* $\mathsf{Traces}(\mathsf{RDTS}^{\mathsf{SSR}}(\mathcal{T}, \nu)) = \mathsf{Traces}(\mathsf{RDTS}^{\mathsf{PIA}}(\mathcal{T}, \nu))$.

Step 4. Bounded-Window Abstraction (BWA)

The fourth abstraction step exploits the observation that configurations need not maintain counters that can no longer influence future updates. Since in $\mathsf{RDTS}^{\mathsf{SSR}}(\mathcal{T}, \nu)$ the size of the relevant window is bounded by $b + 1$ across all configurations, the BWA forgets counters

outside this window and encodes the relevant portion as a sliding window of width $b+1$. In the resulting abstract system $\mathsf{RDTS}^{\mathsf{BWA}}(\mathcal{T},\nu)$, each configuration in S_ν^{BWA} is a triple consisting of a frontier round $r_{\mathsf{max}} \in \mathbb{N}$, a process state $\mathsf{PState} : \mathcal{L} \times [\![0,b]\!] \to \mathbb{N}$ and a network state $\mathsf{NState} : \mathcal{M} \times [\![0,b]\!] \to \mathbb{N}$. The process and network states record counters only for rounds ranging from the current frontier round down to b rounds below. The BWA induces an *action-based bisimulation* [1, Def. 7.15] between $\mathsf{RDTS}^{\mathsf{SSR}}(\mathcal{T},\nu)$ and $\mathsf{RDTS}^{\mathsf{BWA}}(\mathcal{T},\nu)$, establishing the following result:

▶ **Proposition 8** (Soundness and Completeness of BWA). *For every template $\mathcal{T}$ and parameter valuation ν,* $\mathsf{Traces}(\mathsf{RDTS}^{\mathsf{BWA}}(\mathcal{T},\nu)) = \mathsf{Traces}(\mathsf{RDTS}^{\mathsf{SSR}}(\mathcal{T},\nu))$.

The system $\mathsf{RDTS}^{\mathsf{BWA}}(\mathcal{T},\nu)$ is a finite-counter system. Indeed, it maintains a single frontier-round counter r_{max}, and for each depth $d \in [\![0,b]\!]$, there are $|\mathcal{L}| + |\mathcal{M}|$ counters: a location counter $\mathsf{PState}(\ell,d)$ for every $\ell \in \mathcal{L}$ and a message counter $\mathsf{NState}(m,d)$ for every $m \in \mathcal{M}$. Moreover, $\mathsf{RDTS}^{\mathsf{BWA}}(\mathcal{T},\nu)$ can be viewed as a *one-counter system* with a *nondecreasing* counter. This is because the total number of processes $\nu(n)$ is bounded, and each process can send only a bounded number of messages ($\nu(n) \times |\mathcal{L}|$) per round. Consequently, all location and message counters are bounded, and the only unbounded counter is r_{max}, which increases monotonically. Thus, from Prop. 4, 5, 7, and 8, we have:

▶ **Theorem 9.** *For every template $\mathcal{T}$, parameter valuation ν, and* HSC *property $\mathcal{P}$, the system $\mathsf{RDTS}^{\mathsf{BWA}}(\mathcal{T},\nu)$ is a one-counter system with a nondecreasing counter r_{max}, and $\mathsf{RDTS}^{\mathsf{BWA}}(\mathcal{T},\nu) \models \mathcal{P}$ iff $\mathsf{RDTS}(\mathcal{T},\nu) \models \mathcal{P}$.*

We define the parameterized BWA semantics, denoted $\mathsf{PRDTS}^{\mathsf{BWA}}(\mathcal{T})$, as the disjoint union of all fixed-instance systems $\mathsf{RDTS}^{\mathsf{BWA}}(\mathcal{T},\nu)$ over admissible parameter valuations ν. Since the set of counter variables in $\mathsf{RDTS}^{\mathsf{BWA}}(\mathcal{T},\nu)$ is independent of ν, the states of $\mathsf{PRDTS}^{\mathsf{BWA}}(\mathcal{T})$ can be represented using finitely many counters: the frontier-round counter; the location and message counters inherited from $\mathsf{RDTS}^{\mathsf{BWA}}(\mathcal{T},\nu)$; and an additional set of $|\mathsf{P}|$ counters encoding the parameter valuation ν. Note, however, that the location and message counters in parameterized states need not be bounded, since the parameter $\nu(n)$ itself may range over unbounded values. Recall from § 2.2.1 that the parameterized semantics satisfies an HSC property iff all its admissible fixed-instance semantics do. Hence, from Thm. 9, we have:

▶ **Theorem 10.** *For every template $\mathcal{T}$ and* HSC *property $\mathcal{P}$, the system $\mathsf{PRDTS}^{\mathsf{BWA}}(\mathcal{T})$ is a finite-counter system, and $\mathsf{PRDTS}^{\mathsf{BWA}}(\mathcal{T}) \models \mathcal{P}$ iff $\mathsf{PRDTS}(\mathcal{T}) \models \mathcal{P}$.*

3.2 Further Reductions for HSCL Formulas

We now present a further reduction pipeline for HSCL formulas, shown in Fig. 5.

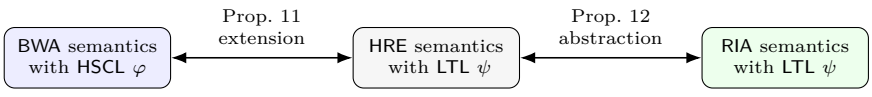

Figure 5 Further reduction pipeline for fixed-instance and parameterized verification of HSCL formulas. All steps are sound and complete. BWA stands for bounded-window abstraction, HRE for history-record extension, and RIA for round-identity abstraction. In the fixed-instance (resp. parameterized) setting, the target semantics RIA is a finite-state (resp. finite-counter) system.

Step 5. History-Record Extension (HRE)

The fifth step evaluates HSCL formulas from the state point of view by extending each BWA configuration with a "history record" that retains exactly those fragments of the execution history that are relevant to HSCL. Given an HSCL formula φ over locations $\mathcal{X}$ and parameters $\mathcal{Y}$, we construct an extended system $\mathsf{RDTS}^{\mathsf{HRE}}(\mathcal{T}, \nu, \varphi)$ by augmenting every state in S^{BWA}_{ν} with two counters for each $\ell \in \mathcal{X}$: a cumulative counter $\mathsf{cumul}(\ell)$, recording the number of visits to ℓ across all rounds; and a local counter $\mathsf{local}(\ell)$, recording the number of visits to ℓ in the current frontier round. Only the frontier round requires a dedicated local counter, since under BWA no update targets a round strictly below the frontier round. Thus, once a round is left behind, its local counts of process visits become fixed.

Observe that local counters are bounded by the total number of processes. In contrast, cumulative counters may grow unboundedly. However, for a fixed parameter valuation ν and formula φ, each cumulative counter can be capped at one more than the largest threshold appearing in the cumulative atoms of φ. Counts beyond this bound are irrelevant to satisfaction of HSCL formulas and can therefore be safely truncated. As a result, the system $\mathsf{RDTS}^{\mathsf{HRE}}(\mathcal{T}, \nu, \varphi)$ remains a single-counter system with a nondecreasing counter r_{max}.

