

1 Characterizing consensus in the Heard-Of model

2 **A. R. Balasubramanian**

3 Technical University of Munich, Germany

4 **Igor Walukiewicz**

5 CNRS, LaBRI, University of Bordeaux, France

6 — Abstract —

7 The Heard-Of model is a simple and relatively expressive model of distributed computation. Because
8 of this, it has gained a considerable attention of the verification community. We give a characterization
9 of all algorithms solving consensus in a fragment of this model. The fragment is big enough to cover
10 many prominent consensus algorithms. The characterization is purely syntactic: it is expressed in
11 terms of some conditions on the text of the algorithm.

12 **2012 ACM Subject Classification** Theory of computation → Distributed computing models; Software
13 and its engineering → Software verification; Theory of computation → Logic and verification

14 **Keywords and phrases** consensus problem, Heard-Of model, verification.

15 **Digital Object Identifier** 10.4230/LIPIcs.CONCUR.2020.9

16 **Related Version** A full version of this paper is available at <https://arxiv.org/abs/2004.09621>.

17 **Funding** *A. R. Balasubramanian*: UMI ReLaX, ERC Advanced Grant 787367 (PaVeS)

18 *Igor Walukiewicz*: ANR FREDDA ANR-17-CE40-0013

19 **1** Introduction

20 Most distributed algorithms solving problems like consensus, leader election, set agreement,
21 or renaming are essentially one iterated loop. Yet, their behavior is difficult to understand due
22 to unbounded number of processes, asynchrony, failures, and other aspects of the execution
23 model. The general context of this work is to be able to say what happens when we change
24 some of the parameters: modify an algorithm or the execution model. Ideally we would like
25 to characterize the space of all algorithms solving a particular problem.

26 To approach this kind of questions, one needs to restrict to a well defined space of all
27 distributed algorithms and execution contexts. In general this is an impossible requirement.
28 Yet the distributed algorithms community has come up with some settings that are expressive
29 enough to represent interesting cases and limited enough to start quantifying over “all
30 possible” distributed algorithms [11, 40, 1].

31 In this work we consider the consensus problem in the Heard-Of model [11]. *Consensus*
32 *problem* is a central problem in the field of distributed algorithms; it requires that all correct
33 processes eventually decide on one of the initial values. *The Heard-Of model* is a round- and
34 message-passing-based model. It can represent many intricacies of various execution models
35 and yet is simple enough to attempt to analyze it algorithmically [9, 14, 15, 28, 27]. Initially,
36 our goal was to continue the quest from [28] of examining what is algorithmically possible
37 to verify in the Heard-Of model. While working on this problem we have realized that a
38 much more ambitious goal can be achieved: to give a simple, and in particular decidable,
39 characterization of all consensus algorithms in well-defined fragments of the Heard-Of model.

40 The Heard-Of model is an open ended model: it does not specify what operations processes
41 can perform and what kinds of communication predicates are allowed. Communication
42 predicates in the Heard-Of model capture in an elegant way both synchrony degree and
43 failure model. In this work we fix the set of atomic communication predicates and atomic



© A. R. Balasubramanian and Igor Walukiewicz;

licensed under Creative Commons License CC-BY

31st International Conference on Concurrency Theory (CONCUR 2020).

Editors: Igor Konnov and Laura Kovács; Article No. 9; pp. 9:1–9:18

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

44 operations. We opted for a set sufficient to express most prominent consensus algorithms (cf.
45 Section 7), but we do not cover all operations found in the literature on the Heard-Of model.

46 Our characterization of algorithms that solve consensus is expressed in terms of syntactic
47 conditions both on the text of the algorithm, and on the constraints in the communication
48 predicate. It exhibits an interesting way all consensus algorithms should behave. One could
49 imagine that there can be a consensus algorithm that makes processes gradually converge
50 to a consensus: more and more processes adopting the same value. This is not the case.
51 A consensus algorithm, in models we study here, should have a fixed number of crucial
52 rounds where precise things are guaranteed to happen. Special rounds have been identified
53 for existing algorithms [33], but not their distribution over different phases. Additionally,
54 here we show that all algorithms should have this structure.

55 As an application of our characterization we can think of using it as an intermediate
56 step in analysis of more complicated settings than the Heard-Of model. An algorithm in
57 a given setting can be abstracted to an algorithm in the Heard-Of model, and then our
58 characterization can be applied. Instead of proving the original algorithm correct it is enough
59 to show that the abstraction is sound. For example, an approach reducing asynchronous
60 semantics to round based semantics under some conditions is developed in [8]. A recent
61 paper [13] gives a reduction methodology in a much larger context, and shows its applicability.
62 The goal language of the reduction is an extension of the Heard-Of model that is not covered
63 by our characterization. As another application, our characterization can be used to quickly
64 see if an algorithm can be improved by taking a less constrained communication predicate,
65 by adapting threshold constants, or by removing parts of code (c.f. Section 7).

66 *Organization of the paper.* In the next section we introduce the Heard-Of model and formulate
67 the consensus problem. In the four consecutive sections we present the characterizations for
68 the core model as well as for the extensions with timestamps, coordinators, and with both
69 timestamps and coordinators at the same time. We then give examples of algorithms that
70 are covered by these characterizations. Proofs can be found in the appendix, as well as in
71 the full version of the paper [3].

72 Related work

73 The celebrated FLP result [18] states that consensus is impossible to achieve in an asyn-
74 chronous system in presence of failures, even in the presence of one crash failure. There is a
75 considerable literature investigating the models in which the consensus problem is solvable.
76 Even closer in spirit to the present paper are results on weakest failure detectors required to
77 solve the problem [6, 19]. Another step closer are works providing generic consensus algo-
78 rithms that can be instantiated to give several known concrete algorithms [31, 22, 21, 5, 34, 33].
79 The present paper considers a relatively simple model, but gives a characterization result of
80 all possible consensus algorithms.

