

Eliminating Message Counters in Synchronous Threshold Automata^{*}

Ilina Stoilkovska^{1,2} (✉), Igor Konnov², Josef Widder², and Florian Zuleger¹

¹ TU Wien, Vienna, Austria

{stoilkov,zuleger}@forsyte.at

² Informal Systems, Vienna, Austria

{igor,josef}@informal.systems

Abstract. In previous work, we introduced synchronous threshold automata for the verification of synchronous fault-tolerant distributed algorithms, and presented a verification method based on bounded model checking. Modeling a distributed algorithm by a threshold automaton requires to correctly deal with the semantics for sending and receiving messages based on the fault assumption. This step was done manually so far, and required human ingenuity. Motivated by similar results for asynchronous threshold automata, in this paper we show that one can start from a faithful model of the distributed algorithm that includes the sending and receiving of messages, and then automatically obtain a threshold automaton by applying quantifier elimination on the receive message counters. In this way, we obtain a fully automated verification pipeline. We present an experimental evaluation, discovering a bug in our previous manual encoding. Interestingly, while quantifier elimination in general produces larger threshold automata than the manual encoding, the verification times are comparable and even faster in several cases, allowing us to verify benchmarks that could not be handled before.

1 Introduction

Formal modeling and automated verification of fault-tolerant distributed algorithms [2,28] received considerable attention recently, e.g., [8,20,29,32,38]. In the more classic approach towards distributed algorithms' correctness, algorithms are described in pseudo code, using send and receive operations whose semantics are typically not formalized, but given in English. As a result, this may lead to ambiguities that are an obstacle both for implementing distributed algorithms faithfully, as well as for computer-aided verification. Threshold automata were introduced as a formalization of fault-tolerant distributed algorithms with precise semantics [5,23,26], and effective automated verification methods have been introduced both for the asynchronous [22] and for the synchronous [36] case. While they are a concise model that allows to capture precisely the non-determinism distributed systems exhibit due to the communication model and

^{*} Partially supported by: Interchain Foundation, Switzerland; Austrian Science Fund (FWF) via doctoral college LogiCS W1255.

partial faults, threshold automata in fact constitute a manual abstraction: a threshold automaton has to capture two major ingredients of a distributed system: (i) the local program control flow that is based on received messages and (ii) the semantics of send and receive operations in a fault-prone environment. For many classical distributed algorithms, this manual abstraction is quite immediate, but as has been observed in [37], more involved distributed algorithms are harder to abstract manually. This manual process consists in understanding how a fault assumption — that typically is well-understood but not formalized — changes the semantics of sending and receiving messages, which is a formalization step that typically requires human ingenuity. The more desirable approach is to have a precise and formal description of (i) and (ii), and to construct the abstraction automatically. This also allows to reuse (ii), that is, the formalization of given distributed computing model for new benchmarks. Indeed, in [37], for asynchronous algorithms, we introduced a method that takes as input formalizations of (i) and (ii) and automatically constructs threshold automata. By this, we have reduced the required expertise of the user, increased the degree of automation on the verification process, and indeed found some bugs in manual abstractions of asynchronous algorithms. However, the approach in [37] focuses on (asynchronous) interleaving semantics, and asynchronous message passing, which pose different challenges than the synchronous setting.

While distributed algorithms are mostly designed for asynchronous systems, there exists a considerable amount of literature that focuses on *synchronous* distributed algorithms. The synchronous computation model is relevant, both theoretically and practically: (a) a well-known impossibility result [18] reveals a class of problems for which a solution in the asynchronous model does not exist, but which can be solved in the synchronous model, (b) some real-time systems are actually built on top of synchronous distributed algorithms [24], and (c) several verification approaches reduce the asynchronous to the synchronous setting [4, 12, 13, 15, 19, 25], enabling the transfer of verification techniques. For these reasons, verification in the synchronous setting received significant interest recently [1, 17, 29]. Applying verification techniques discovered a bug in an already published synchronous consensus algorithm, as reported in [27].

In [36], we proposed a synchronous variant of threshold automata along with an automated verification method based on bounded model checking. We experimentally evaluated our approach on a large number of benchmarks coming from the distributed systems literature. However, the framework in [36] is based on the manual abstraction described above.

Our Contributions. In this paper, we bring the automatic generation of threshold automata to the synchronous setting. We propose a *synchronous* threshold automata (STA) framework that allows us to:

1. model a given algorithm with an STA, whose guards are linear integer arithmetic expressions over the number of *received* messages, such that the obtained STA is in one-to-one correspondence with the pseudo code,

2. model the implicit assumptions imposed by the computation and fault models explicitly, using a so-called *environment assumption*, which is specific to the respective fault model and can be reused for different algorithms,
3. automatically translate the guards over the *local receive* variables into guards over the number of *globally sent* messages, using quantifier elimination,
4. pass the output of the translation as input to the verification tool proposed in [36], which implements a semi-decision procedure for computing the diameter, and performs bounded model checking.

In [36], the STA given as input to the verification tool was produced manually, that is, the steps 1–3 above were done by the user. By automating these steps, we reduce the ingenuity required by the user. We encoded the control flow and the environment assumptions of several synchronous algorithms in our framework and compared the resulting STA with the existing manual encodings from [34]. We confirm that manual abstraction is error-prone, as we discovered glitches in previous manually encoded STA. For all benchmarks, the automatically generated STA are comparable with the manual encodings. For some, the automatically generated STA could be verified faster. Thus in addition to increasing the degree of automation, we also gained in performance.

2 Our approach at a glance

Synchronous Distributed Algorithms. A distributed algorithm is a collection of n processes that perform a common task and exchange messages. At most t of the n processes can be faulty, and f processes are actually faulty. The numbers n, t, f are parameters, where n and t are “known”, that is, they appear in the code (see Fig. 1), while f may differ according to the individual executions. In the synchronous computation model, the actions that a process takes locally depend on the messages that the process has received in the current round by other processes. Often, a process checks whether a quorum has been obtained (e.g., majority, two-thirds, etc.) by counting the number of messages it has received. Obtaining a quorum means that the number of *received* messages has to pass a given threshold, which should guarantee that it is safe for a correct process to take an action, and move to a new local state.

The threshold automata framework [23] is based on the observation that from the viewpoint of enabled transitions in a transition system, we may substitute the check whether a quorum of messages has been *received* with a check whether enough messages have been *sent*. For some algorithms, this substitution is straightforward, but others have more complicated guard expressions over the number of received messages. Consider, for example, the pseudo code of the algorithm *PhaseQueen* [6,9], presented in Fig. 1. The algorithm operates in phases, with two rounds per phase (lines 3–8 and 9–11). In round 1, all processes broadcast their value stored in the variable v (line 3), receive messages from other processes (line 4), and count the number of messages with value 0 (line 5) and value 1 (line 6). If a process received more than $2t$ messages with value 1, then it sets its value to 1 (line 7), otherwise it sets its value to 0 (line 8). In round 2,

```

1 v := input({0, 1})
2 for each phase 1 through t + 1 do
3   broadcast v /* round 1: full message exchange */
4   receive messages from other processes
5   C[0] := number of received 0's
6   C[1] := number of received 1's
7   if C[1] > 2t then v := 1
8   else v := 0
9   if phase = i then broadcast v /* round 2: queen's broadcast */
10  receive queen's message vq
11  if C[v] < n - t then v := vq

```

Fig. 1. The pseudo code of the Byzantine consensus algorithm PhaseQueen

a process i acts as a queen, if the number of the current phase is equal to i (line 9), and it is the only process that broadcasts (line 9). Each process receives the queen's value v_q (line 10), and checks if in round 1, it received less than $n - t$ messages with value equal to its own value v . If this is the case, the process sets its value to the value v_q received from the queen (line 11). This algorithm satisfies the property *agreement*: it ensures that after phase $t + 1$, i.e., after the loop on line 2 terminates, all correct processes have the same value v .

Receive Synchronous Threshold Automata. In Section 3, we propose a *new variant* of synchronous threshold automata, rSTA, with guards expressed over receive variables. Fig. 2 shows the rSTA of the algorithm PhaseQueen. It corresponds to the control flow of the pseudo code in Fig. 1 as follows. The following locations capture local states of correct processes that are currently not a queen:

- vi encodes that a process has the value $i \in \{0, 1\}$,
- $R1vi$ encodes that after the first round a process sets its value to $i \in \{0, 1\}$, and that it has received at least $n - t$ messages that have its value (i.e., the condition from line 11 evaluates to false),
- $R1viQ$ encodes that after the first round a process sets its value to $i \in \{0, 1\}$, and that it has received less than $n - t$ messages that have its value. Such a process will use the queen's message to update its value at the end of the second round (that is, the condition in line 11 evaluates to true),
- $R2vi$ encodes that after the second round a process sets its value to $i \in \{0, 1\}$.

From the location $R2vi$, we have outgoing rules that bring the process back to the beginning of the next phase, i.e., to vi , for $i \in \{0, 1\}$. Additionally, a process might move from the location $R2vi$ to Qvi , for $i \in \{0, 1\}$, and thus become a queen in the next phase. The locations $Qvi, R1Qvi, R2Qvi$, for $i \in \{0, 1\}$, capture the behavior of a correct process acting as a queen in the current phase. The Byzantine processes can act arbitrary, and their behavior is not explicitly modeled in the automaton. However, in some phase, the queen may be Byzantine. To capture this, we introduce locations, populated by a single Byzantine process, namely the locations $F = \{F, \dots, R2QF\}$. The queen is Byzantine in some phase, if the single Byzantine process moves from the location $R2F$ to the location QF .

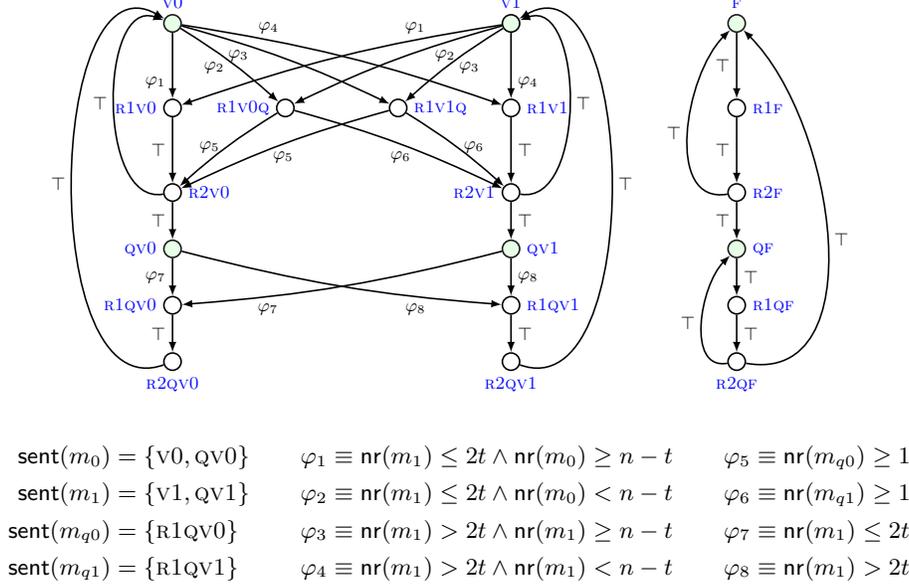


Fig. 2. The rSTA for the algorithm PhaseQueen [6], where $n > 4t \wedge t \geq f$.