Next, we give a recursive procedure to translate the HSCL formula φ into an LTL formula $\mathsf{LTL}(\varphi)$. We use a variant of the standard LTL syntax in which the only temporal operator is "globally", written $\mathbf{G}$, and atomic propositions are linear-arithmetic formulas over the counter variables. For universal round local atoms $\forall_r.\alpha_r$ where $\alpha_r = \sum_{\ell \in \mathcal{X}} c_\ell \cdot \kappa(\ell, r) \leq t$, define the translation $\mathsf{LTL}(\forall_r.\alpha_r) = \mathbf{G}(\sum_{\ell \in \mathcal{X}} c_\ell \cdot \mathsf{local}(\ell) \leq t)$. For cumulative atoms $\beta = \sum_{\ell \in \mathcal{X}} c_\ell \cdot \sum_r \kappa(\ell, r) \leq t$, define the translation $\mathsf{LTL}(\beta) = \mathbf{G}(\sum_{\ell \in \mathcal{X}} c_\ell \cdot \mathsf{cumul}(\ell) \leq t)$. The Boolean cases are handled compositionally: $\mathsf{LTL}(\neg\varphi) = \neg\mathsf{LTL}(\varphi)$ and $\mathsf{LTL}(\varphi \wedge \varphi') = \mathsf{LTL}(\varphi) \wedge \mathsf{LTL}(\varphi')$. Fig. 3 provides LTL translations for core HSCL formulas.

▶ **Proposition 11** (Soundness and Completeness of HRE). *For every* HSCL *formula* φ, $\mathsf{RDTS}^{\mathsf{HRE}}(\mathcal{T}, \nu, \varphi) \models \mathsf{LTL}(\varphi)$ *iff* $\mathsf{RDTS}^{\mathsf{BWA}}(\mathcal{T}, \nu) \models \varphi$.

Step 6. Round-Identity Abstraction (RIA)

The sixth and final step exploits the fact that the properties expressible in HSCL are insensitive to the round identifiers. Consequently, the frontier-round counter $r_{\max}$ can be dropped from HRE configurations. The resulting abstract system $\mathsf{RDTS}^{\mathsf{RIA}}(\mathcal{T}, \nu, \varphi)$ induces a *state-based bisimulation* [1, Def. 7.7] with respect to $\mathsf{RDTS}^{\mathsf{HRE}}(\mathcal{T}, \nu, \varphi)$, giving:

▶ **Proposition 12** (Soundness and Completeness of RIA). *For every* HSCL *formula* φ, $\mathsf{RDTS}^{\mathsf{RIA}}(\mathcal{T}, \nu, \varphi) \models \mathsf{LTL}(\varphi)$ *iff* $\mathsf{RDTS}^{\mathsf{HRE}}(\mathcal{T}, \nu, \varphi) \models \mathsf{LTL}(\varphi)$.

Note that the system $\mathsf{RDTS}^{\mathsf{RIA}}(\mathcal{T}, \nu, \varphi)$ is a finite-counter system, all of whose counters are bounded by a finite value determined by the parameter valuation ν, and is therefore a finite-state system. By Thm. 9 and Prop. 11 and 12, we have:

▶ **Theorem 13.** *For every template* $\mathcal{T}$*, parameter valuation* ν*, and an* HSCL *formula* φ*, the system* $\mathsf{RDTS}^{\mathsf{RIA}}(\mathcal{T}, \nu, \varphi)$ *is a finite-state system, and* $\mathsf{RDTS}^{\mathsf{RIA}}(\mathcal{T}, \nu, \varphi) \models \mathsf{LTL}(\varphi)$ *iff* $\mathsf{RDTS}(\mathcal{T}, \nu) \models \varphi$.

The parameterized RIA semantics, denoted $\mathsf{PRDTS}^{\mathsf{RIA}}(\mathcal{T}, \varphi)$, is defined as the disjoint union of all fixed-instance systems $\mathsf{RDTS}^{\mathsf{RIA}}(\mathcal{T}, \nu, \varphi)$ over all admissible parameter valuations ν. Since the set of counters in $\mathsf{RDTS}^{\mathsf{RIA}}(\mathcal{T}, \nu, \varphi)$ does not depend on ν, each state of $\mathsf{PRDTS}^{\mathsf{RIA}}(\mathcal{T}, \varphi)$ can be represented with finitely many counters: those of $\mathsf{RDTS}^{\mathsf{RIA}}(\mathcal{T}, \nu, \varphi)$ together with an additional $|\mathsf{P}|$ counters encoding the parameter valuation ν. Hence,

$\mathsf{PRDTS}^{\mathsf{RIA}}(\mathcal{T},\varphi)$ is a finite-counter system. The atoms of the LTL formula $\mathsf{LTL}(\varphi)$ are interpreted over these parameter counters and over the local/cumulative visit counters of $\mathsf{PRDTS}^{\mathsf{RIA}}(\mathcal{T},\varphi)$. By the definition of the disjoint union, $\mathsf{PRDTS}^{\mathsf{RIA}}(\mathcal{T},\varphi) \models \mathsf{LTL}(\varphi)$ iff, for all $\nu \models \mathsf{rc}$, $\mathsf{RDTS}^{\mathsf{RIA}}(\mathcal{T},\nu,\varphi) \models \mathsf{LTL}(\varphi)$. Therefore, by Thm. 13, we obtain:

▶ **Theorem 14.** *For every template $\mathcal{T}$ and* HSCL *formula φ, the system $\mathsf{PRDTS}^{\mathsf{RIA}}(\mathcal{T},\varphi)$ is a finite-counter system, and $\mathsf{PRDTS}^{\mathsf{RIA}}(\mathcal{T},\varphi) \models \mathsf{LTL}(\varphi)$ iff $\mathsf{PRDTS}(\mathcal{T}) \models \varphi$.*

The above six transformation steps establish the correctness of our reduction. App. B gives a direct construction of the reduced finite-state and finite-counter system, $\mathsf{RDTS}^{\mathsf{RIA}}(\mathcal{T},\nu,\varphi)$ and $\mathsf{PRDTS}^{\mathsf{RIA}}(\mathcal{T},\varphi)$, from the input template $\mathcal{T}$ and formula φ.