81 The cornerstone idea of the Heard-Of model is to represent both asynchrony and failures
82 by the constraints on the message loss expressed by communication predicates. This greatly
83 simplifies the model, that in turn is very useful for a kind of characterizations we present here.
84 Unavoidably, not all aspects of partial synchrony [17, 12] or failures [7] are covered by the
85 model. For example, after a crash it may be difficult for a process to get into initial state, or
86 in terms of the Heard-of model, do the same round as other processes [38, 8]. Faults are not
87 malicious: a sent value may be lost, but the value may not be modified during transmission.
88 These observations just underline that there is no universal model for distributed algorithms.
89 There exists several other proposals of relatively simple and expressible models [20, 40, 1, 32].


90 The Heard-Of model, while not perfect, is in our opinion representative enough to merit a
91 more detailed study.

92 On the verification side there are at least three approaches to analysis of the Heard-Of or
93 similar models. One is to use automatic theorem provers like Isabelle [10, 9, 14]. Another is
94 deductive verification methods applied to annotated programs [16, 15]. The closest to this
95 work is a model-checking approach [37, 28, 27, 2]. Particularly relevant here is the work
96 of Maric et al. [28]. who show cut-off results for a fragment of the Heard-Of model and
97 then perform verification on a resulting finite state system. Our fragment of the Heard-Of
98 model is incomparable with the one from that work, and arguably it has less restrictions
99 coming from purely technical issues in proofs. While trying to extend the scope of automatic
100 verification methods along the lines in the above papers, we have realized that in our case it
101 is possible to obtain a characterization result.

102 Of course there are also other models of distributed systems that are considered in the
103 context of verification. For example there has been big progress on verification of threshold
104 automata [26, 24, 25, 35, 4]. There are also other methods, as automatically generating
105 invariants for distributed algorithms [23, 39, 36], or verification in Coq proof assistant [41, 42].

106 2 Heard-Of model and the consensus problem

107 In the Heard-Of model a certain number of processes execute the same code synchronously.
108 An algorithm consists of a sequence of *rounds*, every process executes the same round at
109 the same time. The sequence of rounds, called *phase*, is repeated forever. In a round every
110 process sends the value of one of its variables to a communication medium, receives a multiset
of values, and uses it to adopt a new value (cf. Figure 1).



111 **Figure 1** A schema of an execution of a round and of a phase. In every round i every process
112 sends a value of its variable x_i , and sets its variable x_{i+1} depending on the received multiset of
113 values: H_i^j . At the beginning of the phase the value of *inp* is sent, at some round *inp* may be
114 updated; we use i_r for the index of this round. In the last round *dec* may be set. Both *inp* and *dec*
115 are not updated if the value is $?$, standing for undefined.

112 At the beginning every process has its initial value in variable *inp*. Every process is
113 expected to eventually set its decision variable *dec*. Every round is communication closed
114 meaning that a value sent in a round can only be received in the same round; if it is not
115 received it is lost. A *communication predicate* is used to express a constraint on acceptable
116 message losses. Algorithm 1 is a concrete simple example of a 2-round algorithm with two
117 rounds. In the first round the value of *inp* is sent, in the second the value of x_1 . We will
118 explain the algorithm later in the text.

119 We proceed with a description of the syntax and semantics of Heard-Of algorithms. Next
120 we define the consensus problem. In later sections we will extend the core language with

9:4 Consensus in the Heard-Of model

121 timestamps and coordinators.

■ **Algorithm 1** Parametrized OneThird algorithm [11], thr_1, thr_2 are constants from $(0, 1)$

```

122 send (inp)
    | if uni(H) ∧ |H| >  $thr_1 \cdot |\Pi|$  then  $x_1 := inp := \text{smor}(\mathbf{H})$ ;
    | if mult(H) ∧ |H| >  $thr_1 \cdot |\Pi|$  then  $x_1 := inp := \text{smor}(\mathbf{H})$ ;
123 send  $x_1$ 
    | if uni(H) ∧ |H| >  $thr_2 \cdot |\Pi|$  then  $dec := \text{smor}(\mathbf{H})$ ;
Communication predicate:  $F(\psi^1 \wedge F\psi^2)$ 
where:  $\psi^1 := (\varphi_{=} \wedge \varphi_{thr_1}, true)$  and  $\psi^2 := (\varphi_{thr_1}, \varphi_{thr_2})$ 

```

123 Syntax

124 An algorithm has one *phase* that consists of two or more rounds. In the first round each
125 process sends the value of *inp* variable, in the last round it can set the value of *dec* variable.

126 A phase is repeated forever, all processes execute the same round at the same time. A round
127 *i* is a send statement followed by a sequence of conditionals:

```

128 send  $x_{i-1}$ 
    | if  $cond_i^1(\mathbf{H})$  then  $x_i := op_i^1(\mathbf{H})$ ;
    |
    |
    | if  $cond_i^l(\mathbf{H})$  then  $x_i := op_i^l(\mathbf{H})$ ;

```

129 The variables are used in a sequence: first x_0 , which is *inp*, is sent and x_1 is set, then x_1 is
130 sent and x_2 is set, etc. (cf. Figure 1). There should be exactly one round (before the last
131 round) where *inp* is updated; the conditional lines in this round are:

```

132 if  $cond_{ir}^j(\mathbf{H})$  then  $x_{ir} := inp := op_{ir}^j(\mathbf{H})$ 
133

```

134 Since this is a special round, we use the index *ir* to designate this round number. In the last
135 round, only instructions setting variable *dec* can be present:

```

136 if  $cond_r^j(\mathbf{H})$  then  $dec := op_r^j(\mathbf{H})$ 
137

```

138 Because of this special form of the last round, a phase needs to have at least two rounds.
139 Of course one can also have a syntax and a characterization for one round algorithms, but
140 unifying the two hinders readability. Our fragment roughly corresponds to the fragment
141 from [28], without extra restrictions but with a less liberty at the fork point.