Processes in locations v_i, QV_i send messages of type m_i , that is, messages containing the value $i \in \{0, 1\}$. The message types m_{q_i} are used to encode that the queen in the current phase sent a message with value $i \in \{0, 1\}$. When the queen process is Byzantine, it can send messages of type m_{q0} or m_{q1} . We write $\text{sent}(m)$ to denote the set of locations where processes send a message of type m , and $\#\text{sent}(m)$ for the number of sent messages of type m .

The receive guards $\varphi_1, \dots, \varphi_8$ express conditions over the number of received messages of some message type, and capture expressions which appear in the pseudo code. We denote by $\text{nr}(m_i)$ and $\text{nr}(m_{q_i})$ the number of messages containing the value $i \in \{0, 1\}$ that a process received from all processes in the first round of the phase (i.e., the value $C[i]$ in the pseudo code, lines 5, 6) and by the queen in the second round of the phase, respectively. For example, the receive guard φ_1 , occurring on rules that move processes to the location R1V0, checks if a process received at most $2t$ messages of type m_1 (the **else** branch is taken in line 8), and at least $n - t$ messages of type m_0 (the condition in line 11 is false).

We explicitly encode the relationship between the number of received and sent messages using an *environment assumption* Env , which bounds the number of received messages: (i) from below by the number of messages sent by the correct processes, and (ii) from above by the number of messages sent by both the correct and faulty processes. The bound (i) captures the assumptions of the synchronous communication, which requires that all messages sent by

correct processes in a round are received in the same round, and the bound (ii) captures the non-determinism introduced by the faulty processes. E.g., in the algorithm *PhaseQueen*, we have f Byzantine processes, which may send messages of arbitrary types. For the receive variable $\text{nr}(m_i)$, we have the constraint $\#\text{sent}(m_i) \leq \text{nr}(m_i) \leq \#\text{sent}(m_i) + f$ in the environment assumption Env .

The agreement property stated above is a safety property. To check if it holds, it suffices to check that after $t + 1$ phases, either all processes are in locations $\text{v0}, \text{Qv0}$, or in locations $\text{v1}, \text{Qv1}$. The precise formalization of the properties we are interested in verifying can be found in [36].

Our approach. In Section 6, we eliminate the receive variables in an *rSTA* using quantifier elimination for Presburger arithmetic [14, 30, 31]. We strengthen the receive guards by the environment assumption Env that imposes bounds on the values of the receive variables, which are existentially quantified. As a result, a quantifier-free guard expression over the number of sent messages is obtained. For example, the result of applying quantifier elimination to the guard φ_1 over the receive variables from Fig. 2, strengthened by the upper and lower bounds in the environment assumption Env , is the guard $\widehat{\varphi}_1$ with no receive variables:

$$\widehat{\varphi}_1 \equiv \#\text{sent}(m_1) \leq 2t \wedge \#\text{sent}(m_0) + f \geq n - t \wedge \widehat{\text{Env}}$$

where $\widehat{\text{Env}}$ are the residual constraints from eliminating the receive variables from the environment assumption Env . The condition $\text{nr}(m_1) \leq 2t$ in the guard φ_1 is translated to $\#\text{sent}(m_1) \leq 2t$, and the condition $\text{nr}(m_0) \geq n - t$ to $\#\text{sent}(m_0) + f \geq n - t$. That is, when translating the guards, the number of the faulty processes f is used in guards that check if the number of sent messages passes a threshold, whereas f is not used in guards that check if the number of sent messages is below a threshold. (Byzantine processes send messages arbitrarily.)

The *STA* where all guards over the receive variables are replaced by the automatically generated guards over the number of sent messages constitutes a valid input to the bounded model checking technique for *STA* from [36], which we use to verify their safety properties. We show that this method is sound and complete by showing the existence of a bisimulation between the composition of n copies of *rSTA* and the composition of n copies of the produced *STA*. Thus, eliminating the receive message counters preserves temporal properties. We implemented this technique and used it to automatically generate *STA* for a set of benchmarks, and compared them to the existing manually encoded *STA* for the same benchmarks. We discuss our the experimental results in Section 7.

3 Synchronous Threshold Automata

We recall synchronous threshold automata from [36] and extend them with receive variables below. A *synchronous threshold automaton (STA)* is the tuple $\text{STA} = (\mathcal{L}, \mathcal{I}, \mathcal{R}, \Pi, RC, \text{Env})$, whose locations \mathcal{L} , initial locations \mathcal{I} , rules \mathcal{R} , parameters Π , and resilience condition RC are defined below. We define the environment assumption Env in Section 3.2.

Parameters Π , Resilience Condition RC . We assume that the set Π of *parameters* contains at least the parameter n , denoting the total number of processes. The *resilience condition RC* is a linear arithmetic expression over the parameters from Π . We call the vector $\boldsymbol{\pi} = \langle \pi_1, \dots, \pi_{|\Pi|} \rangle$ the *parameter vector*, and the vector $\mathbf{p} = \langle p_1, \dots, p_{|\Pi|} \rangle \in \mathbb{N}^{|\Pi|}$ a *valuation* of $\boldsymbol{\pi}$. The set $\mathbf{P}_{RC} = \{\mathbf{p} \in \mathbb{N}^{|\Pi|} \mid \mathbf{p} \text{ is a valuation of } \boldsymbol{\pi} \text{ and } \mathbf{p} \text{ satisfies } RC\}$ contains the *admissible valuations* of $\boldsymbol{\pi}$. The mapping $N : \mathbf{P}_{RC} \rightarrow \mathbb{N}$ maps an admissible valuation $\mathbf{p} \in \mathbf{P}_{RC}$ to the number $N(\mathbf{p}) \in \mathbb{N}$ of *participating processes*, i.e., the number of processes whose behavior is modeled using the STA. We denote by $N(\boldsymbol{\pi})$ the linear combination of parameters that defines the number of participating processes.

Locations \mathcal{L}, \mathcal{I} . The *locations* $\ell \in \mathcal{L}$ encode the current value of the local variables of a process, together with information about the program counter. We assume that each local variable and the program counter ranges over a finite set of values, that is, we assume that the set \mathcal{L} of locations is a finite set. The *initial locations* in $\mathcal{I} \subseteq \mathcal{L}$ encode the initial values of the local variables.

Message Types \mathcal{M} . Let \mathcal{M} denote the set of *message types*. To encode sending messages in the STA, we define a mapping $\text{sent} : \mathcal{M} \rightarrow 2^{\mathcal{L}}$, that maps a message type $m \in \mathcal{M}$ to a set $\text{sent}(m) \subseteq \mathcal{L}$ of locations, such that $\text{sent}(m) = \{\ell \in \mathcal{L} \mid \text{a process in } \ell \text{ sends message of type } m\}$.

Let $L \subseteq \mathcal{L}$ denote a set of locations, and let $\#L$ denote the number of processes in locations from the set L . To define guards over the sent messages and express temporal properties, we define *c-propositions*:

$$\#L \geq \mathbf{a} \cdot \boldsymbol{\pi} + b \text{ for } L \subseteq \mathcal{L}, \mathbf{a} \in \mathbb{Z}^{|\Pi|}, \text{ and } b \in \mathbb{Z}$$

We denote by CP the set of *c-propositions*. If the set L of locations in the *c-proposition* is equal to the set $\text{sent}(m)$, for some $m \in \mathcal{M}$, the *c-proposition* is used to check whether the number of messages of type $m \in \mathcal{M}$ is greater than or equal to a linear combination of the parameters, also called a *threshold*. Formally, the *c-propositions* are evaluated in tuples $(\boldsymbol{\kappa}, \mathbf{p})$, where $\boldsymbol{\kappa} \in \mathbb{N}^{|\mathcal{L}|}$ is an $|\mathcal{L}|$ -dimensional vector of *counters*, and $\mathbf{p} \in \mathbf{P}_{RC}$ is an admissible valuation:

$$(\boldsymbol{\kappa}, \mathbf{p}) \models \#L \geq \mathbf{a} \cdot \boldsymbol{\pi} + b \quad \text{iff} \quad \sum_{\ell \in L} \boldsymbol{\kappa}[\ell] \geq \mathbf{a} \cdot \mathbf{p} + b \quad (1)$$

Rules \mathcal{R} . A rule $r \in \mathcal{R}$ is a tuple $(\text{from}, \text{to}, \varphi)$, where: $\text{from}, \text{to} \in \mathcal{L}$ are locations, and φ is a *guard*, i.e., a Boolean combination of *c-propositions*. The guards $r.\varphi$, for $r \in \mathcal{R}$, analogously to (1), are evaluated in tuples $(\boldsymbol{\kappa}, \mathbf{p})$, and the semantics of the Boolean connectives is standard.

3.1 Receive Synchronous Threshold Automata

A *receive STA* is the tuple $\text{rSTA} = (\mathcal{L}, \mathcal{I}, \mathcal{R}^\Delta, \Delta, \Pi, RC, \text{Env}^\Delta)$, whose locations \mathcal{L} , initial locations \mathcal{I} , parameters Π , and resilience condition RC are defined as for STA. We define the receive variables Δ and rules \mathcal{R}^Δ below, and the environment assumption Env^Δ in Section 3.2.

Receive Variables Δ . The set Δ contains *receive variables* $\text{nr}(m)$ that store the number of messages of type $m \in \mathcal{M}$ that were received by a process. Thus, $|\Delta| = |\mathcal{M}|$, as in Δ there is exactly one receive variable $\text{nr}(m)$ per message type $m \in \mathcal{M}$. The values of the receive variables depend on the number of messages sent in a given round (discussed in more detail in Section 3.2).