4 Case Studies

We evaluate the practical usefulness of our reduction on four case studies: (i) Ben-Or's consensus with crash faults [3, Algo A], (ii) Ben-Or's consensus with Byzantine faults [3, Algo B], (iii) Bracha's consensus with Byzantine faults [9, Fig. 4], and (iv) Raft's leader-election [35, Sec. 5.2]. The GitHub repository [21] provides source files and instructions to reproduce all experiments. For each case study, the repository includes an `.smv` file whose initial commented block specifies the corresponding process template, followed by the RIA semantics used as input to nuXmv. The models abstract randomization by nondeterminism; as a result, termination is no longer guaranteed, although the remaining properties are unaffected. Byzantine faults are modeled explicitly using dedicated fault locations, from which faulty processes may nondeterministically broadcast messages of arbitrary types. The model checker nuXmv is called via its IC3 engine (`check_ltlspec_ic3`) to check whether the translated LTL specifications hold on the RIA semantics. Table 2 reports, for each case study and property, the verification time and the maximal IC3 depth. Across all experiments, IC3 reaches shallow depths and terminates in under 15 seconds. To stress-test our reduction, we introduced intentional faults into the algorithm models by weakening the resilience condition or altering key guard predicates. As expected, for these faulty variants nuXmv produced counterexamples within seconds.

Table 2 Case-study results with the IC3 engine of nuXmv. Columns b, $|\mathcal{L}|$, |Rules|, and rc give the jump bound, number of locations, number of rules, and resilience condition of the process template. For each property in Fig. 3, the rest of columns report verification time in seconds, with the maximal IC3 depth in parentheses. For T, the entry reports the time and depth to find a counterexample.

Protocol	b	$\lvert\mathcal{L}\rvert$	|Rules|	rc	**A**	**V**	**RT**	**T**	**LU**
Ben-Or (Crash)	1	9	26	$n > 2t$	1.4s (13)	0.4s (9)	3.1s (8)	0.5s (3)	–
Ben-Or (Byzantine)	1	10	27	$n > 5t$	7.0s (11)	1.2s (7)	4.3s (7)	0.6s (3)	–
Bracha (Byzantine)	1	12	31	$n > 3t$	14.0s (14)	1.8s (8)	6.5s (11)	0.7s (3)	–
Raft leader election	2	11	25	$n > 2t$	–	–	–	–	1.8s (8)

5 Conclusion

Asynchronous round-based distributed algorithms underpin many modern distributed systems, including many consensus algorithms at the core of blockchain technologies [44]. Our reduction offers a formal step toward their verification via mature symbolic model checkers and suggests a number of challenging extensions.

First, an important direction is to extend our reduction to incorporate probabilistic behavior and so allow verification of almost-sure termination. While existing probabilistic model checkers [29, 15] can efficiently verify almost-sure reachability conditions on *finite* MDPs, obtaining such a model is far from immediate, even for a fixed parameter valuation. The main challenge arises in the SSR step as a probabilistic semantics requires to reorder not only a trace, but the whole execution tree.

Second, a compelling challenge is to extend our reduction to handle algorithms whose number of locations within a round grows with the round index, as is the case for variants of Paxos [30, 31] and DAG-based consensus [2, 23, 24, 19, 18] (where the number of locations also grows with the number of processes). Interestingly, the reduction we presented in this paper does not assume that the set of locations within each round is finite; however, the reduced system would then have infinitely many counters.

References

1 Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking*. The MIT Press, 2008.

2 Leemon Baird and Atul Luykx. The hashgraph protocol: Efficient asynchronous BFT for high-throughput distributed ledgers. In *Proceedings of the 2020 International Conference on Omni-layer Intelligent Systems (COINS 2020)*, pages 1–7. IEEE, 2020. doi:10.1109/COINS49042.2020.9191430.

3 Michael Ben-Or. Another advantage of free choice (extended abstract): Completely asynchronous agreement protocols. In *Procedings of the 2nd Annual ACM Symposium on Principles of Distributed Computing (PODC 1983)*, pages 27–30. Association for Computing Machinery, 1983. doi:10.1145/800221.806707.

4 Nathalie Bertrand, Pranav Ghorpade, and Sasha Rubin. Parameterized verification of asynchronous round-based distributed algorithms via reduction to finite-counter systems, 2026. arXiv:2606.27867.

5 Nathalie Bertrand, Pranav Ghorpade, Sasha Rubin, Bernhard Scholz, and Pavle Subotić. Reusable formal verification of dag-based consensus protocols. In *Proceedings of the 17th International Symposium on NASA Formal Methods (NFM 2025)*, pages 138–158. Springer-Verlag, 2025. doi:10.1007/978-3-031-93706-4_9.

6 Nathalie Bertrand, Igor Konnov, Marijana Lazić, and Josef Widder. Verification of Randomized Consensus Algorithms Under Round-Rigid Adversaries. In *Proceedings of the 30th International Conference on Concurrency Theory (CONCUR 2019)*, volume 140 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 33:1–33:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.CONCUR.2019.33.

7 Nathalie Bertrand, Bastien Thomas, and Josef Widder. Guard Automata for the Verification of Safety and Liveness of Distributed Algorithms. In *Proceeding of the 32nd International Conference on Concurrency Theory (CONCUR 2021)*, volume 203 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 15:1–15:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.CONCUR.2021.15.

8 Roderick Bloem, Swen Jacobs, Ayrat Khalimov, Igor Konnov, Sasha Rubin, Helmut Veith, and Josef Widder. *Decidability of Parameterized Verification*. Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool Publishers, 2015. doi:10.2200/S00658ED1V01Y201508DCT013.

9 Gabriel Bracha. Asynchronous Byzantine agreement protocols. *Information and Computation*, 75(2):130–143, 1987. doi:10.1016/0890-5401(87)90054-X.

10 Roberto Cavada, Alessandro Cimatti, Michele Dorigatti, Alberto Griggio, Alessandro Mariotti, Andrea Micheli, Sergio Mover, Marco Roveri, and Stefano Tonetta. The nuXmv symbolic model checker. In *Proceedings of the 26th International Conference on Computer Aided Verification (CAV 2014)*, pages 334–342. Springer International Publishing, 2014. doi:10.1007/978-3-319-08867-9_22.

11 Saksham Chand, Yanhong A. Liu, and Scott D. Stoller. Formal verification of multi-paxos for distributed consensus. In *Proceedings of the 21st International Symposium on Formal Methods (FM 2016)*, volume 9995 of *Lecture Notes in Computer Science*, pages 119–136, 2016. doi:10.1007/978-3-319-48989-6_8.

12 Tushar Deepak Chandra and Sam Toueg. Unreliable failure detectors for reliable distributed systems. *J. ACM*, 43(2):225–267, March 1996. doi:10.1145/226643.226647.

13 Mouna Chaouch-Saad, Bernadette Charron-Bost, and Stephan Merz. A reduction theorem for the verification of round-based distributed algorithms. In *Proceedings of the 3rd International Workshop on Reachability Problems (RP 2009)*, pages 93–106. Springer Berlin Heidelberg, 2009. doi:10.1007/978-3-642-04420-5_10.

14 Bernadette Charron-Bost and André Schiper. The heard-of model: computing in distributed systems with benign faults. *Distributed Computing*, 22(1):49–71, April 2009. doi:10.1007/s00446-009-0084-6.

15 Christian Dehnert, Sebastian Junges, Joost-Pieter Katoen, and Matthias Volk. A storm is coming: A modern probabilistic model checker, 2017. arXiv:1702.04311.