142 The intuition behind the syntax is that in the i^{th} round, after a process sends its value of
143 the x_{i-1} variable and receives a multiset **H**, it finds the first instruction whose condition is
144 satisfied by **H** and performs the corresponding assignment. Hence, even if multiple conditions
145 are satisfied by a multi-set, only the first such condition is executed.

146 As an example, consider Algorithm 1. It has two rounds, each begins with a **send**
147 statement. In the first round both x_1 and *inp* are set, in the second round *dec* is set. The
148 conditions talk about properties of the received **H** multiset; we describe them below.

149 In round *i* every process first sends the value of variable x_{i-1} , and then receives a multiset
150 of values **H** that it uses to set the value of the variable x_i . The possible tests on the received
151 set **H** are **uni**, **mult**, and $|\mathbf{H}| > thr \cdot |\Pi|$ saying respectively that: the multiset has only
152 one value; has more than one value; and that is of size $> thr \cdot n$ where n is the number of
153 processes and $0 \leq thr < 1$. The possible operations are **min**(**H**) resulting in the minimal
154 value in **H**, and **smor**(**H**) resulting in the minimal most frequent value in **H**. For example, the

155 first conditional line in Algorithm 1 tests if there is only one value in H , and if this value has
 156 multiplicity at least $thr_1 \cdot n$ in H ; if so inp and x_1 are set to this value, it does not matter if
 157 min or smor operation is used in this case. The test in the second line holds when received H
 158 set has at least two values and is of size at least $thr_1 \cdot n$. In this case x_1 is set to the smallest
 159 most frequent value in H .

160 In addition to description of rounds, an algorithm has also a communication predicate
 161 putting constraints on the behavior of the communication medium. A *communication*
 162 *predicate for a phase* with r rounds is a tuple $\psi = (\psi_1, \dots, \psi_r)$, where each ψ_i is a conjunction
 163 of atomic communication predicates that we specify later. A *communication predicate for an*
 164 *algorithm* is

$$165 \quad (\overline{G\psi}) \wedge (F(\psi^1 \wedge F(\psi^2 \wedge \dots (F\psi^k) \dots)))$$

166 where $\overline{\psi}$ and ψ^i are communication predicates for a phase. Predicate $\overline{\psi}$ is a *global predicate*,
 167 and ψ^1, \dots, ψ^k are *sporadic predicates*. So the global predicate specifies constraints on every
 168 phase of execution, while sporadic predicates specify a *sequence* of special phases that should
 169 happen: first ψ_1 , followed later by ψ_2 , etc. We have two types of atomic communication
 170 predicates: $\varphi_=_$ says that every process receives the same multiset; φ_{thr} says that every
 171 process receives a multiset of size at least $thr \cdot n$ where n is the number of processes. In
 172 Algorithm 1 the global predicate is trivial, and we require two special phases. In the first of
 173 them, in its first round every process should receive exactly the same H multiset, and the
 174 multiset should contain values from at least thr_1 fraction of all processes.

175 Semantics

176 The values of variables come from a fixed linearly ordered set D . Additionally, we take a
 177 special value $? \notin D$ standing for undefined. We write $D_?$ for $D \cup \{?\}$.

178 We describe the semantics of an algorithm for n processes. A *state of an algorithm* is a
 179 pair of n -tuples of values; denoted (f, d) . Intuitively, f specifies the value of the *inp* variable
 180 for each process, and d specifies the value of the *dec* variable. The value of *inp* can never be
 181 $?$, while initially the value of *dec* is $?$ for every process. We denote by $mset(f)$ the multiset
 182 of values appearing in the tuple f . Only values of *inp* and *dec* survive between phases. All
 183 the other variables are reset to $?$ at the beginning of each phase.

184 There are two kinds of transitions:

$$185 \quad (f, d) \xrightarrow{\psi} (f', d') \quad \text{a phase transition}$$

$$186 \quad f \xrightarrow{\varphi}_i f' \quad \text{a transition for round } i$$

188 Phase transitions will be defined as a composition of round transitions. In a transition for
 189 round i , tuple f describes the values of x_{i-1} , and f' the values of x_i . Phase transition is
 190 labeled with a phase communication predicate, while a round transition has a round number
 191 and a conjunction of atomic predicates as labels.

192 Before defining these transitions we need to describe the semantics of communication
 193 predicates. At every round processes send values of their variable to a communication medium,
 194 and then receive a multiset of values from the medium (cf. Figure 1). Communication medium
 195 is not assumed to be perfect, it can send a different multiset of values to every process,
 196 provided it is a sub-multiset of received values. An atomic communication predicate puts
 197 constraints on multisets that every process receives. In other words, such a predicate specifies
 198 constraints on a tuple of multisets $\vec{H} = (H_1, \dots, H_n)$. Predicate $\varphi_=_$ requires that all the
 199 multisets are the same. Predicate φ_{thr} requires that every multiset is bigger than $thr \cdot n$ for

200 some number $0 \leq thr < 1$. Predicate *true* does not put any restrictions. We write $\vec{H} \models \varphi$
 201 when the tuple of multisets \vec{H} satisfies the conjunction of atomic predicates φ .

202 Once a process p receives a multiset H_p , it uses it to do an update of one of its variables.
 203 For this it finds the first condition that H_p satisfies and performs the operation from the
 204 corresponding assignment.

205 Recall that a *condition* is a conjunction of atomic conditions: **uni**, **mult**, $|H| > thr \cdot |\Pi|$.
 206 A multiset H satisfies **uni** when it contains just one value; it satisfies **mult** if it contains more
 207 than one value. A multiset H satisfies $|H| > thr \cdot |\Pi|$ when the size of H is bigger than $thr \cdot n$,
 208 where n is the number of processes. Observe that only predicates of the last type take into
 209 account possible repetitions of the same value.