Let $M \subseteq \mathcal{M}$ denote a set of message types, and let $\#M$ denote the total number of messages of types $m \in M$, received by some process. Observe that the notation $\#M$ is a shorthand for $\sum_{m \in M} \text{nr}(m)$. We will use these two notations interchangeably. Further, when M is a singleton set, that is, when $M = \{m\}$, we will simply use the notation $\text{nr}(m)$ to denote $\#\{m\}$. For the purpose of expressing guards over the receive variables $\text{nr}(m)$, for $m \in \mathcal{M}$, we define *r-propositions*:

$$\#M \geq \mathbf{a} \cdot \boldsymbol{\pi} + b, \text{ such that } M \subseteq \mathcal{M}, \mathbf{a} \in \mathbb{Z}^{|\Pi|}, b \in \mathbb{Z}$$

We denote by RP the set of *r-propositions*. The intended meaning of the *r-propositions* is to check whether the total number of messages of types $m \in M$ received by some process i passes some threshold. Formally, they are evaluated in tuples (\mathbf{d}, \mathbf{p}) , where $\mathbf{d} \in \mathbb{N}^{|\mathcal{M}|}$ is a vector of values assigned to each receive variable $\text{nr}(m)$, for $m \in \mathcal{M}$, and $\mathbf{p} \in \mathbf{P}_{RC}$. We define:

$$(\mathbf{d}, \mathbf{p}) \models \#M \geq \mathbf{a} \cdot \boldsymbol{\pi} + b \quad \text{iff} \quad \sum_{m \in M} \mathbf{d}[m] \geq \mathbf{a} \cdot \mathbf{p} + b \quad (2)$$

Rules \mathcal{R}^Δ . Similarly to the way we defined rules of STA above, the rules $r^\Delta \in \mathcal{R}^\Delta$ in rSTA are tuples $r^\Delta = (\text{from}, \text{to}, \varphi)$, where $r^\Delta.\text{from}, r^\Delta.\text{to} \in \mathcal{L}$ are locations, and $r^\Delta.\varphi$ is a *receive guard*, which is a Boolean combination of *c-propositions* and *r-propositions*. The receive guards $r^\Delta.\varphi$, for $r^\Delta \in \mathcal{R}^\Delta$, are evaluated in tuples $(\mathbf{d}, \boldsymbol{\kappa}, \mathbf{p})$. Given a tuple $(\mathbf{d}, \boldsymbol{\kappa}, \mathbf{p})$, where $\mathbf{d} \in \mathbb{N}^{|\mathcal{M}|}$ is a vector of valuations of the receive variables $\text{nr}(m)$, for $m \in \mathcal{M}$, $\boldsymbol{\kappa} \in \mathbb{N}^{|\mathcal{L}|}$ is an $|\mathcal{L}|$ -dimensional vector of counters, and $\mathbf{p} \in \mathbf{P}_{RC}$ is an admissible valuation, we evaluate *c-propositions* and *r-propositions* (the semantics of the Boolean connectives is standard):

$$(\mathbf{d}, \boldsymbol{\kappa}, \mathbf{p}) \models \#L \geq \mathbf{a} \cdot \boldsymbol{\pi} + b \quad \text{iff} \quad (\boldsymbol{\kappa}, \mathbf{p}) \models \#L \geq \mathbf{a} \cdot \boldsymbol{\pi} + b \quad (\text{cf. (1)})$$

$$(\mathbf{d}, \boldsymbol{\kappa}, \mathbf{p}) \models \#M \geq \mathbf{a} \cdot \boldsymbol{\pi} + b \quad \text{iff} \quad (\mathbf{d}, \mathbf{p}) \models \#M \geq \mathbf{a} \cdot \boldsymbol{\pi} + b \quad (\text{cf. (2)})$$

3.2 Environment Assumption and Modeling Faults

Depending on the fault model, when constructing a (receive) STA that models the behavior of a process running a given algorithm, we typically need to introduce additional locations or rules that are used to capture the behavior of the faulty processes. Additionally, to faithfully model the faulty environment, we will introduce constraints on the number of processes in given locations in both STA and rSTA, expressed using *c-propositions*, as well as constraints on the values of the receive variables of the rSTA, expressed using *e-propositions*:

$$\#M \geq \#L + \mathbf{a} \cdot \boldsymbol{\pi} + b, \text{ such that } M \subseteq \mathcal{M}, L \subseteq \mathcal{L}, \mathbf{a} \in \mathbb{Z}^{|\Pi|}, b \in \mathbb{Z}$$

We denote by EP the set of e -propositions. The e -propositions are evaluated in tuples $(\mathbf{d}, \boldsymbol{\kappa}, \mathbf{p})$ where $\mathbf{d} \in \mathbb{N}^{|\mathcal{M}|}$ is a vector of valuations of the receive variables, $\boldsymbol{\kappa} \in \mathbb{N}^{|\mathcal{L}|}$ is an $|\mathcal{L}|$ -dimensional vector of counters, and $\mathbf{p} \in \mathbf{P}_{RC}$. We say that:

$$(\mathbf{d}, \boldsymbol{\kappa}, \mathbf{p}) \models \#M \geq \#L + \mathbf{a} \cdot \boldsymbol{\pi} + b \quad \text{iff} \quad \sum_{m \in \mathcal{M}} \mathbf{d}[m] \geq \sum_{\ell \in \mathcal{L}} \boldsymbol{\kappa}[\ell] + \mathbf{a} \cdot \mathbf{p} + b$$

The e -propositions will be used to express that the number of received messages is in the range from the number of messages sent by *correct* processes to the total number of sent messages (sent by both correct and faulty processes).

For STA, the environment assumption Env is a conjunction of c -propositions and their negations. For rSTA, the environment assumption Env^Δ is a conjunction of c -propositions, e -propositions and their negations. The c -propositions restrict the number of processes in certain locations, while the e -propositions restrict the values of the receive variables by relating them to the number of sent messages of the same type. We define the environment assumptions Env and Env^Δ of the STA and rSTA, respectively, as $\text{Env} \equiv \text{Env}_{\text{CP}}$ and $\text{Env}^\Delta \equiv \text{Env}_{\text{CP}} \wedge \text{Env}_{\text{EP}}$, where Env_{CP} and Env_{EP} are conjunctions of c -propositions and e -propositions and their negations, respectively, such that:

$$\text{Env}_{\text{CP}} \equiv \text{C1} \wedge \text{C2} \wedge \text{Env}_{\text{CP},*} \quad \text{and} \quad \text{Env}_{\text{EP}} \equiv \text{E1} \wedge \text{Env}_{\text{EP},*}$$

where, irrespective of the fault model, we have the following constraints:

- (C1) $\bigwedge_{\ell \in \mathcal{L}} \#\{\ell\} \geq 0$, i.e., the number of processes in a location ℓ is non-negative,
- (C2) $\#\mathcal{L} = N(\boldsymbol{\pi})$, i.e., the number of processes in all locations \mathcal{L} is equal to the number of participating processes,
- (E1) $\bigwedge_{m \in \mathcal{M}} \#\text{sent}(m) \leq \text{nr}(m)$, i.e., the number $\text{nr}(m)$ of received messages of each message type $m \in \mathcal{M}$ is bounded from below by the number $\#\text{sent}(m)$ of messages of type m , sent by correct processes.

The formulas $\text{Env}_{\text{CP},*}$ and $\text{Env}_{\text{EP},*}$ for $* \in \{\text{cr}, \text{so}, \text{byz}\}$, depend on the fault model, i.e., on whether we model crash, send omission, or Byzantine faults.

Crash Faults. Crash-faulty processes stop executing the algorithm prematurely and cannot restart. To model the behavior of the crash-faulty processes, the set \mathcal{L} of locations of the (receive) STA is the set: $\mathcal{L} = \mathcal{L}_{\text{corr}} \cup \mathcal{L}_{\text{cr}} \cup \{\ell_{\text{fld}}\}$, where $\mathcal{L}_{\text{corr}}$ is a set of *correct* locations, $\mathcal{L}_{\text{cr}} = \{\ell_{\text{cr}} \mid \ell_{\text{cr}} \text{ is a fresh copy of } \ell \in \mathcal{L}_{\text{corr}}\}$ is a set of *crash* locations, and ℓ_{fld} is a *failed* location. The crash locations $\ell_{\text{cr}} \in \mathcal{L}_{\text{cr}}$ model the same values of the local variables and program counter as their correct counterpart $\ell \in \mathcal{L}_{\text{corr}}$. The difference is that processes in the crash locations $\ell_{\text{cr}} \in \mathcal{L}_{\text{cr}}$ are flagged by the environment to crash in the current round. After crashing, they move to the failed location ℓ_{fld} , where they remain forever. This models that the crashed processes cannot restart.

A crash-faulty process may send a message to a subset of the other processes in the round in which it crashes. To model this, we introduce the mapping $\text{sent}_{\text{cr}} : \mathcal{M} \rightarrow 2^{\mathcal{L}_{\text{cr}}}$, which defines, for each $m \in \mathcal{M}$, the set of crash locations $\text{sent}_{\text{cr}}(m) \subseteq$

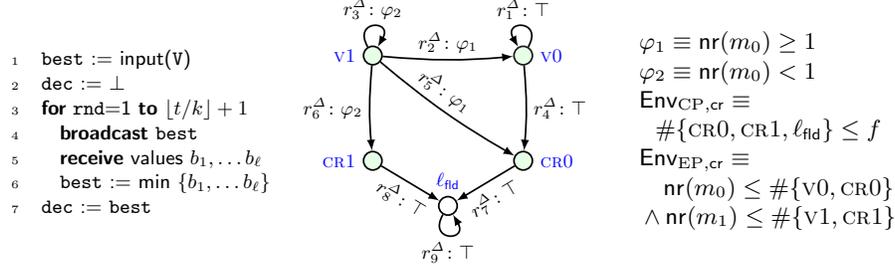


Fig. 3. The pseudo code of the algorithm FloodMin for $k = 1$ [28], which tolerates crash faults, and the receive STA encoding its loop body.

\mathcal{L}_{cr} where processes send a message of type m . Then, $\#(\text{sent}(m) \cup \text{sent}_{\text{cr}}(m))$ denotes the number of messages sent by correct and crash-faulty processes. In addition to the new locations, we add the following new rules:

- (cr1) for every rule $r \in \mathcal{R}$, if $r.\text{from} \in \mathcal{L}_{\text{corr}}$ and $r.\text{to} \in \mathcal{L}_{\text{corr}}$, then we add the rule $(r.\text{from}, \ell_{\text{cr}}, r.\varphi)$, where $\ell_{\text{cr}} \in \mathcal{L}_{\text{cr}}$ is the crash location corresponding to $r.\text{to}$,
- (cr2) for every crash location $\ell_{\text{cr}} \in \mathcal{L}_{\text{cr}}$, we add the rule $(\ell_{\text{cr}}, \ell_{\text{fld}}, \top)$,
- (cr3) for the failed location ℓ_{fld} , we add the rule $(\ell_{\text{fld}}, \ell_{\text{fld}}, \top)$.

The rules (cr1) move processes from the correct to the crash locations, in rounds where the environment flags them as crashed. The rules (cr2) move processes from the crashed locations to the failed location, where they can only apply the self-loop rule (cr3), which keeps them in the failed location.