16 Giorgio Delzanno, Michele Tatarek, and Riccardo Traverso. Model checking paxos in spin. In *Proceedings 5th International Symposium on Games, Automata, Logics and Formal Verification (GandALF 2014)*, volume 161 of *EPTCS*, pages 131–146, 2014. doi:10.4204/EPTCS.161.13.

17 Stéphane Demri, Alain Finkel, Valentin Goranko, and Govert van Drimmelen. Model-checking ctl* over flat presburger counter systems. *J. Appl. Non Class. Logics*, 20(4):313–344, 2010. doi:10.3166/JANCL.20.313-344.

18 Fantom Foundation. Lachesis aBFT, 2024. URL: https://docs.fantom.foundation/technology/lachesis-abft.

19 Adam Gagol, Damian Lesniak, Damian Straszak, and Michal Swietek. Aleph: Efficient atomic broadcast in asynchronous networks with Byzantine nodes. In *Proceedings of the 1st ACM Conference on Advances in Financial Technologies (AFT 2019)*, pages 214–228. ACM, 2019. doi:10.1145/3318041.3355467.

20 Steven M. German and A. Prasad Sistla. Reasoning about systems with many processes. *Journal of the ACM*, 39(3):675–735, 1992. doi:10.1145/146637.146681.

21 Pranav Ghorpade. Github case studies repository, 2026. URL: https://github.com/pranavg5526/Parameterized-Verification-of-Round-Based-Distributed-Algorithms.

22 Aman Goel and Karem Sakallah. On symmetry and quantification: A new approach to verify distributed protocols. In *Proceedings of the 13th International Symposium on NASA Formal Methods (NFM 2021)*, pages 131–150, Berlin, Heidelberg, 2021. Springer-Verlag. doi:10.1007/978-3-030-76384-8_9.

23 Idit Keidar, Eleftherios Kokoris-Kogias, Oded Naor, and Alexander Spiegelman. All you need is DAG. In *Proceedings of 40th ACM Symposium on Principles of Distributed Computing (PODC 2021)*, pages 165–175. ACM, 2021. doi:10.1145/3465084.3467905.

24 Idit Keidar, Oded Naor, Ouri Poupko, and Ehud Shapiro. Cordial miners: Fast and efficient consensus for every eventuality. In *Proceedings of the 37th International Symposium on Distributed Computing, (DISC 2023)*, volume 281 of *LIPIcs*, pages 26:1–26:22. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.DISC.2023.26.

25 Jason R. Koenig, Oded Padon, Sharon Shoham, and Alex Aiken. Inferring invariants with quantifier alternations: Taming the search space explosion. In *TACAS (1)*, volume 13243 of *Lecture Notes in Computer Science*, pages 338–356. Springer, 2022. doi:10.1007/978-3-030-99524-9_18.

26 Igor Konnov, Marijana Lazic, Ilina Stoilkovska, and Josef Widder. Survey on parameterized verification with threshold automata and the Byzantine model checker. *Logical Methods in Computer Science*, 19(1), 2023. doi:10.46298/LMCS-19(1:5)2023.

27 Igor Konnov, Marijana Lazić, Helmut Veith, and Josef Widder. A short counterexample property for safety and liveness verification of fault-tolerant distributed algorithms. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages (POPL 2017)*, pages 719–734. Association for Computing Machinery, 2017. doi:10.1145/3009837.3009860.

28 Igor Konnov, Helmut Veith, and Josef Widder. SMT and POR beat counter abstraction: Parameterized model checking of threshold-based distributed algorithms. In *CAV (1)*, volume 9206 of *Lecture Notes in Computer Science*, pages 85–102. Springer, 2015. doi:10.1007/978-3-319-21690-4_6.

29 M. Kwiatkowska, G. Norman, and D. Parker. PRISM 4.0: Verification of probabilistic real-time systems. In *Proceedings of the 23rd International Conference on Computer Aided Verification (CAV 2011)*, volume 6806 of *LNCS*, pages 585–591. Springer, 2011. doi:10.1007/978-3-642-22110-1_47.

30 Leslie Lamport. The part-time parliament. *ACM Transactions on Computer Systems*, 16(2):133–169, 1998. doi:10.1145/279227.279229.

31 Leslie Lamport. Fast Paxos. *Distributed Computing*, 19(2):79–103, 2006. doi:10.1007/s00446-006-0005-x.

32 Nancy Lynch and Frits Vaandrager. Forward and backward simulations i.: untimed systems. *Information and Computation*, 121(2):214–233, 1995. doi:10.1006/inco.1995.1134.

33 Zohar Manna and Amir Pnueli. A hierarchy of temporal properties. In *PODC*, pages 377–410. ACM, 1990. doi:10.1145/93385.93442.

34 Tatsuya Noguchi, Tatsuhiro Tsuchiya, and Tohru Kikuno. Safety verification of asynchronous consensus algorithms with model checking. In *IEEE 18th Pacific Rim International Symposium on Dependable Computing (PRDC 2012)*, pages 80–88. IEEE Computer Society, 2012. doi:10.1109/PRDC.2012.24.

35 Diego Ongaro and John Ousterhout. In search of an understandable consensus algorithm. In *Proceedings of the 2014 USENIX Annual Technical Conference*, pages 305–320, 2014. URL: https://www.usenix.org/conference/atc14/technical-sessions/presentation/ongaro.

36 Oded Padon, Kenneth L. McMillan, Aurojit Panda, Shmuel Sagiv, and Sharon Shoham. Ivy: safety verification by interactive generalization. *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2016)*, 2016. URL: https://api.semanticscholar.org/CorpusID:17141487.

37 William Schultz, Edward Ashton, Heidi Howard, and Stavros Tripakis. Scalable, interpretable distributed protocol verification by inductive proof slicing. *CoRR*, abs/2404.18048, 2024. doi:10.48550/arXiv.2404.18048.

38 Yee Jiun Song and Robbert van Renesse. Bosco: One-Step Byzantine Asynchronous Consensus. In Gadi Taubenfeld, editor, *Distributed Computing*, pages 438–450, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg. doi:10.1007/978-3-540-87779-0_30.

39 T. K. Srikanth and Sam Toueg. Simulating authenticated broadcasts to derive simple fault-tolerant algorithms. *Distributed Comput.*, 2(2):80–94, 1987. doi:10.1007/BF01667080.

40 Ilina Stoilkovska, Igor Konnov, Josef Widder, and Florian Zuleger. Verifying safety of synchronous fault-tolerant algorithms by bounded model checking. In Tomás Vojnar and Lijun Zhang, editors, *Proceedings of the 25th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2019)*, Lecture Notes in Computer Science, pages 357–374. Springer, 2019. doi:10.1007/978-3-030-17465-1_20.

41 Ilina Stoilkovska, Igor Konnov, Josef Widder, and Florian Zuleger. Eliminating message counters in threshold automata. In *Proceedings of the 18th International Symposium on Automated Technology for Verification and Analysis (ATVA 2020)*, pages 196–212. Springer-Verlag, 2020. doi:10.1007/978-3-030-59152-6_11.