210 We can now define the *update value* $\text{update}_i(H)$, describing to which value the process
 211 sets its variable in round i upon receiving the multiset H . For this the process finds the first
 212 conditional statement in the sequence of instructions for round i whose condition is satisfied
 213 by $H - \{?\}$ and looks at the operation in the statement:

- 214 ■ if it is $x := \min(H)$ then $\text{update}_i(H)$ is the minimal value in $H - \{?\}$;
- 215 ■ if it is $x := \text{smor}(H)$ then $\text{update}_i(H)$ is the smallest most frequent value in $H - \{?\}$;
- 216 ■ if no condition is satisfied then $\text{update}_i(H) = ?$.

217 A transition $f \xrightarrow{\varphi}_i f'$ is possible when there exists a tuple of multisets $(H_1, \dots, H_n) \models \varphi$
 218 such that for all $p = 1, \dots, n$: $H_p \subseteq \text{mset}(f)$, and $f'(p) = \text{update}_i(H_p)$. Observe that ? value
 219 in H_p is ignored by the **update** function, but not by the communication predicate.

220 A transition $(f, d) \xrightarrow{\psi} (f', d')$, for $\psi = (\varphi_1, \dots, \varphi_n)$, is possible when there is a sequence:

$$221 \quad f_0 \xrightarrow{\varphi_1}_1 f_1 \xrightarrow{\varphi_2}_2 \dots \xrightarrow{\varphi_{r-1}}_{r-1} f_{r-1} \xrightarrow{\varphi_r}_r f_r \quad \text{where}$$

- 222 ■ $f_0 = f$;
- 223 ■ $f'(p) = f_{\text{ir}}(p)$ if $f_{\text{ir}}(p) \neq ?$, and $f'(p) = f_{\text{ir}}(p)$ otherwise;
- 224 ■ $d'(p) = d(p)$ if $d(p) \neq ?$, and $d'(p) = f_r(p)$ otherwise.

225 This means that if in round **ir**, the value of $f_{\text{ir}}(p)$ was ?, then the process p retains its value
 226 of the *inp* onto the next phase; otherwise the process p updates its value of *inp* to $f_{\text{ir}}(p)$.
 227 The value of *dec* cannot be updated, it can only be set if it has not been set before. For
 228 setting the value of *dec*, the value from the last round is used.

229 An *execution* is a sequence of phase transitions. An *execution of an algorithm respecting*
 230 *a communication predicate* $(G\bar{\psi}) \wedge (\mathbf{F}(\psi^1 \wedge \mathbf{F}(\psi^2 \wedge \dots (\mathbf{F}\psi^k) \dots)))$ is an infinite sequence:

$$231 \quad (f_0, d_0) \xrightarrow{\bar{\psi}^*} (f_1, d_1) \xrightarrow{\psi \wedge \psi^1} (f'_1, d'_1) \dots \xrightarrow{\bar{\psi}^*} (f_k, d_k) \xrightarrow{\psi \wedge \psi^k} (f'_k, d'_k) \xrightarrow{\bar{\psi}^\omega} \dots$$

232 where $\xrightarrow{\bar{\psi}^*}$ stands for a finite sequence of $\xrightarrow{\bar{\psi}}$ transitions, and $\xrightarrow{\bar{\psi}^\omega}$ for an infinite sequence.
 233 For every execution there is some fixed n determining the number of processes, f_0 is any
 234 n -tuple of values without ?, and d_0 is the n -tuple of ? values. Observe that the size of the
 235 first tuple determines the size of every other tuple. By definition of transitions, there is
 236 always a transition from every configuration, so an execution cannot block. Thus we can
 237 think of every execution as being infinite.

238 ► **Definition 1** (Consensus problem). *An algorithm has agreement property if for every*
 239 *number of processes n , and for every state (f, d) reachable by an execution of the algorithm,*
 240 *for all processes p_1 and p_2 , either $d(p_1) = d(p_2)$ or one of the two values is ?. An algorithm*
 241 *has termination property if for every n , and for every execution there is a state (f, d) on this*
 242 *execution with $d(p) \neq ?$ for all $p = 1, \dots, n$. An algorithm solves consensus if it has agreement*
 243 *and termination properties.*

244 ▶ Remark 2. Normally, the consensus problem also requires irrevocability and integrity
 245 properties, but these are always guaranteed by the semantics: once set, a process cannot
 246 change its *dec* value, and a variable can be set only to one of the values that has been
 247 received.

248 ▶ Remark 3. The original definition of the Heard-Of model is open ended: it does not limit
 249 possible forms of a communication predicate, conditions, or operations. Clearly, for the kind
 250 of result we present here, we need to fix them.

251 ▶ Remark 4. In the original definition processes are allowed to have identifiers. We do not
 252 need them for the set of operations we consider. Later we will add coordinators without
 253 referring to identifiers. This is a relatively standard way of avoiding identifiers while having
 254 reasonable expressivity.

255 3 A characterization for the core language

256 We present a characterization of all the algorithms in our language that solve consensus. In
 257 later sections we will extend it to include timestamps and coordinators. As it will turn out,
 258 for our analysis we will need to consider only two values a, b with a fixed order between them:
 259 we take a smaller than b . This order influences the semantics of instructions: the result of
 260 `min` is a on a multiset containing at least one a ; the result of `smor` is a on a multiset with
 261 the same number of a 's and b 's. Because of this asymmetry we mostly focus on the number
 262 of b 's in a tuple. In our analysis we will consider tuples of the form $bias(\theta)$ for $\theta < 1$, i.e., a
 263 tuple where we have n processes (for some large enough n), out of which $\theta \cdot n$ of them have
 264 their value set to b ; and the remaining ones to a . The tuple containing only b 's (resp. only
 265 a 's) is called *solo* (resp. *solo^a*).