We model the behavior of crash-faulty processes explicitly, that is, we have $N(\pi) = n$. The constraints $\text{Env}_{\text{CP},\text{cr}}$ and $\text{Env}_{\text{EP},\text{cr}}$ for the crash fault model are:

$$\begin{aligned} \text{Env}_{\text{CP},\text{cr}} &= \#(\mathcal{L}_{\text{cr}} \cup \{\ell_{\text{fld}}\}) \leq f \\ \text{Env}_{\text{EP},\text{cr}} &\equiv \bigwedge_{m \in \mathcal{M}} \text{nr}(m) \leq \#(\text{sent}(m) \cup \text{sent}_{\text{cr}}(m)) \end{aligned}$$

The formula $\text{Env}_{\text{CP},\text{cr}}$ ensures that there are no more than f faults. The formula $\text{Env}_{\text{EP},\text{cr}}$ restricts the values of the receive variables by ensuring that the number of received messages of type $m \in \mathcal{M}$ for each process is a value, bounded from above by the number $\#(\text{sent}(m) \cup \text{sent}_{\text{cr}}(m))$ of messages of type m , sent by the correct processes and the processes flagged as crashed in the current round.

Fig. 3 depicts the pseudo code and the rSTA of the crash-tolerant k -set agreement algorithm FloodMin, for $k = 1$ [28]. We identify the sets $\mathcal{L}_{\text{corr}} = \{\text{v0}, \text{v1}\}$ of correct locations, $\mathcal{L}_{\text{cr}} = \{\text{CR0}, \text{CR1}\}$ of crash locations, $\mathcal{M} = \{m_0, m_1\}$ of message types. The location vi encodes that a correct process has its variable **best** set to $i \in \{0, 1\}$, the location CRi encodes that the value of **best** of a crashed process is $i \in \{0, 1\}$, and the message type m_i encodes a message containing the value $i \in \{0, 1\}$. The failed location is ℓ_{fld} . We define $\text{sent}(m_i) = \{\text{vi}\}$ and $\text{sent}_{\text{cr}}(m_i) = \{\text{CRi}\}$, for $i \in \{0, 1\}$. The two receive guards $\varphi_1 \equiv \text{nr}(m_0) \geq 1$ and $\varphi_2 \equiv \text{nr}(m_0) < 1$ check if a process received at least one message of type m_0

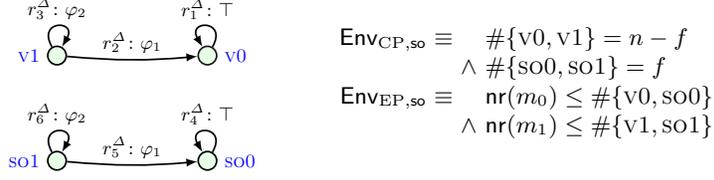


Fig. 4. The receive STA encoding the loop body of the algorithm FMinOmit for $k = 1$, which tolerates send omission faults and whose pseudo code is given in Fig. 3.

(i.e., if the minimal value 0 has been received in line 5 of the pseudo code) and no message of type m_0 , respectively. The constraint $\text{Env}_{\text{CP},\text{cr}}$ ensures that there are not more than f processes in the locations CR0 , CR1 , and ℓ_{fld} together. The constraint $\text{Env}_{\text{EP},\text{cr}}$ bounds the values of the receive variables $\text{nr}(m_i)$ from above by the number of processes in locations v_i , $\text{CR}i$, for $i \in \{0, 1\}$.

Send Omission Faults. A send-omission-faulty process may omit to send a message, but acts as a correct process on the receiving side. We model algorithms tolerating send omission faults similarly to crash faults: the set \mathcal{L} of locations is $\mathcal{L} = \mathcal{L}_{\text{corr}} \cup \mathcal{L}_{\text{so}}$, where $\mathcal{L}_{\text{corr}}$ is a set of *correct* locations and $\mathcal{L}_{\text{so}} = \{\ell_{\text{so}} \mid \ell_{\text{so}} \text{ is a fresh copy of } \ell \in \mathcal{L}_{\text{corr}}\}$ is a set of *send-omission* locations. For every rule $r \in \mathcal{R}$ connecting two locations $\ell, \ell' \in \mathcal{L}_{\text{corr}}$, there exists a rule $(\ell_{\text{so}}, \ell'_{\text{so}}, r, \varphi) \in \mathcal{R}$, connecting their two corresponding send-omission locations $\ell_{\text{so}}, \ell'_{\text{so}} \in \mathcal{L}_{\text{so}}$. We introduce the mapping $\text{sent}_{\text{so}} : \mathcal{M} \rightarrow 2^{\mathcal{L}_{\text{so}}}$, which defines the set of send-omission locations where processes send a message of type $m \in \mathcal{M}$.

As there are no rules that connect the locations from $\mathcal{L}_{\text{corr}}$ to the locations from \mathcal{L}_{so} , the automaton consists of two parts: one used by the correct processes, and one used by the send-omission-faulty processes. The behavior of the send-omission-faulty processes is encoded explicitly, using locations and rules in the automaton, hence, we define $N(\pi) = n$. The constraint $\text{Env}_{\text{CP},\text{so}}$ ensures that the number of processes populating the correct locations is $n - f$, and the number of processes populating the send-omission locations is f . The constraint $\text{Env}_{\text{EP},\text{so}}$ ensures that the number of received messages of type $m \in \mathcal{M}$ for each process is bounded from above by the number $\#(\text{sent}(m) \cup \text{sent}_{\text{so}}(m))$ of messages of type m , sent by the correct and the send-omission-faulty processes. Formally:

$$\begin{aligned}
 \text{Env}_{\text{CP},\text{so}} &= \#\mathcal{L}_{\text{corr}} = n - f \wedge \#\mathcal{L}_{\text{so}} = f \\
 \text{Env}_{\text{EP},\text{so}} &\equiv \bigwedge_{m \in \mathcal{M}} \text{nr}(m) \leq \#(\text{sent}(m) \cup \text{sent}_{\text{so}}(m))
 \end{aligned}$$

Fig. 4 depicts the rSTA for the k -set agreement algorithm FMinOmit, for $k = 1$, which is a variant of the algorithm FloodMin (Fig. 3) that tolerates send omission faults. We identify the sets $\mathcal{L}_{\text{corr}} = \{v0, v1\}$ of correct locations, $\mathcal{L}_{\text{so}} = \{\text{so0}, \text{so1}\}$ of send-omission locations, and $\mathcal{M} = \{m_0, m_1\}$ of message types. We define $\text{sent}(m_i) = \{v_i\}$ and $\text{sent}_{\text{so}}(m_i) = \{\text{so}i\}$, for $i \in \{0, 1\}$. The constraint $\text{Env}_{\text{CP},\text{so}}$ ensures that there are exactly $n - f$ processes in the correct

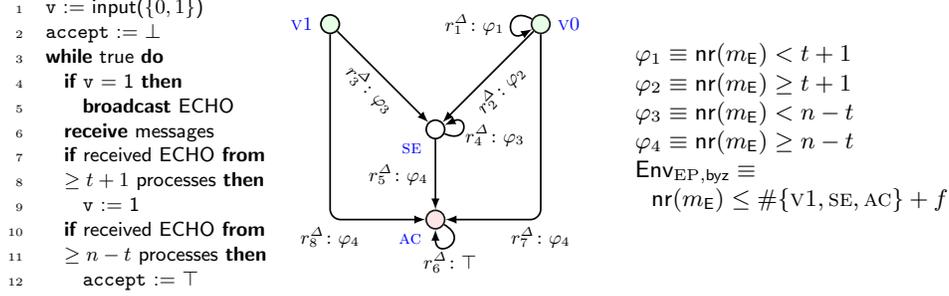


Fig. 5. The pseudo code of the algorithm RB [21], which tolerates Byzantine faults, and the receive STA encoding its loop body.

locations $v0, v1$, and exactly f processes in the send-omission locations $so0, so1$. The receive guards φ_1 and φ_2 are the syntactically same as in the $r\text{STA}$ for the crash-tolerant version of the algorithm FloodMin, for $k = 1$. However, the environment constraint $\text{Env}_{\text{EP},so}$ differs from $\text{Env}_{\text{EP},cr}$: it restricts the number $\text{nr}(m_i)$ of received messages of type m_i to a value which is less than or equal to the number of processes in locations v_i, so_i , for $i \in \{0, 1\}$.

Byzantine Faults. To model the behavior of the Byzantine-faulty processes, which can act arbitrary, no new locations and rules are introduced in the (receive) STA. Instead, the (receive) STA is used to model the behavior of the correct processes, and the effect that the Byzantine-faulty processes have on the correct ones is captured in the guards (and environment assumption). The number of messages sent by Byzantine-faulty processes is overapproximated by the parameter f , which denotes the number of faults. That is, for a message type $m \in \mathcal{M}$, the number $\#\text{sent}(m) + f$ is the upper bound on the number of messages sent by correct and Byzantine-faulty processes.

The (receive) STA for Byzantine faults is used to model the behavior of the correct processes, hence $N(\boldsymbol{\pi}) = n - f$. As we do not introduce new locations or rules, we have $\text{Env}_{\text{CP},\text{byz}} \equiv \top$. The constraint $\text{Env}_{\text{EP},\text{byz}}$ encodes the effect that the Byzantine-faulty processes have on the correct processes, by bounding the receive variables $\text{nr}(m)$ by $\text{sent}(m) + f$ from above, for $m \in \mathcal{M}$:

$$\text{Env}_{\text{EP},\text{byz}} \equiv \bigwedge_{m \in \mathcal{M}} \text{nr}(m) \leq \text{sent}(m) + f$$

Fig. 5 shows the pseudo code of the Byzantine reliable broadcast algorithm RB [21]. The locations $\mathcal{L} = \{v0, v1, \text{SE}, \text{AC}\}$ model the behavior of the correct processes. The location v_i encodes that a process has value $i \in \{0, 1\}$, the location SE that a process has sent an ECHO message, and the location AC that a process sets its value to 1 in line 12. There is a single message type, m_E , which encodes a message containing the value ECHO. There are four receive guards, $\varphi_1, \dots, \varphi_4$. The guard φ_2 , for example, checks that at least $t + 1$ ECHO messages

are received, capturing line 8 of the pseudo code. The set of processes that send an ECHO message is $\text{sent}(m_E) = \{V1, SE, AC\}$. The constraint $\text{Env}_{EP,byz}$ ensures that there are not more than $\#\{V1, SE, AC\} + f$ received messages of type m_E .

Remark on Algorithms with a Coordinator. When modeling Byzantine-tolerant algorithms where a process acts as a coordinator (such as, e.g., the algorithm `PhaseQueen` in Fig. 1), we need to take into account that at some point, the coordinator will be Byzantine. Thus, we add locations $\mathcal{L}_{byz} \subseteq \mathcal{L}$ for a single Byzantine process, disjoint from the locations that are used by the correct processes. The new locations do not encode any values of the local variables; they ensure that the Byzantine process (which may become a coordinator) moves synchronously with the other processes. In the rSTA for the algorithm `PhaseQueen` (Fig. 2), we defined $\mathcal{L}_{byz} = F = \{F, \dots, R2QF\}$. As we model the behavior of a single Byzantine process explicitly, we have $N(\boldsymbol{\pi}) = n - f + 1$.