42 Bastien Thomas and Ocan Sankur. PyLTA: A verification tool for parameterized distributed algorithms. In *Proceedings of the 29th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2023)*, pages 28–35. Springer Nature Switzerland, 2023. doi:10.1007/978-3-031-30820-8_4.

43 Doug Woos, James R. Wilcox, Steve Anton, Zachary Tatlock, Michael D. Ernst, and Thomas Anderson. Planning for change in a formal verification of the Raft consensus protocol. In *Proceedings of the 5th ACM SIGPLAN Conference on Certified Programs and Proofs (CPP 2016)*, pages 154–165. ACM, 2016. doi:10.1145/2854065.2854081.

44 Yang Xiao, Ning Zhang, Wenjing Lou, and Y. Thomas Hou. A survey of distributed consensus protocols for blockchain networks. *IEEE Communications Surveys & Tutorials*, 22(2):1432–1465, 2020. doi:10.1109/COMST.2020.2969706.

45 Jianan Yao, Runzhou Tao, Ronghui Gu, Jason Nieh, Suman Jana, and Gabriel Ryan. DistAI: Data-Driven automated invariant learning for distributed protocols. In *15th USENIX Symposium on Operating Systems Design and Implementation (OSDI 21)*, pages 405–421. USENIX Association, July 2021. URL: https://www.usenix.org/conference/osdi21/presentation/yao.

A Detailed fixed-instance and parameterized semantics

A.1 Fixed-instance semantics

Let $\mathcal{T} = \langle \mathsf{P}, \mathsf{rc}, \mathcal{L}, \mathcal{I}, \mathcal{M}, \mathsf{Bcast}, \mathsf{Rules} \rangle$ be a template and $\nu : \mathsf{P} \to \mathbb{N}$ a *parameter valuation* such that $\nu \models \mathsf{rc}$. In the system consisting of $\nu(n)$ identical processes, each process execute the behavior specified by $\langle \mathcal{L}, \mathcal{I}, \mathcal{M}, \mathsf{Bcast}, \mathsf{Rules} \rangle$.

The *fixed-instance semantics* of the template $\mathcal{T}$ for a parameter valuation ν, denoted $\mathsf{RDTS}(\mathcal{T}, \nu)$, is an action-labeled transition system: $\mathsf{RDTS}(\mathcal{T}, \nu) = \langle S_\nu, I_\nu, Act_\nu, Tr_\nu \rangle$, where S_ν denotes the set of *configurations*, $I_\nu \subseteq S_\nu$ the set of *initial configurations*, Act_ν the *action set*, and $Tr_\nu \subseteq S_\nu \times Act_\nu \times S_\nu$ the *transition relation*.

Configurations. A *configuration* is a pair $\mathsf{cfg} = \langle \mathsf{PState}, \mathsf{NState} \rangle$ consisting of a process state and a network state. The process state is a function $\mathsf{PState} : [\![1, \nu(n)]\!] \to \langle \mathsf{loc} : \mathcal{L}, \mathsf{rd} : \mathbb{N}, \mathsf{rcvMsg} : \mathcal{M} \times \mathbb{N} \to \mathbb{N} \rangle$ that assigns to each process a location in $\mathcal{L}$, a round number in $\mathbb{N}$, and a multiset of received messages without sender identities.

The network state is a function $\mathsf{NState} : \mathcal{M} \times \mathbb{N} \to \mathbb{N}$ that records how many messages of each type and round tag were broadcast, irrespective of the sender's identity. The acyclicity assumption in Definition 1 ensures that no process visits the same location more than once per round. Consequently, each process broadcasts at most $|\mathcal{L}|$ messages in a round, and the range of NState is therefore bounded above by $\nu(n) \times |\mathcal{L}|$. Moreover, if all processes are in a round with number less than r (formally, $\forall i : \mathsf{PState}(i).\mathsf{rd} < r$), then $\mathsf{NState}(\cdot, r)$ must be the zero function $\mathbf{0}$, mapping every argument to 0. Finally, in every configuration, the number of received messages is *pointwise bounded* by the number of broadcast messages:

$$\forall i,\ \forall (m, r) \in \mathcal{M} \times \mathbb{N} : \quad \mathsf{PState}(i).\mathsf{rcvMsg}(m, r) \ \leq \ \mathsf{NState}(m, r). \tag{1}$$

A configuration $\langle \mathsf{PState}, \mathsf{NState} \rangle$ is *initial* if for all $i \in [\![1, \nu(n)]\!]$, $\mathsf{PState}(i).\mathsf{loc} \in \mathcal{I}$, $\mathsf{PState}(i).\mathsf{rd} = 0$, and $\mathsf{PState}(i).\mathsf{rcvMsg} = \mathbf{0}$; and $\mathsf{NState} = \mathbf{0}$.

Actions and Transitions. A process $i \in [\![1, \nu(n)]\!]$ can perform two types of actions: (i) $\mathsf{Recv}(i, \langle m, r \rangle)$, which represents the reception of a message of type m tagged with round r; (ii) $\mathsf{Update}(i, \rho, r)$, which represents updating its location by applying rule $\rho \in \mathsf{Rules}$ at round $r \in \mathbb{N}$. These form the action set Act_ν, in which each action represents a step executed by a single process. The transition relation Tr_ν consists of all triples $(\mathsf{cfg}, a, \mathsf{cfg}')$ such that $a \in Act_\nu$ is *enabled* in the configuration $\mathsf{cfg} \in S_\nu$ (denoted $\mathsf{enabled}(\mathsf{cfg}, a)$), and executing a at cfg yields cfg' (denoted $\mathsf{effect}(\mathsf{cfg}, a) = \mathsf{cfg}'$). We now formally define $\mathsf{enabled}$ and effect for *receive* and *update* actions.

Receive. In a configuration $\mathsf{cfg} = \langle \mathsf{PState}, \mathsf{NState} \rangle$, the action $\mathsf{Recv}(i, \langle m, r \rangle)$ is enabled when process i has not yet received all broadcasts of type m tagged with round r. Formally,

$$\mathsf{enabled}(\mathsf{cfg}, \mathsf{Recv}(i, \langle m, r \rangle)) := \mathsf{PState}(i).\mathsf{rcvMsg}(m, r) < \mathsf{NState}(m, r).$$

Executing $\mathsf{Recv}(i, \langle m, r \rangle)$ delivers one message of type m tagged with round r to process i and leaves all other components unchanged. Formally,
$\mathsf{effect}(\mathsf{cfg}, \mathsf{Recv}(i, \langle m, r \rangle)) := \langle \mathsf{PState}', \mathsf{NState}' \rangle$, where:

- For every process $j \neq i$, the state remains unchanged: $\mathsf{PState}'(j) = \mathsf{PState}(j)$.
- Location and round component of process i state remain unchanged:
  $\mathsf{PState}'(i).\mathsf{loc} = \mathsf{PState}(i).\mathsf{loc}$ and $\mathsf{PState}'(i).\mathsf{rd} = \mathsf{PState}(i).\mathsf{rd}$.
- The receive component of process i state is incremented at $\langle m, r \rangle$:
  $\mathsf{PState}'(i).\mathsf{rcvMsg}(m, r) = \mathsf{PState}(i).\mathsf{rcvMsg}(m, r) + 1$ and for all
  $\langle m', r' \rangle \neq \langle m, r \rangle$, $\mathsf{PState}'(i).\mathsf{rcvMsg}(m', r') = \mathsf{PState}(i).\mathsf{rcvMsg}(m', r')$.
- The network state remains unchanged: $\mathsf{NState}' = \mathsf{NState}$.