266 We show that there is essentially one way to solve consensus. The text of the algorithm
 267 together with the form of the global predicate determines a threshold \overline{thr} . We prove that in
 268 the language we consider here, there should be a *unifier phase* which guarantees that the
 269 tuple of *inp* values after the phase belongs to one of the following four types: *solo*, *solo^a*,
 270 *bias*(θ), or *bias*($1 - \theta$) where $\theta \geq \overline{thr}$. Intuitively, this means that there is a dominant value in
 271 the tuple. This phase should be followed by a *decider phase* which guarantees that if the tuple
 272 of *inp* values is of one of the above mentioned types, then all the processes decide. While
 273 this ensures termination, agreement is ensured by proving that some structural properties on
 274 the algorithm should always hold.

275 Before stating the characterization, we will make some observations that allow us to
 276 simplify the structure of an algorithm, and in consequence simplify the statements.

277 It is easy to see that in our language we can assume that the list of conditional instructions
 278 in each round can have at most one **uni** conditional followed by a sequence of **mult** conditionals
 279 with non-increasing thresholds:

```

280   if uni(H) ∧ |H| > thrui · |Π| then x := opui(H)
281   if mult(H) ∧ |H| > thrmi,1 · |Π| then x := opmi(H)
282   ⋮
283   if mult(H) ∧ |H| > thrmi,k · |Π| then x := opmi(H)
284
```

285 We use superscript i to denote the round number: so thr_u^1 is a threshold associated to **uni**
 286 instruction in the first round, etc. If round i does not have a **uni** instruction, then thr_u^1 will
 287 be -1 . For the sake of brevity, $thr_m^{i,k}$ will always denote the minimal threshold appearing in
 288 any of the **mult** instructions in round i and -1 if no **mult** instructions exist in round i .

289 We fix a *communication predicate*:

$$290 \quad (\overline{G\psi}) \wedge (F(\psi^1 \wedge F(\psi^2 \wedge \dots (F\psi^k) \dots))) \quad (1)$$

291 Without loss of generality we can assume that every sporadic predicate implies the global
 292 predicate; in consequence, $\overline{\psi} \wedge \psi^i$ is equivalent to ψ^i . Recall that each of $\overline{\psi}, \psi^1, \dots, \psi^k$ is
 293 an r -tuple of conjunctions of atomic predicates. We write $\psi|_i$ for the i -th element of the
 294 tuple and so ψ is $(\psi|_1, \dots, \psi|_r)$. By $thr_i(\psi)$ we denote the threshold constant appearing in
 295 the predicate $\psi|_i$, i.e., if $\psi|_i$ has φ_{thr} as a conjunct, then $thr_i(\psi) = thr$, if it has no such
 296 conjunct then $thr_i(\psi) = -1$. We call $\psi|_i$ an *equalizer* if it has $\varphi_=-$ as a conjunct. In this case
 297 we also say that ψ has an equalizer.

298 Recall (cf. page 2) that a transition $f \xrightarrow{\psi}_i f'$ for a round i under a phase predicate ψ is
 299 possible when there is a tuple of multisets $(H_1, \dots, H_n) \models \psi|_i$ such that for all $p = 1, \dots, n$:
 300 $H_p \in mset(f)$ and $f'(p) = \text{update}_i(H_p)$.

301 ► **Definition 5.** A round i is preserving w.r.t. ψ iff one of the three conditions hold: (i) it
 302 does not have an **uni** instruction, (ii) it does not have a **mult** instruction, or (iii) $thr_i(\psi) <$
 303 $\max(thr_u^i, thr_m^{i,k})$. Otherwise the round is non-preserving. The round is solo safe w.r.t. ψ if
 304 $0 \leq thr_u^i \leq thr_i(\psi)$.

305 If i is a preserving round, then there exists a tuple f such that $? \notin mset(f)$ and such that
 306 a transition $f \xrightarrow{\psi}_i f'$ is possible for f' a tuple consisting solely of $?$. The consequence of
 307 such a transition is that *inp* is not updated in the phase, i.e., old values of *inp* are *preserved*.
 308 On the other extreme, if all transitions in the phase are non-preserving then all *inp* values
 309 are necessarily updated by the phase. Finally, a solo safe round cannot alter the *solo* state,
 310 i.e., $solo \xrightarrow{\psi}_i solo$ is the only transition possible from *solo*.

311 ► **Remark 6.** Suppose rounds $1, \dots, i-1$ are non-preserving under $\overline{\psi}$, the global predicate.
 312 In this situation, since $? \notin mset(f)$, if $f \xrightarrow{\overline{\psi}|_1}_1 f_1 \xrightarrow{\overline{\psi}|_2}_2 \dots \xrightarrow{\overline{\psi}|_{i-1}}_{i-1} f_{i-1}$ then $? \notin mset(f_{i-1})$.
 313 Hence, no heard-of multi-set H constructed from f_{i-1} can have $?$ value. Notice that every
 314 process is bound to receive a heard-of set of size at least $thr_i(\overline{\psi})$ in round i . For a sake of
 315 example, suppose $thr_i(\overline{\psi}) > thr_m^{i,2}$. The semantics then guarantees that every heard-of set
 316 sent during the i^{th} round either satisfies the **uni** instruction, or one of the first two **mult**
 317 instructions, or no instruction at all. Hence, in such a case all the **mult** instructions except
 318 the first two can be removed from the description of round i as they will be never executed.
 319 This implies that we can adopt the following assumption.

320 ► **Assumption 1.** For every round i , if rounds $1, \dots, i-1$ are non-preserving under $\overline{\psi}$ then

$$321 \quad \begin{cases} thr_u^i \geq thr_i(\overline{\psi}) & \text{if round } i \text{ has } \mathbf{uni} \text{ instruction} \\ thr_m^{i,k} \geq thr_i(\overline{\psi}) & \text{if round } i \text{ has } \mathbf{mult} \text{ instruction} \end{cases} \quad (2)$$

322 We put some restrictions on the form of algorithms we consider in our characterization.
 323 They greatly simplify the statements, and as we argue, are removing cases that are not that
 324 interesting anyway.