In this case, we define the constraints $\text{Env}_{CP,co}$, which restrict the number of processes in given locations. We also identify locations $\mathcal{L}_{co} \subseteq \mathcal{L}$, which only a (correct or Byzantine) coordinator is allowed to populate. The environment constraint $\text{Env}_{CP,co}$ for Byzantine-tolerant algorithms with a coordinator is:

$$\text{Env}_{CP,co} \equiv \#\mathcal{L}_{co} = 1 \wedge \#\mathcal{L}_{byz} = 1$$

where $\#\mathcal{L}_{co} = 1$ (resp. $\#\mathcal{L}_{byz} = 1$) ensures that there is exactly one process in the coordinator locations \mathcal{L}_{co} (resp. in the Byzantine locations \mathcal{L}_{byz}).

Additionally, we have message types $m_{co} \in \mathcal{M}$ that model the coordinator messages, and denote by ℓ_F the location where the Byzantine process performs the coordinator broadcast. The constraint $\text{Env}_{EP,co}$ states that the number of received coordinator messages of type m_{co} does not exceed the total number of coordinator messages of type m_{co} sent by the correct and Byzantine coordinators:

$$\text{Env}_{EP,co} \equiv \text{Env}_{EP,byz} \wedge \bigwedge_{m_{co} \in \mathcal{M}} \text{nr}(m_{co}) \leq \#(\text{sent}(m_{co}) \cup \{\ell_F\})$$

Thus, for the algorithm `PhaseQueen`, whose rSTA we depicted in Fig. 2:

$$\begin{aligned} \text{Env}_{CP,co} &\equiv \#\{QV0, \dots, R2QV1, QF, \dots, R2QF\} = 1 \wedge \#\{F, \dots, R2QF\} = 1 \\ \text{Env}_{EP,co} &\equiv \bigwedge_{i \in \{0,1\}} (\text{nr}(m_i) \leq \#\text{sent}(m_i) + f \wedge \text{nr}(m_{qi}) \leq \#(\text{sent}(m_{qi}) \cup \{R1QF\})) \end{aligned}$$

4 Counter Systems

For an STA $(\mathcal{L}, \mathcal{I}, \mathcal{R}, II, RC, \text{Env})$ and an admissible valuation $\mathbf{p} \in \mathbf{P}_{RC}$, we recall the definition of a counter system from [36]. A *counter system* w.r.t. an admissible valuation $\mathbf{p} \in \mathbf{P}_{RC}$ and an STA $(\mathcal{L}, \mathcal{I}, \mathcal{R}, II, RC, \text{Env})$ is the tuple $\text{CS}(\text{STA}, \mathbf{p}) = (\Sigma(\mathbf{p}), I(\mathbf{p}), R(\mathbf{p}))$, representing a system of $N(\mathbf{p})$ processes whose behavior is modeled using the STA, where $\Sigma(\mathbf{p})$ is the set of *configurations*, $I(\mathbf{p})$ is the set of *initial configurations*, and $R(\mathbf{p})$ is the *transition relation*.

A *configuration* $\sigma \in \Sigma(\mathbf{p})$ is a tuple $(\boldsymbol{\kappa}, \mathbf{p})$, where $\mathbf{p} \in \mathbf{P}_{RC}$ is an admissible valuation, and $\boldsymbol{\kappa} \in \mathbb{N}^{|\mathcal{L}|}$ is an $|\mathcal{L}|$ -dimensional vector of *counters*, such that $\sigma \models \text{Env}$. For every $\sigma \in \Sigma(\mathbf{p})$, we have $\sum_{\ell \in \mathcal{L}} \sigma.\boldsymbol{\kappa}[\ell] = N(\mathbf{p})$. This follows from $\sigma \models \text{Env}$, in particular from $\sigma \models \#\mathcal{L} = N(\boldsymbol{\pi})$, the definition of $N(\mathbf{p})$, and the semantics of the c -propositions. A configuration $\sigma \in \Sigma(\mathbf{p})$ is *initial*, i.e., $\sigma \in I(\mathbf{p}) \subseteq \Sigma(\mathbf{p})$, iff $\sigma.\boldsymbol{\kappa}[\ell] = 0$, for every $\ell \in \mathcal{L} \setminus \mathcal{I}$. That is, the value $\sigma.\boldsymbol{\kappa}[\ell]$ of the counter for each non-initial location $\ell \in \mathcal{L} \setminus \mathcal{I}$ is set to 0 in $\sigma \in \mathcal{I}$.

To define the transition relation $R(\mathbf{p})$, we first define the notion of a transition. A *transition* is a function $tr : \mathcal{R} \rightarrow \mathbb{N}$ that maps each rule $r \in \mathcal{R}$ to a *factor* $tr(r) \in \mathbb{N}$. Given a valuation \mathbf{p} of $\boldsymbol{\pi}$, the set $Tr(\mathbf{p}) = \{tr \mid \sum_{r \in \mathcal{R}} tr(r) = N(\mathbf{p})\}$ contains transitions whose factors sum up to $N(\mathbf{p})$. For a transition tr and a rule $r \in \mathcal{R}$, the factor $tr(r)$ denotes the number of processes that act upon this rule. By restricting the set $Tr(\mathbf{p})$ to contain transitions whose factors sum up to $N(\mathbf{p})$, we ensure that in a transition, every process takes a step. This captures the semantics of synchronous computation. A transition $tr \in Tr(\mathbf{p})$ is *enabled* in a tuple $(\boldsymbol{\kappa}, \mathbf{p})$, where $\boldsymbol{\kappa}$ is an $|\mathcal{L}|$ -dimensional vector of counters and $\mathbf{p} \in \mathbf{P}_{RC}$ an admissible valuation, iff for every $r \in \mathcal{R}$, such that $tr(r) > 0$, it holds that $(\boldsymbol{\kappa}, \mathbf{p}) \models r.\varphi$, and for every $\ell \in \mathcal{L}$, we have $\boldsymbol{\kappa}[\ell] = \sum_{r \in \mathcal{R} \wedge r.\text{from}=\ell} tr(r)$. The former condition ensures that processes only use rules whose guards are satisfied, and the latter that every process moves in an enabled transition.

Given a transition $tr \in Tr(\mathbf{p})$, we define the *origin* $o(tr) = (\boldsymbol{\kappa}, \mathbf{p})$ of tr , where for every location $\ell \in \mathcal{L}$, we have $\boldsymbol{\kappa}[\ell] = \sum_{r \in \mathcal{R} \wedge r.\text{from}=\ell} tr(r)$, and the *goal* $g(tr) = (\boldsymbol{\kappa}', \mathbf{p})$ of tr , where for every location $\ell \in \mathcal{L}$, we have $\boldsymbol{\kappa}'[\ell] = \sum_{r \in \mathcal{R} \wedge r.\text{to}=\ell} tr(r)$. The origin $o(tr)$ is the unique tuple $(\boldsymbol{\kappa}, \mathbf{p})$ where the transition tr is enabled, while its goal $g(tr)$ is the unique tuple $(\boldsymbol{\kappa}', \mathbf{p})$ that is obtained by applying the transition tr to its origin $o(tr)$. The *transition relation* $R(\mathbf{p})$ is the relation $R(\mathbf{p}) \subseteq \Sigma(\mathbf{p}) \times Tr(\mathbf{p}) \times \Sigma(\mathbf{p})$, such that $\langle \sigma, tr, \sigma' \rangle \in R(\mathbf{p})$ iff $\sigma = o(tr)$ is the origin and $\sigma' = g(tr)$ is the goal of the transition tr .

5 Synchronous Transition Systems

Let $\text{rSTA} = (\mathcal{L}, \mathcal{I}, \mathcal{R}^\Delta, \Delta, \Pi, RC, \text{Env}^\Delta)$ be a receive STA, and $\mathbf{p} \in \mathbf{P}_{RC}$ an admissible valuation of the parameter vector $\boldsymbol{\pi}$. A *synchronous transition system* (or *system*), w.r.t. an admissible valuation $\mathbf{p} \in \mathbf{P}_{RC}$ and an rSTA is the triple $\text{STS}(\text{rSTA}, \mathbf{p}) = \langle S(\mathbf{p}), S_0(\mathbf{p}), T(\mathbf{p}) \rangle$, representing a system of $N(\mathbf{p})$ processes whose behavior is modeled using the rSTA, where $S(\mathbf{p})$ is the set of *states*, $S_0(\mathbf{p})$ is the set of *initial states*, and $T(\mathbf{p})$ is the *transition relation*.

Recall that the environment assumption Env^Δ of the rSTA is the conjunction $\text{Env}^\Delta \equiv \text{Env}_{\text{CP}} \wedge \text{Env}_{\text{EP}}$. A *state* $s \in S(\mathbf{p})$ is a tuple $s = \langle \boldsymbol{\ell}, \mathbf{nr}_1, \dots, \mathbf{nr}_{N(\mathbf{p})}, \mathbf{p} \rangle$, where $\boldsymbol{\ell} \in \mathcal{L}^{N(\mathbf{p})}$ is an $N(\mathbf{p})$ -dimensional vector of locations, and $\mathbf{nr}_i \in \mathbb{N}^{|\mathcal{M}|}$, for $1 \leq i \leq N(\mathbf{p})$, is a vector of valuations of the receive variables $\text{nr}(m)$, with $m \in \mathcal{M}$, for each process i , such that $s \models \text{Env}_{\text{CP}}$. In a state $s \in S(\mathbf{p})$, the vector $\boldsymbol{\ell}$ of locations is used to store the current location $s.\boldsymbol{\ell}[i] \in \mathcal{L}$ for each process i , while the vector $\mathbf{nr}_i \in \mathbb{N}^{|\mathcal{M}|}$ stores the values of the receive variables for each process i , with $1 \leq i \leq N(\mathbf{p})$. Further, each state $s \in S(\mathbf{p})$ satisfies Env_{CP} .