Update. In a configuration $\mathsf{cfg} = \langle \mathsf{PState}, \mathsf{NState} \rangle$, the action $\mathsf{Update}(i, \rho, r)$ is enabled when process i is in location $\rho.\mathsf{frm}$ at round r, and its received messages in round r satisfy the guard $\rho.\mathsf{guard}$. Formally,

$$\begin{aligned}\mathsf{enabled}(\mathsf{cfg}, \mathsf{Update}(i, \rho, r)) := {} & \mathsf{PState}(i).(\mathsf{loc}, \mathsf{rd}) = (\rho.\mathsf{frm}, r) \text{ and} \\ & \mathsf{PState}(i).\mathsf{rcvMsg}(\_, r) \models \rho.\mathsf{guard}[\mathsf{P} \leftarrow \nu].\end{aligned}$$

Executing $\mathsf{Update}(i, \rho, r)$ has two effects. First, process i moves from location $\rho.\mathsf{frm}$ to $\rho.\mathsf{to}$ and updates its round counter to $r + \rho.\mathsf{type}$. Second, it broadcasts a message of type $\mathsf{Bcast}(\rho.\mathsf{to})$, tagged with round $r + \rho.\mathsf{type}$. Formally,

$\mathsf{effect}(\mathsf{cfg}, \mathsf{Update}(i, \rho, r)) = \langle \mathsf{PState}', \mathsf{NState}' \rangle$, where:

- For every process $j \neq i$, the state remains unchanged: $\mathsf{PState}'(j) = \mathsf{PState}(j)$.
- The location component of process i is updated to $\rho.\mathsf{to}$:
  $\mathsf{PState}'(i).\mathsf{loc} = \rho.\mathsf{to}$.
- The round component of process i state is incremented by $\rho.\mathsf{type}$:
  $\mathsf{PState}(i)'.\mathsf{rd} = r + \rho.\mathsf{type}$
- The receive component of process i remains unchanged:
  $\mathsf{PState}'(i).\mathsf{rcvMsg}(m, r) = \mathsf{PState}(i).\mathsf{rcvMsg}(m, r)$.
- The network state is incremented at $\langle \mathsf{Bcast}(\rho.\mathsf{to}), r + \rho.\mathsf{type} \rangle$:
  $\mathsf{NState}'(m', r') = \mathsf{NState}(m', r') + 1$ if $\langle m', r' \rangle = \langle \mathsf{Bcast}(\rho.\mathsf{to}), r + \rho.\mathsf{type} \rangle$, and
  $\mathsf{NState}'(m', r') = \mathsf{NState}(m', r')$ otherwise.

A.2 Parameterized semantics

The parameterized semantics is a disjoint union of all fixed-instance semantics $\mathsf{RDTS}(\mathcal{T}, \nu)$ for $\nu \models \mathsf{rc}$. Fix a template $\mathcal{T}$, and let $\mathsf{Val} = \{\nu : \mathsf{P} \to \mathbb{N} \mid \nu \models \mathsf{rc}\}$ denote the set of all admissible parameter valuations. The parameterized semantics of $\mathcal{T}$ is an action-labeled transition system, called the *parameterized* RDTS (PRDTS), defined as $\mathsf{PRDTS}(\mathcal{T}) = \langle S,\ I,\ Act,\ Tr \rangle$, which is the *disjoint union* of $\mathsf{RDTS}(\mathcal{T}, \nu)$ over all $\nu \in \mathsf{Val}$, i.e.:

$$S = \bigcup_{\nu \in \mathsf{Val}} \{\nu\} \times S_\nu, \quad I = \bigcup_{\nu \in \mathsf{Val}} \{\nu\} \times I_\nu, \quad Act = \bigcup_{\nu \in \mathsf{Val}} \{\nu\} \times Act_\nu,$$

$$Tr = \left\{ (\nu, \mathsf{cfg}) \xrightarrow{(\nu, a)} (\nu, \mathsf{cfg}') \;\middle|\; \nu \in \mathsf{Val},\ (\mathsf{cfg}, a, \mathsf{cfg}') \in Tr_\nu \right\}.$$

For every execution $\pi = (\nu, \mathsf{cfg}_0)(\nu, a_1)(\nu, \mathsf{cfg}_1)\dots$ of $\mathsf{PRDTS}(\mathcal{T})$, let $\pi_{\downarrow_\nu} = \mathsf{cfg}_0 a_1 \mathsf{cfg}_1 \dots$. Clearly $\pi_{\downarrow_\nu}$ is an execution of $\mathsf{RDTS}(\mathcal{T}, \nu)$. Moreover, for every execution π of $\mathsf{RDTS}(\mathcal{T}, \nu)$, there exists execution π' of $\mathsf{PRDTS}(\mathcal{T})$ such that $\pi = \pi'_{\downarrow_\nu}$.

A.3 Finite-Counter Systems

We use a notion of finite-counter system that can be directly represented as a nuXmv model. A *finite-counter system* $\mathcal{C}$ is a tuple $(V, \mathsf{INVAR}, \mathsf{INIT}, \mathsf{TRAN})$, where $V = \{x_1, \dots, x_k\}$ is a finite set of variables (also called counters) ranging over $\mathbb{N}$, INVAR and INIT are linear-arithmetic formulas over $x_1, \dots, x_k$, and TRAN is a linear-arithmetic formula over $x_1, \dots, x_k, x'_1, \dots, x'_k$. Here, x'_i denotes the next-state copy of x_i.

A finite-counter system $\mathcal{C} = (V, \mathsf{INVAR}, \mathsf{INIT}, \mathsf{TRAN})$ induces a transition system $\mathcal{S}_C := (Q, Q_0, \to)$, where Q is the set of valuations $s : V \to \mathbb{N}$ satisfying INVAR, $Q_0 \subseteq Q$ is the set of valuations satisfying INIT, and $\to \subseteq Q \times Q$ contains exactly the pairs (s, s') satisfying TRAN, i.e., those for which TRAN is true after substituting each x_i by $s(x_i)$ and each x'_i by $s'(x_i)$. Intuitively, INVAR defines the admissible counter values, INIT defines the admissible initial counter values, and TRAN relates current counter values to their possible next values.