325 ► **Proviso 1.** We adopt the following additional syntactic restrictions:

- 326 ■ We require that the global predicate does not have an equalizer.
- 327 ■ We assume that there is no **mult** instruction in the round $\mathbf{ir} + 1$.

328 Concerning the first of the above requirements, if the global predicate has an equalizer
 329 then it is quite easy to construct an algorithm for consensus because equalizer guarantees

330 that in a given round all the processes receive the same value. The characterization below
 331 can be extended to this case but would require to mention it separately in all the statements.
 332 Concerning the second requirement, We can show that if such a `mult` instruction exists then
 333 either the algorithm violates consensus, or the instruction will never be fired in any execution
 334 of the algorithm and so it can be removed without making an algorithm incorrect.

335 In order to state our characterization we need to give formal definitions of concepts we
 336 have discussed at the beginning of the section.

337 ► **Definition 7.** *The border threshold is $\overline{thr} = \max(1 - thr_u^1, 1 - thr_m^{1,k}/2)$.*

338 ► **Definition 8.** *A predicate ψ is a*

- 339 ■ Decider, if all rounds are solo safe w.r.t. ψ
- 340 ■ Unifier, if the three conditions hold:
 - 341 ■ $thr_1(\psi) \geq thr_m^{1,k}$ and either $thr_1(\psi) \geq thr_u^1$ or $thr_1(\psi) \geq \overline{thr}$,
 - 342 ■ there exists i such that $1 \leq i \leq \mathbf{ir}$ and $\psi|_i$ is an equalizer,
 - 343 ■ rounds $2, \dots, i$ are non-preserving w.r.t. ψ and rounds $i+1, \dots, \mathbf{ir}$ are solo-safe w.r.t. ψ

344 Finally, we list some syntactic properties of algorithms that, as we will see later, imply
 345 the agreement property.

346 ► **Definition 9.** *An algorithm is syntactically safe when:*

- 347 1. First round has a `mult` instruction.
- 348 2. Every round has a `uni` instruction.
- 349 3. In the first round the operation in every `mult` instruction is `smor`.
- 350 4. $thr_m^{1,k}/2 \geq 1 - thr_u^{\mathbf{ir}+1}$, and $thr_u^1 \geq 1 - thr_u^{\mathbf{ir}+1}$.

351 Recall that ψ^1, \dots, ψ^k are the set of sporadic predicates from the communication predicate.
 352 Without loss of generality we can assume that there is at least one sporadic predicate; at
 353 a degenerate case it is always possible to take a sporadic predicate that is the same as the
 354 global predicate. With these definitions we can state our characterization:

355 ► **Theorem 10.** *Consider algorithms in the core language satisfying syntactic constraints
 356 from Assumption 1 and Proviso 1. An algorithm solves consensus iff it is syntactically safe
 357 according to Definition 9, and it satisfies the condition:*

358 **T** *There is $i \leq j$ such that ψ^i is a unifier and ψ^j is a decider.*

359 A two value principle is a corollary from the proof of the above theorem: an algorithm
 360 solves consensus iff it solves consensus for two values. Indeed, it turns out that it is enough
 361 to work with three values a, b , and $?$ standing for undefined. The proof considers separately
 362 safety and liveness aspects of the consensus problem. Notice that the properties from
 363 Definition 9 intervene also in the proof of termination.

364 ► **Lemma 11.** *An algorithm violating structural properties from Definition 9 cannot solve
 365 consensus. An algorithm with the structural properties has the agreement property.*

366 ► **Lemma 12.** *An algorithm with the structural properties from Definition 9 has the
 367 termination property iff it satisfies condition T from Theorem 10.*

368 **4** A characterization for algorithms with timestamps

369 We extend our characterization to algorithms with timestamps. Now, variable *inp* stores not
 370 only the value but also a timestamp, that is the number of the last phase at which *inp* was
 371 updated. These timestamps are used in the first round, as a process considers only values
 372 with the most recent timestamp. The syntax is the same as before except that we introduce
 373 a new operation, called *maxts*, that must be used in the first round and nowhere else. So
 374 the form of the first round becomes:

```

375   send (inp, ts)
      | if cond1l(H) then x1 := maxts(H);
      | :
      | if cond1l(H) then x1 := maxts(H);
  
```

376 The semantics of transitions for rounds and phases needs to take into account timestamps.
 377 The semantics changes only for the first round; its form becomes $(f, t) \xrightarrow{\varphi} f'$, where t is a
 378 vector of timestamps (n -tuple of natural numbers). Timestamps are ignored by communication
 379 predicates and conditions, but are used in the update operation. The operation *maxts*(H)
 380 returns the smallest among values with the most recent timestamp in H.

381 The form of a phase transition changes to $(f, t, d) \xrightarrow{\psi} (f', t', d')$. Value $t(p)$ is the
 382 timestamp of the last update of *inp* of process p (whose value is $f(p)$). We do not need
 383 to keep timestamps for d since the value of *dec* can be set only once. Phase transitions
 384 are defined as before, taking into account the above mentioned change for the first round
 385 transition, and the fact that in the round *ir* when *inp* is updated then so is its timestamp.
 386 Some examples of algorithms with timestamps are presented in Section 7.

387 As in the case of the core language, without loss of generality we can assume conditions
 388 from Assumption 1. Concerning Proviso 1, we assume almost the same conditions, but now
 389 the second one refers to the round *ir* and not to the round *ir* + 1, and is a bit stronger.