To formally define that a state $s \in S(\mathbf{p})$ satisfies the environment constraint Env_{CP} , we define the semantics of c -propositions w.r.t. states $s \in S(\mathbf{p})$. Let $\text{counters}_{\mathbf{p}} : S(\mathbf{p}) \times \mathcal{L} \rightarrow \mathbb{N}$ denote a mapping that maps a state $s \in S(\mathbf{p})$ and a location $\ell \in \mathcal{L}$ to the number of processes that are in location ℓ in the state s , that is, $\text{counters}_{\mathbf{p}}(s, \ell) = |\{i \mid 1 \leq i \leq N(\mathbf{p}) \wedge s.\ell[i] = \ell\}|$. Further, let $\kappa(s) \in \mathbb{N}^{|\mathcal{L}|}$ denote the $|\mathcal{L}|$ -dimensional vector of counters w.r.t. the state $s \in S(\mathbf{p})$, where for every location $\ell \in \mathcal{L}$, we have that $\kappa(s)[\ell]$ stores the number of processes that are in location ℓ in the state s , that is, $\kappa(s)[\ell] = \text{counters}_{\mathbf{p}}(s, \ell)$. We say that $s \models \#L \geq \mathbf{a} \cdot \boldsymbol{\pi} + b$ iff $(\kappa(s), s.\mathbf{p}) \models \#L \geq \mathbf{a} \cdot \boldsymbol{\pi} + b$. A state $s \in S(\mathbf{p})$ satisfies the environment constraints Env_{CP} , that is, $s \models \text{Env}_{\text{CP}}$ iff $(\kappa(s), s.\mathbf{p}) \models \text{Env}_{\text{CP}}$.

In an *initial state* $s_0 \in S_0(\mathbf{p})$, the vector ℓ of locations stores only initial locations, i.e., $\ell[i] \in \mathcal{I}$, for $1 \leq i \leq N(\mathbf{p})$, and all receive variables of all processes are initialized to 0. Formally, a state $s_0 = \langle \ell, \mathbf{nr}_1, \dots, \mathbf{nr}_{N(\mathbf{p})}, \mathbf{p} \rangle$ is *initial*, i.e., $s_0 \in S_0(\mathbf{p})$, if $s_0.\ell \in \mathcal{I}^{N(\mathbf{p})}$ and $s_0.\mathbf{nr}_i[m] = 0$, for $1 \leq i \leq N(\mathbf{p})$ and $m \in \mathcal{M}$.

We now define the transition relation $T(\mathbf{p}) \subseteq S(\mathbf{p}) \times S(\mathbf{p})$, where we will use the environment constraint Env_{EP} to restrict the values of the receive variables. A transition $(s, s') \in T(\mathbf{p})$ encodes one round in the execution of the distributed algorithm. In a round, the processes send and receive messages, and update their variables based on the received messages. Further, all the messages sent in the current round are received in the same round. The process variable updates are captured by moving processes from one location to another, based on the values of the receive variables. The *transition relation* $T(\mathbf{p})$ is a binary relation $T(\mathbf{p}) \subseteq S(\mathbf{p}) \times S(\mathbf{p})$, where $(s, s') \in T(\mathbf{p})$ iff for every process i , with $1 \leq i \leq N(\mathbf{p})$:

1. $0 \leq s'.\mathbf{nr}_i[m] \leq N(\mathbf{p})$, such that $(s'.\mathbf{nr}_i, \kappa(s), s.\mathbf{p}) \models \text{Env}_{\text{EP}}$, for $m \in \mathcal{M}$,
2. there exists $r^\Delta \in \mathcal{R}^\Delta$ such that:
 - $s.\ell[i] = r^\Delta.\text{from}$,
 - $(s'.\mathbf{nr}_i, \kappa(s), s.\mathbf{p}) \models r^\Delta.\varphi$,
 - $s'.\ell[i] = r^\Delta.\text{to}$.
3. $s'.\mathbf{p} = s.\mathbf{p}$ and $s' \models \text{Env}_{\text{CP}}$.

In a transition $(s, s') \in T(\mathbf{p})$, the receive variables and locations of each process are updated. That is, the value $s'.\mathbf{nr}_i[m]$ of the receive variable $\mathbf{nr}(m)$ of process i is assigned a value in the range from 0 to $N(\mathbf{p})$ non-deterministically, such that the environment constraint Env_{EP} is satisfied. This ensures that the number of received messages of type m is non-negative, that it does not exceed the number of participating processes, and that the receive variables of each process are assigned values that satisfy the constraints of the environment assumption. In the case of the synchronous computation model, this captures that all messages sent by correct processes in a round are received in the same round, and that the number of messages of type m , received by process i , is bounded by above by the total number of messages of type m , sent by both correct and faulty processes. To update the locations, each process i picks a rule $r^\Delta \in \mathcal{R}^\Delta$ that it applies to update its location, if the process i is in location $r^\Delta.\text{from}$ in the state s , and if the newly assigned values of the receive variables of process i in the state s' satisfy the receive guard $r^\Delta.\varphi$. If this is the case, the process i updates

its location to $r^\Delta.to$ in the state s' . The parameter values remain unchanged, and we require that the state s' satisfies Env_{CP} , i.e., it is a valid state.

6 Abstracting rSTA to STA

Given an rSTA, our goal is to construct an STA, which differs from the rSTA only in the guards on its rules and the environment assumption. For each rule $r^\Delta \in \mathcal{R}^\Delta$ in the rSTA, whose guard $r^\Delta.\varphi$ is a receive guard, we will construct a rule $r \in \mathcal{R}$ in the STA, such that the guard $r.\varphi$ is a Boolean combination of c -propositions. We will perform the abstraction in two steps: (i) we will strengthen each receive guard $r^\Delta.\varphi$, occurring on the rules $r^\Delta \in \mathcal{R}^\Delta$ of the rSTA, with the constraints imposed by the faulty environment and the synchronous computation model, encoded in the environment assumption Env^Δ , and (ii) we will eliminate the receive variables from the receive guards and environment assumptions of rSTA to obtain the guards and environment assumption of STA.

6.1 Guard Strengthening

Let $\text{rSTA} = (\mathcal{L}, \mathcal{I}, \mathcal{R}^\Delta, \Delta, \Pi, RC, \text{Env}^\Delta)$ be a receive STA, where the rules $r^\Delta \in \mathcal{R}^\Delta$ have guards containing expressions over the receive variables $\text{nr}(m) \in \Delta$, and where the environment assumption $\text{Env}^\Delta \equiv \text{Env}_{\text{CP}} \wedge \text{Env}_{\text{EP}}$ is a conjunction of two environment constraints, Env_{CP} and Env_{EP} , where the latter restricts the values of the receive variables. Recall that in Section 3.2, we defined different environment constraints Env_{EP} for the different fault models. In general, these constraints express that for each message type $m \in \mathcal{M}$, the receive variable $\text{nr}(m)$ is assigned a value which is greater or equal to the number of messages of type m sent by correct processes, and which is smaller or equal to the total number of messages of type m , sent by both correct and faulty processes (e.g., $\#\text{sent}(m) \leq \text{nr}(m) \leq \#\text{sent}(m) + \#\text{sent}_{\text{cr}}(m)$ for crash faults). As a first step towards eliminating the receive variables from the receive guards, we strengthen the rules from the set \mathcal{R}^Δ , such that we add the environment constraints Env_{EP} to their guards in order to bound the values of the receive variables.

Definition 1. *Given $r^\Delta \in \mathcal{R}^\Delta$, its strengthened rule is $\hat{r}^\Delta = \text{strengthen}(r^\Delta)$, such that: $\hat{r}^\Delta.from = r^\Delta.from$, $\hat{r}^\Delta.to = r^\Delta.to$, $\hat{r}^\Delta.\varphi = r^\Delta.\varphi \wedge \text{Env}_{\text{EP}}$.*

We denote by $\hat{\mathcal{R}}^\Delta = \{\text{strengthen}(r^\Delta) \mid r^\Delta \in \mathcal{R}^\Delta\}$ the set of strengthened rules in $\text{rSTA} = (\mathcal{L}, \mathcal{I}, \mathcal{R}^\Delta, \Delta, \Pi, RC, \text{Env}^\Delta)$, where $\text{Env}^\Delta \equiv \text{Env}_{\text{CP}} \wedge \text{Env}_{\text{EP}}$.

6.2 Eliminating the Receive Variables

Let $\text{rSTA} = (\mathcal{L}, \mathcal{I}, \mathcal{R}^\Delta, \Delta, \Pi, RC, \text{Env}^\Delta)$ be a receive STA, and let $\hat{\mathcal{R}}^\Delta$ be the set of strengthened rules (Definition 1). We define an STA $= (\mathcal{L}, \mathcal{I}, \mathcal{R}, \Pi, RC, \text{Env})$ whose locations, initial locations, and parameters are the same as in rSTA, while we construct the rules \mathcal{R} and the environment assumption Env of the STA below.

Recall that $\text{Env}^\Delta \equiv \text{Env}_{\text{CP}} \wedge \text{Env}_{\text{EP}}$. To define the environment assumption Env of the constructed STA, we set $\text{Env} \equiv \text{Env}_{\text{CP}}$. Before we define the rules of the constructed STA, we define the mapping eliminate .

Definition 2. Let ϕ be a propositional formula over r -, c -, and e -propositions. Let $\delta = \langle \text{nr}(m_1), \dots, \text{nr}(m_{|\mathcal{M}|}) \rangle$ denote the $|\mathcal{M}|$ -dimensional receive variables vector, and QE denote the quantifier elimination procedure for Presburger arithmetic. The formula $\text{eliminate}(\phi) = \text{QE}(\exists \delta \phi)$ is a quantifier-free formula, with no occurrence of receive variables $\text{nr}(m) \in \Delta$, which is logically equivalent to $\exists \delta \phi$.

To construct a rule $r \in \mathcal{R}$ of an STA, given a rule $r^\Delta \in \mathcal{R}^\Delta$ of an rSTA, we will apply the mapping eliminate to each guard of the strengthened rule $\hat{r}^\Delta \in \hat{\mathcal{R}}^\Delta$, where $\hat{r}^\Delta = \text{strengthen}(r^\Delta)$. The result of quantifier elimination is a quantifier-free formula over c -propositions, which is logically equivalent to $\exists \delta \hat{r}^\Delta$.

Definition 3. Given $r^\Delta \in \mathcal{R}^\Delta$, its constructed rule is $r = \text{construct}(r^\Delta) \in \mathcal{R}$, such that: $r.\text{from} = r^\Delta.\text{from}$, $r.\text{to} = r^\Delta.\text{to}$, $r.\varphi = \text{eliminate}(\hat{r}^\Delta.\varphi)$, where $\hat{r}^\Delta = \text{strengthen}(r^\Delta)$.