B Explicit Encoding into Finite-Counter System

This appendix gives the explicit finite-counter-system construction obtained after the reduction presented in Section 3. The input is a process template $\mathcal{T} = \langle \mathsf{P}, \mathsf{rc}, \mathcal{L}, \mathcal{I}, \mathcal{M}, \mathsf{Bcast}, \mathsf{Rules} \rangle$ and an HSCL formula φ. For fixed-instance verification, the input also includes a parameter valuation $\nu : \mathsf{P} \to \mathbb{N}$.

The output is a finite-counter system $(V, \mathsf{INVAR}, \mathsf{INIT}, \mathsf{TRAN})$ together with the LTL formula obtained from φ, whose atomic propositions are linear-arithmetic formulas over V. Thus, the construction produces an instance of LTL model checking over a finite-counter system, which can be encoded in symbolic model checkers such as nuXmv.

In the fixed-instance case, the invariant formula INVAR additionally fixes the parameter variables according to ν. This implies that only finitely many valuations satisfy INVAR. Hence, for fixed-instance verification, the generated finite-counter system is in fact finite-state.

For a rule $\rho \in \mathsf{Rules}$, write $\rho = (\ell, \ell', k, g)$, where $\ell = \rho.\mathsf{from}$, $\ell' = \rho.\mathsf{to}$, $k = \rho.\mathsf{type}$, and $g = \rho.\mathsf{guard}$. Let $b = \max\{\rho.\mathsf{type} \mid \rho \in \mathsf{Rules}\}$. Further, let $p_1, \dots, p_q$ be the parameter variables from P that occur in g, and let $m_1, \dots, m_s$ be the message variables from $\mathcal{M}$ that occur in g. Let $\widehat{g}(p_1, \dots, p_q, m_1, \dots, m_s)$ be a quantifier-free Presburger formula equivalent to $\exists r_1, \dots, r_s.\ (\bigwedge_{i=1}^{s} 0 \le r_i \le m_i) \wedge g[m_1 \mapsto r_1, \dots, m_s \mapsto r_s]$.

Let $X \subseteq \mathcal{L}$ be the set of locations mentioned in φ and let K be one plus the largest threshold appearing in cumulative atoms of φ, see Def. 2 ($K = 0$ if no such atom exists).

Below, we describe each of the components of the obtained finite-counter system and the LTL formula, including some brief explanations to help with understanding.

Variables V

The set V consists of the following non-negative integer variables: par_p ($p \in \mathsf{P}$); $\mathsf{loc}_{\ell,d}$ ($\ell \in \mathcal{L},\ d \in \{0, \dots, b-1\}$); $\mathsf{msg}_{m,d}$ ($m \in \mathcal{M},\ d \in \{0, \dots, b-1\}$); and $\mathsf{local}_\ell, \mathsf{cumul}_\ell$ ($\ell \in X$). These are called, respectively, *parameter variables*, *location variables*, *message variables*, and *history variables*.

Invariant formula INVAR

The formula INVAR is the conjunction of the constraints below.

1. The parameter variables satisfy the resilience condition: $\mathsf{rc}\big[p \mapsto \mathsf{par}_p\big]$. For fixed-instance verification with parameter valuation $\nu : \mathsf{P} \to \mathbb{N}$, INVAR additionally includes the constraints $\mathsf{par}_p = \nu(p)$ $(p \in \mathsf{P})$.
2. For $H = n$ when failures are modeled explicitly by dedicated fault locations and $H = n - t$ when failures are modeled implicitly by allowing processes to stop participating, the location variables are bounded by the total number of processes: $H \leq \sum_{\ell \in \mathcal{L}} \sum_{d=0}^{b-1} \mathsf{loc}_{\ell,d} \leq \mathsf{par}_n$.
3. The message variables are bounded by the maximum number of messages that can be broadcast in one round: $0 \leq \mathsf{msg}_{m,d} \leq |\mathcal{L}| \cdot \mathsf{par}_n$ $(m \in \mathcal{M},\ d \in \{0, \ldots, b-1\})$.
4. The history variables are bounded as follows: $(0 \leq \mathsf{local}_\ell \leq \mathsf{par}_n) \wedge (0 \leq \mathsf{cumul}_\ell \leq K)$ $(\ell \in X)$.

Initial formula INIT

The formula INIT is the conjunction of the constraints below.

1. The initial location variables satisfy $\sum_{\ell \in \mathcal{I}} \mathsf{loc}_{\ell,0} = \mathsf{par}_n$; $\mathsf{loc}_{\ell,0} = 0$ $(\ell \notin \mathcal{I})$; and $\mathsf{loc}_{\ell,d} = 0$ $(\ell \in \mathcal{L},\ d \in \{1, \ldots, b-1\})$.
2. The initial message variables satisfy $\mathsf{msg}_{m,d} = 0$ $(m \in \mathcal{M},\ d \in \{0, \ldots, b-1\})$.
3. The initial history variables satisfy $\mathsf{local}_\ell = \mathsf{loc}_{\ell,0}$ and $\mathsf{cumul}_\ell = \mathsf{loc}_{\ell,0}$ if $\ell \in \mathcal{I}$ else $\mathsf{local}_\ell = 0$ and $\mathsf{cumul}_\ell = 0$.

Transition formula TRAN

The transition formula is the disjunction of the local-transition formula and the jump-transition formula: $\mathsf{TRAN} = \mathsf{TRAN}_{\mathsf{loc}} \vee \mathsf{TRAN}_{\mathsf{jump}}$.

Local-transition formula $\mathsf{TRAN}_{\mathsf{loc}}$. Intuitively, the formula $\mathsf{TRAN}_{\mathsf{loc}}$ encodes one process taking a type-0 rule in the current frontier round. For each type-0 rule $\rho = (\ell, \ell', 0, g)$, we define a clause $\mathsf{LocClause}_\rho$, and set $\mathsf{TRAN}_{\mathsf{loc}} = \bigvee_{\rho=(\ell,\ell',0,g) \in \mathsf{Rules}} \mathsf{LocClause}_\rho$. The clause $\mathsf{LocClause}_\rho$ is the conjunction of the following constraints.

1. The enabledness constrain: $\mathsf{loc}_{\ell,0} > 0 \ \wedge\ \widehat{g}\big[p \mapsto \mathsf{par}_p,\ m \mapsto \mathsf{msg}_{m,0}\big]$.
2. The location variables are updated by moving one process from ℓ to ℓ' at depth 0: $\mathsf{loc}'_{\ell,0} = \mathsf{loc}_{\ell,0} - 1$; $\mathsf{loc}'_{\ell',0} = \mathsf{loc}_{\ell',0} + 1$; and all other location variables are unchanged.
3. If $\mathsf{Bcast}(\ell') = m \in \mathcal{M}$, then the corresponding message variable at depth 0 is incremented: $\mathsf{msg}'_{m,0} = \mathsf{msg}_{m,0} + 1$, and all other message variables are unchanged. If $\mathsf{Bcast}(\ell') = \bot$, then all message variables are unchanged.
4. If $\ell' \in X$, then the history variables for ℓ' is incremented as follows: $\mathsf{local}'_{\ell'} = \mathsf{local}_{\ell'} + 1$; $\mathsf{cumul}'_{\ell'} = \mathsf{cumul}_{\ell'} + 1$ if $\mathsf{cumul}_{\ell'} < K$ else $\mathsf{cumul}'_{\ell'} = \mathsf{cumul}_{\ell'}$; and all other history variables are unchanged. If $\ell' \notin X$, then all history variables are unchanged.
5. Finally, all parameter variables are unchanged.