390 ► **Proviso 2.** *We adopt the following syntactic restrictions:*

- 391 ■ We require that the global predicate does not have an equalizer.
- 392 ■ We assume that there is no *mult* instruction in the round *ir*, and that $\text{thr}_u^{\text{ir}} \geq 1/2$.

393 The justification for the first restriction is as before. Concerning the second restriction,
 394 we can prove that if these two assumptions do not hold then either the algorithm violates
 395 consensus, or we can remove the *mult* instruction and increase thr_u^{ir} without making an
 396 algorithm incorrect.

397 Our characterization resembles the one for the core language. The structural conditions
 398 get slightly modified: the condition on constants is weakened, and there is no need to talk
 399 about *smor* operations in the first round.

400 ► **Definition 13.** *An algorithm is syntactically t-safe when:*

- 401 1. Every round has a *uni* instruction.
- 402 2. First round has a *mult* instruction.
- 403 3. $\text{thr}_m^{1,k} \geq 1 - \text{thr}_u^{\text{ir}+1}$ and $\text{thr}_u^1 \geq 1 - \text{thr}_u^{\text{ir}+1}$.

404 We consider the same shape of a communication predicate as in the case of the core
 405 language (1). A characterization for the case with timestamps uses a stronger version of a
 406 unifier that we define now. The intuition is that we do not have $\overline{\text{thr}}$ constant because of
 407 *maxts* operations in the first round. In other words, the conditions are the same as before
 408 but when taking $\overline{\text{thr}} > 1$.

409 ► **Definition 14.** A predicate ψ is a strong unifier ψ if it is a unifier in a sense of Definition 8
410 and $\text{thr}_u^1 \leq \text{thr}_1(\psi)$.

411 Modulo the above two changes, the characterization stays the same.

412 ► **Theorem 15.** Consider algorithms in the language with timestamps satisfying syntactic
413 constraints from Assumption 1 and Proviso 2. An algorithm satisfies consensus iff it is
414 syntactically t -safe according to Definition 13, and it satisfies:

415 **sT** There are $i \leq j$ such that ψ^i is a strong unifier and ψ^j is a decider.

416 **5 A characterization for algorithms with coordinators**

417 We consider algorithms equipped with coordinators. The novelty is that we can now have
418 rounds where there is a unique process that receives values from other processes, as well as
419 rounds where there is a unique process that sends values to other processes. For this we
420 extend the syntax by introducing a round type that can be: **every**, **lr** (leader receive), or
421 **ls** (leader-send):

- 422 ■ A round of type **every** behaves as before.
- 423 ■ In a round of type **lr** only one arbitrarily selected process receives values.
- 424 ■ In a round of type **ls**, the process selected in the immediately preceding **lr** round sends
425 its value to all other processes.

426 If an **ls** round is not preceded by an **lr** round then an arbitrarily chosen process sends its
427 value. We assume that every **lr** round is immediately followed by an **ls** round, because
428 otherwise the **lr** round would be useless. We also assume that *inp* and *dec* are not updated
429 during **lr** rounds, as only one process is active in these rounds.

430 For **ls** rounds we introduce a new communication predicate. The predicate φ_{1s} says that
431 the leader successfully sends its message to everybody; it makes sense only for **ls** rounds.

432 These extensions of the syntax are reflected in the semantics. For convenience we introduce
433 two new names for tuples: one^b is a tuple where all the entries are ? except for one entry
434 which is b ; similarly for one^a . Abusing the notation we also write $one^?$ for *solo*[?], namely the
435 tuple consisting only of ? values.

436 Let us define the semantics of **lr** and **ls** rounds. If i -th round is of type **lr**, we have
437 a transition $f \xrightarrow{\psi}_i one^d$ for every $d \in \text{fire}_i(f, \psi)$. In particular, if $? \in \text{fire}_i(f, \psi)$ then
438 $f \xrightarrow{\varphi}_i solo^?$ is possible.

439 Suppose i -th round is of type **ls**. If $\psi|_i$ contains φ_{1s} as a conjunct then

$$440 \quad one^d \xrightarrow{\psi}_i solo^d \quad \text{if round } (i-1) \text{ is of type } \mathbf{lr}$$

$$441 \quad f \xrightarrow{\psi}_i solo^d \quad \text{for } d \in \text{set}(f) \quad \text{otherwise}$$

443 When $\psi|_i$ does not contain φ_{1s} then independently of the type of the round $(i-1)$ we have
444 $f \xrightarrow{\psi}_i f'$ for every $d \in \text{set}(f)$ and f' such that $\text{set}(f') \subseteq \{d, ?\}$.

445 We consider the same shape of a communication predicate as in the case of the core
446 language (1).

447 The semantics allows us to adopt some more simplifying assumptions about the syntax
448 of the algorithm, and the form of the communication predicate.

449 ► **Assumption 2.** We assume that **ls** rounds do not have a **mult** instruction. Indeed, from
450 the above semantics it follows that **mult** instruction is never used in a round of type **ls**. It
451 also does not make much sense to use φ_{1s} in rounds other than of type **ls**. So to shorten some

9:12 Consensus in the Heard-Of model

452 definitions we require that φ_{1s} can appear only in communication predicates for $1s$ -rounds.
 453 For similar reasons we require that $\varphi_{=}$ predicate is not used in $1s$ -rounds. As we have
 454 observed in the first paragraph, we can assume that neither round ir nor the last round are
 455 of type $1r$.

456 The notions of preserving and solo-safe rounds get adapted to the new syntax.

457 ► **Definition 16.** A round of type $1s$ is c -solo-safe w.r.t. ψ if ψ_i has φ_{1s} as a conjunct, it is
 458 c -preserving otherwise. A round of type other than $1s$ is c -preserving or c -solo-safe w.r.t ψ
 459 if it is so in the sense of Definition 5.

460 ► **Definition 17.** A c -equalizer is a conjunction containing a term of the form $\varphi_{=}$ or φ_{1s} .