Proposition 1. For every strengthened rule $\hat{r}^\Delta \in \hat{\mathcal{R}}^\Delta$ and every tuple $(\mathbf{d}, \boldsymbol{\kappa}, \mathbf{p})$, where $\mathbf{d} \in \mathbb{N}^{|\mathcal{M}|}$, $\boldsymbol{\kappa} \in \mathbb{N}^{|\mathcal{C}|}$, and $\mathbf{p} \in \mathbf{P}_{RC}$, we have:

$$(\mathbf{d}, \boldsymbol{\kappa}, \mathbf{p}) \models \hat{r}^\Delta.\varphi \quad \text{implies} \quad (\boldsymbol{\kappa}, \mathbf{p}) \models \text{eliminate}(\hat{r}^\Delta.\varphi)$$

Proposition 1 is a consequence of quantifier elimination. Note that the converse of this proposition does not hold in general. That is, $(\boldsymbol{\kappa}, \mathbf{p}) \models \text{eliminate}(\hat{r}^\Delta.\varphi)$ does not imply that $(\mathbf{d}, \boldsymbol{\kappa}, \mathbf{p}) \models \hat{r}^\Delta.\varphi$, for every $\mathbf{d} \in \mathbb{N}^{|\mathcal{M}|}$. However, by quantifier elimination, we have that $(\boldsymbol{\kappa}, \mathbf{p}) \models \text{eliminate}(\hat{r}^\Delta.\varphi)$ implies $(\boldsymbol{\kappa}, \mathbf{p}) \models \exists \delta \hat{r}^\Delta.\varphi$.

6.3 Soundness and Completeness

This construction of an STA is sound and complete. That is, given a rSTA and an admissible valuation $\mathbf{p} \in \mathbf{P}_{RC}$, we show that there exists a bisimulation relation between the system $\text{STS}(\text{rSTA}, \mathbf{p})$, induced by rSTA and \mathbf{p} , and a counter system $\text{CS}(\text{STA}, \mathbf{p})$, induced by the constructed STA and \mathbf{p} . The existence of a bisimulation implies that $\text{STS}(\text{rSTA}, \mathbf{p})$ and $\text{CS}(\text{STA}, \mathbf{p})$ satisfy the same CTL* formulas [3]. To express temporal formulas, as atomic propositions we use the c -propositions from the set CP. We define two labeling functions, $\lambda_{S(\mathbf{p})}$ and $\lambda_{\Sigma(\mathbf{p})}$, where $\lambda_{S(\mathbf{p})} : S(\mathbf{p}) \rightarrow 2^{\text{CP}}$ assigns to a state $s \in S(\mathbf{p})$ the set of c -propositions that hold in it (the function $\lambda_{\Sigma(\mathbf{p})} : \Sigma(\mathbf{p}) \rightarrow 2^{\text{CP}}$ is defined analogously).

We introduce an *abstraction mapping* $\alpha_{\mathbf{p}} : S(\mathbf{p}) \rightarrow \Sigma(\mathbf{p})$ that maps states $s \in S(\mathbf{p})$ of $\text{STS}(\text{rSTA}, \mathbf{p})$ to configurations $\sigma \in \Sigma(\mathbf{p})$ of $\text{CS}(\text{STA}, \mathbf{p})$, such that $\sigma = \alpha_{\mathbf{p}}(s)$ iff $\sigma = (\boldsymbol{\kappa}(s), s.\mathbf{p})$. By the definition of the abstraction mapping $\alpha_{\mathbf{p}}$ and the semantics of c -propositions, we have that a state and its abstraction satisfy the same c -propositions. Further, given a configuration $\sigma \in \Sigma(\mathbf{p})$, we can construct a state $s \in S(\mathbf{p})$, such that $\sigma = \alpha_{\mathbf{p}}(s)$. While this is always possible, the constructed state s might not be reachable in any execution of the system $\text{STS}(\text{rSTA}, \mathbf{p})$. However, we can use the constraint Env_{EP} to restrict the value of the receive variables in the constructed state s , such that it is a valid state in the system $\text{STS}(\text{rSTA}, \mathbf{p})$. The main result of this section is stated below. The detailed proof of this result can be found in the first author's PhD thesis.

Theorem 1. The binary relation $B(\mathbf{p}) = \{(s, \sigma) \mid s \in S(\mathbf{p}), \sigma \in \Sigma(\mathbf{p}), \sigma = \alpha_{\mathbf{p}}(s)\}$ is a bisimulation relation.

7 Experimental Evaluation

To show the usefulness of translating rSTA to STA, we: (i) encoded synchronous fault-tolerant distributed algorithms using rSTA, (ii) implemented the method from Section 6 in a prototype, (iii) compared the output to the existing manual encodings from [34], some of which are artifacts of the experimental evaluation from [36] and were given as examples throughout this paper, and (iv) verified the properties of the generated STA using the technique from [36].

Encoding Algorithms as rSTA. We extended the STA encoding from [36], to support (i) declarations of receive variables and (ii) constraints given by the environment assumption. The algorithms we encoded are listed in Table 1, and their rSTA can be found in [35]. For each of them, there already existed a manually produced STA [34]. The manually produced rSTA and STA have the same structure w.r.t. locations and rules, and differ only in the guards that occur on the rules: in the rSTA, we have receive guards, which are Boolean combinations of r -propositions and c -propositions, while in the manually encoded STA, the guards are Boolean combinations of c -propositions.

Applying Quantifier Elimination. We implemented a script that parses the input rSTA and creates an STA whose rules have guards that are Boolean combinations of c -propositions, according to the abstraction from Section 6. To automate the quantifier elimination step, we applied Z3 [16] tactics for quantifier elimination [10, 11], to formulas of the form $\exists \delta \hat{r}^\Delta.\varphi$, where $\hat{r}^\Delta.\varphi \equiv r^\Delta.\varphi \wedge \text{Env}_{\text{EP}}$ is the strengthened guard of the receive guard $r^\Delta.\varphi$, for $r^\Delta \in \mathcal{R}^\Delta$. For all our benchmarks, the STA is generated within seconds, as reported in Table 1.

Analyzing the Automatically Generated STA. We compared the guards of the automatically generated STA (autoSTA) to the manually encoded STA (manSTA). Syntactically, the guards of autoSTA are larger in general, as they contain additional constraints that result from quantifier elimination. Semantically, we check whether the guards for the autoSTA imply the guards of the manSTA. For each automatically generated guard φ_{auto} , we check whether its corresponding guard φ_{man} from the manual encoding is implied by φ_{auto} , for all values of the parameters and number of sent messages by checking the validity of the formula:

$$\forall \mathbf{p} \in \mathbf{P}_{RC} \forall L_1 \dots \forall L_{|\mathcal{M}|} \varphi_{\text{auto}}(L_1, \dots, L_{|\mathcal{M}|}) \rightarrow \varphi_{\text{man}}(L_1, \dots, L_{|\mathcal{M}|}) \quad (3)$$

where $L_j = \text{sent}(m_j)$, for $m_j \in \mathcal{M}$ and $1 \leq j \leq |\mathcal{M}|$, denotes the set of locations where processes send messages of type m_j . We automate the validity check of (3) using an SMT solver, such as Z3, to check the unsatisfiability of its negation. With this check we are able to either verify that the earlier manSTA faithfully model the benchmark algorithms, or detect discrepancies, which we investigated further. Our translation technique produces the strongest possible guards, due to the soundness and completeness result. Hence, we expected that the implication holds for all the guards of all the benchmarks we considered. This

Table 1. The algorithms we encoded as `rSTA` and the results of applying the verification technique from [36]. The column `QE` states the time needed to produce an `autoSTA` from an `rSTA`. The column \Rightarrow states if (3) is valid all, some, or none of guards. We report on the time it took the solvers `Z3` and `CVC4` to (i) check the guard implications (only `Z3`), (ii) compute the diameter for the `autoSTA`, and (iii) check the safety properties of the `autoSTA`, (iv) compute the diameter for the `manSTA`, (v) check the safety properties of the `manSTA`, using the SMT-based procedure from [36].

algorithm	QE		(i)		(ii) autoSTA				(iii) autoSTA		(iv) manSTA		(v) manSTA	
	Z3	\Rightarrow	\Rightarrow time		d time		BMC time		d time		BMC time			
			Z3	d	Z3	CVC4	Z3	CVC4	Z3	CVC4	Z3	CVC4		
RB	0.16s	all	0.18s	2	0.09s	0.26s	0.03s	0.03s	0.07s	0.27s	0.02s	0.03s		
HybridRB	0.39s	all	0.41s	2	0.14s	0.75s	0.03s	0.06s	0.09s	0.67s	0.03s	0.05s		
OmitRB	0.34s	all	0.36s	2	0.11s	0.69s	0.03s	0.05s	0.09s	0.67s	0.02s	0.04s		
FairCons	0.25s	all	0.44s	2	0.17s	2.82s	0.07s	0.16s	0.14s	2.68s	0.06s	0.14s		
FloodMin, $k = 1$	0.10s	all	0.19s	2	0.07s	0.25s	0.06s	0.11s	0.06s	0.25s	0.06s	0.09s		
FloodMin, $k = 2$	0.26s	all	0.35s	2	0.13s	1.72s	0.07s	0.19s	0.15s	2.22s	0.06s	0.17s		
FMinOmit, $k = 1$	0.10s	all	0.13s	1	0.03s	0.03s	0.01s	0.01s	0.06s	0.04s	0.01s	0.01s		
FMinOmit, $k = 2$	0.27s	all	0.26s	1	0.05s	0.08s	0.01s	0.03s	0.05s	0.08s	0.01s	0.03s		
FloodSet	0.20s	all	0.31s	2	0.11s	0.71s	0.07s	0.17s	0.10s	0.90s	0.06s	0.15s		
kSetOmit, $k = 1$	0.59s	all	0.52s	3	2.71s	53.36s	0.22s	0.85s	1.09s	1m8s	0.23s	0.81s		
kSetOmit, $k = 2$	1.43s	all	1.18s	–	t.o.	t.o.	–	–	t.o.	t.o.	–	–		
PhaseKing	1.19s	all	1.57s	4	3.53s	16.51s	0.24s	1.57s	3.67s	15.80s	0.25s	1.47s		
ByzKing	1.16s	all	1.58s	4	1.92s	1m19s	0.27s	1.97s	3.73s	38.50s	0.24s	2.26s		
HybridKing	3.59s	some	3.03s	4	0.33s	6.34s	0.18s	1.11s	t.o.	t.o.	–	–		
OmitKing	3.09s	all	2.79s	4	0.26s	6.12s	0.15s	0.91s	1h15m	t.o.	9.08s	1m27s		
PhaseQueen	0.42s	all	0.90s	3	0.37s	4.46s	0.04s	0.61s	0.40s	4.72s	0.06s	0.50s		
ByzQueen	0.42s	all	0.91s	3	0.39s	17.15s	0.09s	0.58s	0.53s	10.6s	0.08s	0.61s		
HybridQueen	1.34s	some	1.77s	3	0.13s	2.04s	0.05s	0.37s	t.o.	t.o.	–	–		
OmitQueen	1.13s	all	1.56s	3	0.13s	2.18s	0.20s	0.46s	0.57s	8.87s	0.27s	1.21s		

is however not the case for the algorithms `HybridKing` and `HybridQueen` which are designed to tolerate hybrid faults, in particular, send omissions and Byzantine faults. There, we found that one automatically generated guard does not imply its corresponding manual guard, and concluded that this is due to a flaw in the manual encoding by manual inspection. We found a similar problem with a missing rule in the (purely) Byzantine versions of these algorithms, namely `ByzKing` and `ByzQueen`. By adding these rules and correcting the appropriate manual guards, we were able to establish the validity of (3) for all guards.