Jump-transition formula $\mathsf{TRAN}_{\mathsf{jump}}$. Intuitively, the formula $\mathsf{TRAN}_{\mathsf{jump}}$ encodes a jump of the frontier round. For each jump length $h \in \{1, \ldots, b\}$, we define a clause $\mathsf{JumpClause}_h$, and set $\mathsf{TRAN}_{\mathsf{jump}} = \bigvee_{h=1}^{b} \mathsf{JumpClause}_h$. The clause $\mathsf{JumpClause}_h$ uses auxiliary non-negative integer variables $J_{d,\rho}$ $(d \in \{0, \ldots, b-1\},\ \rho \in \mathsf{Rules},\ \rho.\mathsf{type} > 0)$. The variable $J_{d,\rho}$ denotes the number of processes at depth d that take the positive-type rule ρ during the jump. Formally, $\mathsf{JumpClause}_h$ can be viewed as existentially quantifying the auxiliary variables $J_{d,\rho}$;

the resulting formula can then be made linear-arithmetic by Presburger quantifier elimination. In our encoding, we use the equivalent and more direct presentation: the variables $J_{d,\rho}$ are included in V as auxiliary variables and are updated nondeterministically by jump transitions, so that their current values serve as witnesses for the jump.

The clause $\mathsf{JumpClause}_h$ is the conjunction of the following constraints. First for every $\lambda \in \mathcal{L}$ and $d \in \{0, \dots, b-1\}$ let

$$J^{\mathsf{out}}_{\lambda,d} = \sum_{\substack{\rho=(\lambda,\ell',k,g)\in \mathsf{Rules} \\ k>0}} J_{d,\rho}, \qquad J^{\mathsf{in}}_{\lambda} = \sum_{\substack{d\in\{0,\dots,b-1\} \\ \rho=(\ell,\lambda,k,g)\in \mathsf{Rules} \\ k>0}} J_{d,\rho}.$$

1. The auxiliary variables cannot move more processes than are available: $J^{\mathsf{out}}_{\lambda,d} \le \mathsf{loc}_{\lambda,d}$ $(\lambda \in \mathcal{L},\ d \in \{0,\dots,b-1\})$.
2. A positive-type rule may be taken only if it lands exactly in the new frontier round. Thus, for every $\rho = (\ell, \ell', k, g)$ with $k > 0$, and every $d \in \{0, \dots, b-1\}$, we require $J_{d,\rho} > 0 \implies d + h = k$.
3. Whenever some processes take a rule, its guard must be enabled at the corresponding depth: $J_{d,\rho} > 0 \implies \widehat{g}\big[p \mapsto \mathsf{par}_p,\ m \mapsto \mathsf{msg}_{m,d}\big]$.
4. The location variables are updated by shifting the old window by length h and placing the processes that jump in the new frontier round: $\mathsf{loc}'_{\lambda,0} = J^{\mathsf{in}}_{\lambda}$ $(\lambda \in \mathcal{L})$. For depths skipped by the jump, we set $\mathsf{loc}'_{\lambda,d} = 0$ $(\lambda \in \mathcal{L},\ 1 \le d < h)$. For the remaining depths, we shift the old variables and remove the processes that jumped: $\mathsf{loc}'_{\lambda,d} = \mathsf{loc}_{\lambda,d-h} - J^{\mathsf{out}}_{\lambda,d-h}$ $(\lambda \in \mathcal{L},\ h \le d \le b-1)$.
5. The message variables are updated analogously. The messages in the new frontier round are exactly those broadcast by the processes that jump into that round. Thus, for every $m \in \mathcal{M}$, $\mathsf{msg}'_{m,0} = \sum_{\lambda\in\mathcal{L},\ \mathsf{Bcast}(\lambda)=m} J^{\mathsf{in}}_{\lambda}$. For depths skipped by the jump, we set $\mathsf{msg}'_{m,d} = 0$ $(m \in \mathcal{M},\ 1 \le d < h)$. For the remaining depths, we shift the old message variables: $\mathsf{msg}'_{m,d} = \mathsf{msg}_{m,d-h}$ $(m \in \mathcal{M},\ h \le d \le b-1)$.
6. The history variables are updated according to the processes that enter locations mentioned in the formula. For every $\lambda \in X$, $\mathsf{local}'_{\lambda} = J^{\mathsf{in}}_{\lambda}$; $\mathsf{cumul}'_{\lambda} = \mathsf{cumul}_{\lambda} + J^{\mathsf{in}}_{\lambda}$ if $\mathsf{cumul}_{\lambda} + J^{\mathsf{in}}_{\lambda} \le K$ else $\mathsf{cumul}'_{\lambda} = K$.
7. Finally, all parameter variables are unchanged.

Remark. The nuXmv encoding additionally requires the transition relation to be total, i.e., to have no deadlock states. We ensure this manually by adding stuttering transitions wherever necessary.

LTL Specification

For completeness, we recall the HSCL-to-LTL translation presented in Section 3.2. From an HSCL formula φ, we construct an LTL formula $\mathsf{LTL}(\varphi)$ recursively (see Def. 2 for the syntax of HSCL). The atoms of the translated formula are linear-arithmetic formulas over the history variables of the finite-counter system, namely formulas of the form $\sum_{\ell\in X} c_\ell \cdot \mathsf{local}_\ell \le t$ or $\sum_{\ell\in X} c_\ell \cdot \mathsf{cumul}_\ell \le t$, where $c_\ell \in \mathbb{N}$ are constants, $t \in \mathsf{LinTrm}(\mathsf{P})$, and $\mathsf{local}_\ell, \mathsf{cumul}_\ell$ are history variables. The translation is defined as follows. For a universal round-local atom $\forall r.\ \alpha_r$, where $\alpha_r = \sum_{\ell\in X} c_\ell \cdot \kappa(\ell, r) \le t$, we define $\mathsf{LTL}(\forall r.\ \alpha_r) = \mathbf{G}\big(\sum_{\ell\in X} c_\ell \cdot \mathsf{local}_\ell \le t\big)$. For a cumulative atom $\beta = \sum_{\ell\in X} c_\ell \cdot \sum_r \kappa(\ell, r) \le t$, we define $\mathsf{LTL}(\beta) = \mathbf{G}\big(\sum_{\ell\in X} c_\ell \cdot \mathsf{cumul}_\ell \le t\big)$. Boolean connectives are translated compositionally: $\mathsf{LTL}(\neg\psi) = \neg\mathsf{LTL}(\psi)$ and $\mathsf{LTL}(\psi \wedge \psi') = \mathsf{LTL}(\psi) \wedge \mathsf{LTL}(\psi')$.