Model Checking of Safety Properties. We gave the `STA` we obtained as output of our translation procedure as input to the bounded model checking tool from [36], which computes a diameter of a counter system and performs bounded model checking for safety properties. The experiments were run on a machine with 2.8 GHz Quad-Core Intel(R) Core(TM) i7 CPU and 16GB. The results of applying the SMT-based procedure from [36] to the `autoSTA`, as well as to the extended set [34] of `manSTA` from [36], are presented in Table 1. The timeout, denoted by t.o. in the table, was set to 24 hours. For all algorithms, we note that bounded model checking with both `Z3` and `CVC4` performs similarly for both `autoSTA` and `manSTA`. For computing the diameter, we observe that for the algorithms: `RB` [21]

(Fig. 5), HybridRB, OmitRB [9], FairCons [33], FloodMin, for $k = 1$ (Fig. 3) and $k = 2$ [28], FMinOmit, for $k = 1$ (Fig. 4) and $k = 2$ [28], kSetOmit, for $k = 2$ [33], FloodSet [28], PhaseKing [7], and PhaseQueen [6] (Fig. 1), we obtain comparable results on both the autoSTA and manSTA. For the other algorithms, we found:

- computing the diameter for the autoSTA of kSetOmit, with $k = 1$ [33], is slightly slower with Z3 and slightly faster with CVC4 than for the manSTA;
- Z3 performs better when computing the diameter for the autoSTA than for the manSTA of both ByzKing and ByzQueen [9], while CVC4 performs worse. Note that in Table 1 we report the times for the manSTA of ByzKing and ByzQueen that have missing rules. After adding the rules to the manSTA, computing the diameter on the autoSTA is still faster with both solvers;
- Z3 and CVC4 compute the diameter for the autoSTA of HybridKing and HybridQueen [9] within seconds, in contrast to both timing out for the manSTA;
- computing the diameter with Z3 is significantly faster for the autoSTA than for the manSTA of OmitKing [9]. CVC4 computes the diameter for autoSTA of OmitKing, while for manSTA it times out. The computed diameter $d = 4$ for autoSTA is smaller than the diameter 8, computed for manSTA;
- Z3 and CVC4 compute the diameter for the autoSTA of OmitQueen [9] faster than for manSTA.

8 Conclusions

We established a fully automated pipeline that for a synchronous distributed algorithm: (1) starts from a formal model that captures its pseudo code, (2) produces a formal model suitable for verification, and (3) automatically verifies its safety properties. Our technique thus closes the gap between the original description of an algorithm (using received messages) and the synchronous threshold automaton of the algorithm given as an input to a verification tool.

There are two major differences to the asynchronous case considered in [37]. First, the asynchronous model uses interleaving semantics, while in the synchronous model all processes take a step in a transition. Second, in the asynchronous model, there are no limitations when a message will be delivered. The lower bound on the number of received messages, given in the synchronous model by the number of sent messages by correct processes, is only *eventually* satisfied in the asynchronous model, and thus is not used in the process of eliminating the receive variables from the receive guards.

We did extensive experimental evaluation of our method. We attribute the better performance of the bounded model checking technique from [36] on the automatically generated STA to the fact that the automatically generated guards contain more additional constraints, coming from the environment assumption, which help guide the SMT solvers. Moreover, not only do we obtain the diameter bounds faster, we also obtain better bounds for the automatically generated STA of some benchmarks. These findings confirm the conjecture that manual encoding of distributed algorithms is a tedious and error-prone task and suggest that there is a real benefit of producing guards automatically.

References

1. Aminof, B., Rubin, S., Stoilkovska, I., Widder, J., Zuleger, F.: Parameterized model checking of synchronous distributed algorithms by abstraction. In: VMCAI. LNCS, vol. 10747, pp. 1–24. Springer (2018)
2. Attiya, H., Welch, J.: Distributed Computing. Wiley, 2nd edn. (2004)
3. Baier, C., Katoen, J.P.: Principles of model checking. MITP (2008)
4. Bakst, A., von Gleissenthall, K., Kici, R.G., Jhala, R.: Verifying distributed programs via canonical sequentialization. PACMPL **1**(OOPSLA), 110:1–110:27 (2017)
5. Balasubramanian, A.R., Esparza, J., Lazić, M.: Complexity of Verification and Synthesis of Threshold Automata. In: ATVA (2020), (to appear)
6. Berman, P., Garay, J.A., Perry, K.J.: Asymptotically Optimal Distributed Consensus. Tech. rep., Bell Labs (1989), plan9.bell-labs.co/who/garay/asopt.ps
7. Berman, P., Garay, J.A., Perry, K.J.: Towards Optimal Distributed Consensus (Extended Abstract). In: FOCS. pp. 410–415 (1989)
8. Bertrand, N., Konnov, I., Lazić, M., Widder, J.: Verification of randomized consensus algorithms under round-rigid adversaries. In: CONCUR. pp. 33:1–33:15 (2019)
9. Biely, M., Schmid, U., Weiss, B.: Synchronous consensus under hybrid process and link failures. Theoretical Computer Science **412**(40), 5602–5630 (2011)
10. Bjørner, N.: Linear quantifier elimination as an abstract decision procedure. In: IJCAR. pp. 316–330 (2010)
11. Bjørner, N., Janota, M.: Playing with quantified satisfaction. In: LPAR. pp. 15–27 (2015)
12. Bouajjani, A., Enea, C., Ji, K., Qadeer, S.: On the completeness of verifying message passing programs under bounded asynchrony. In: CAV. pp. 372–391 (2018)
13. Chaouch-Saad, M., Charron-Bost, B., Merz, S.: A reduction theorem for the verification of round-based distributed algorithms. In: RP. LNCS, vol. 5797, pp. 93–106 (2009)
14. Cooper, D.C.: Theorem proving in arithmetic without multiplication. Machine intelligence **7**(91-99), 300 (1972)
15. Damian, A., Drăgoi, C., Militaru, A., Widder, J.: Communication-closed asynchronous protocols. In: CAV (2). LNCS, vol. 11562, pp. 344–363. Springer (2019)
16. De Moura, L., Bjørner, N.: Z3: An efficient SMT solver. In: TACAS, LNCS, vol. 1579, pp. 337–340 (2008)
17. Drăgoi, C., Henzinger, T.A., Veith, H., Widder, J., Zufferey, D.: A logic-based framework for verifying consensus algorithms. In: VMCAI. LNCS, vol. 8318, pp. 161–181 (2014)
18. Fischer, M.J., Lynch, N.A., Paterson, M.S.: Impossibility of distributed consensus with one faulty process. J. ACM **32**(2), 374–382 (1985)
19. v. Gleissenthall, K., Gökhan Kici, R., Bakst, A., Stefan, D., Jhala, R.: Pretend synchrony. In: POPL (2019), (to appear)
20. Hawblitzel, C., Howell, J., Kapritsos, M., Lorch, J.R., Parno, B., Roberts, M.L., Setty, S., Zill, B.: Ironfleet: Proving safety and liveness of practical distributed systems. Commun. ACM **60**(7), 83–92 (Jun 2017)
21. K., Srikanth, T., Toueg, S.: Optimal clock synchronization. J. ACM **34**(3), 626–645 (1987)
22. Konnov, I., Lazić, M., Veith, H., Widder, J.: A short counterexample property for safety and liveness verification of fault-tolerant distributed algorithms. In: POPL. pp. 719–734 (2017)

23. Konnov, I., Veith, H., Widder, J.: On the completeness of bounded model checking for threshold-based distributed algorithms: Reachability. *Information and Computation* **252**, 95–109 (2017), <http://dx.doi.org/10.1016/j.ic.2016.03.006>
24. Kopetz, H., Grünsteidl, G.: TTP – A Protocol for Fault-Tolerant Real-Time Systems. *IEEE Computer* **27**(1), 14–23 (1994), <http://doi.ieeecomputersociety.org/10.1109/2.248873>
25. Kragl, B., Qadeer, S., Henzinger, T.A.: Synchronizing the asynchronous. In: CONCUR. pp. 21:1–21:17 (2018)
26. Kukovec, J., Konnov, I., Widder, J.: Reachability in parameterized systems: All flavors of threshold automata. In: CONCUR. LIPIcs, vol. 118, pp. 19:1–19:17 (2018)
27. Lincoln, P., Rushby, J.: A formally verified algorithm for interactive consistency under a hybrid fault model. In: FTCS. pp. 402–411 (1993)
28. Lynch, N.: *Distributed Algorithms*. Morgan Kaufman (1996)
29. Maric, O., Sprenger, C., Basin, D.A.: Cutoff bounds for consensus algorithms. In: CAV. pp. 217–237 (2017)
30. Presburger, M.: Über die vollständigkeit eines gewissen systems der arithmetik ganzer zahlen, in welchem die addition als einzige operation hervortritt. *Comptes Rendus du I congres de Mathématiciens des Pays Slaves* pp. 92–101 (1929)
31. Pugh, W.: A practical algorithm for exact array dependence analysis. *Communications of the ACM* **35**(8), 102–114 (1992)
32. Rahli, V., Guaspari, D., Bickford, M., Constable, R.L.: Formal specification, verification, and implementation of fault-tolerant systems using EventML. *ECEASST* **72** (2015)
33. Raynal, M.: *Fault-Tolerant Agreement in Synchronous Message-Passing Systems*. Synthesis Lectures on Distributed Computing Theory, Morgan & Claypool Publishers (2010)
34. Stoilkovska, I.: Manually Encoded Synchronous Threshold Automata. <https://github.com/istoilkovska/syncTA/algorithms>, [Online; accessed October 2020]
35. Stoilkovska, I.: Receive Synchronous Threshold Automata. <https://github.com/istoilkovska/syncTA/receiveSTA>, [Online; accessed October 2020]
36. Stoilkovska, I., Konnov, I., Widder, J., Zuleger, F.: Verifying Safety of Synchronous Fault-Tolerant Algorithms by Bounded Model Checking. In: TACAS. pp. 357–374 (2019)
37. Stoilkovska, I., Konnov, I., Widder, J., Zuleger, F.: Eliminating Message Counters in Threshold Automata. In: ATVA (2020), (to appear)
38. Wilcox, J.R., Woos, D., Panchekha, P., Tatlock, Z., Wang, X., Ernst, M.D., Anderson, T.E.: Verdi: a framework for implementing and formally verifying distributed systems. In: PLDI. pp. 357–368 (2015)