461 ► **Proviso 3.** We assume the same conditions as in Proviso 8, but using the concepts of
 462 c -equalizers instead of equalizers.

463 To justify the proviso we prove that **mult** instruction in round $ir + 1$ cannot be useful.
 464 Assumption 1 is also updated to using the notion of c -preserving instead of preserving. We
 465 restate it for convenience.

466 ► **Assumption 3.** For every round i , if rounds $1, \dots, i - 1$ are non- c -preserving under $\bar{\psi}$
 467 then

$$468 \begin{cases} thr_u^i \geq thr_i(\bar{\psi}) & \text{if round } i \text{ has } \mathbf{uni} \text{ instruction} \\ thr_m^{i,k} \geq thr_i(\bar{\psi}) & \text{if round } i \text{ has } \mathbf{mult} \text{ instruction} \end{cases} \quad (3)$$

469 Finally, the above modifications imply modifications of terms from Definition 8.

470 ► **Definition 18.** A predicate ψ is called a

471 ■ c -decider, if all rounds are c -solo safe w.r.t. ψ .

472 ■ c -unifier, if

473 ■ $thr_1(\psi) \geq thr_m^{1,k}$ and either $thr_1(\psi) \geq thr_u^1$ or $thr_1(\psi) \geq \bar{thr}$,

474 ■ there exists i such that $1 \leq i \leq ir$ and $\psi|_i$ is an c -equalizer,

475 ■ rounds $2, \dots, i$ are non- c -preserving w.r.t. ψ and rounds $i + 1, \dots, ir$ are c -solo-safe
 476 w.r.t. ψ .

477 With these modifications, we get an analog of Theorem 10 for the case with coordinators
 478 subject to the modified provisos as explained above.

479 ► **Theorem 19.** Consider algorithms in the language with timestamps satisfying syntactic
 480 constraints from Assumptions 2, 3 and Proviso 3. An algorithm satisfies consensus iff the
 481 first round and the $(ir + 1)^{th}$ round are not of type $1s$, it is syntactically safe according to
 482 Definition 9, and it satisfies the condition:

483 **cT** There are $i \leq j$ such that ψ^i is a c -unifier and ψ^j is a c -decider.

484 6 A characterization for algorithms with coordinators and timestamps

485 Finally, we consider an extension of the core language with both coordinators and timestamps.
 486 Formally, we extend the coordinator model with timestamps in the same way we have extended
 487 the core model. So now *inp* variables store pairs (value, timestamp), and all the instructions
 488 in the first round are *maxts* (cf. page 10).

489 ► **Proviso 4.** We assume the same proviso as for timestamps; namely, Proviso 2, but using
 490 the notion of c -equalizer.

491 As in the previous cases we justify our proviso by showing that the algorithm violating
492 the second condition would not be correct or the condition could be removed.

493 The characterization is a mix of conditions from timestamps and coordinator cases.

494 ► **Definition 20.** A predicate ψ is a strong c -unifier if it is a c -unifier (cf. Definition 14)
495 and $thr_u^1 \leq thr_1(\psi)$.

496 ► **Theorem 21.** Consider algorithms in the language with timestamps satisfying syntactic
497 constraints from Assumptions 2, 3 and Proviso 4. An algorithm satisfies consensus iff the
498 first round and the $(ir + 1)^{th}$ round are not of type $1s$, it has the structural properties from
499 Definition 13, and it satisfies:

500 **scT** There are $i \leq j$ such that ψ^i is a strong c -unifier and ψ^j is a c -decider.

501 7 Examples

502 We apply the characterizations from the previous sections to some consensus algorithms
503 studied in the literature, and their variants.

504 First, we can revisit the parametrized Algorithm 1 from page 4. This is an algorithm
505 in the core language, and it depends on two thresholds. Theorem 10 implies that it solves
506 consensus iff $thr_1/2 \geq 1 - thr_2$. In case of $thr_1 = thr_2 = 2/3$ we obtain the well known
507 OneThird algorithm. But, for example, $thr_1 = 1/2$ and $thr_2 = 3/4$ are also possible solutions
508 for this inequality. So Algorithm 1 solves consensus for these values of thresholds.

509 Because of the conditions on constants, $thr_m^{1,k}/2 \geq 1 - thr_u^{ir+1}$ coming from Definition 9,
510 it is not possible to have an algorithm in the core language where all constants are at most
511 $1/2$. This answers a question from [11] for the language we consider here.

512 The above condition on constants is weakened to $thr_m^{1,k} \geq 1 - thr_u^{ir+1}$ when we have
513 timestamps. In this case indeed it is possible to use only $1/2$ thresholds [27].

514 When we have both timestamps and coordinators, we get variants of Paxos algorithm.

```

send (inp, ts) lr
  | if uni(H) ∧ |H| > 1/2 · |Π| then x1 := maxts(H);
  | if mult(H) ∧ |H| > 1/2 · |Π| then x1 := maxts(H);
send x1 ls
  | if uni(H) then x2 := inp := smor(H);
515 send x2 lr
  | if uni(H) ∧ |H| > 1/2 · |Π| then x3 := smor(H);
send x3 ls
  | if uni(H) then dec := smor(H);
Communication predicate: F(ψ1) where ψ1 := (φ1/2, φ1s, φ1/2, φ1s)

```

► Algorithm 2 Paxos algorithm

516 The algorithm is correct by Theorem 21. One can observe that without modifying the
517 code there is not much room for improvement in this algorithm. A decider phase is needed to
518 solve consensus, and ψ_1 is a minimal requirement for a decider phase. A possible modification
519 is to change the thresholds in the first round to, say, $1/3$ and in the third round to $2/3$ (both

520 in the algorithm and in the communication predicate).

```

520   send (inp, ts) lr
      | if uni(H) ∧ |H| > 1/2 · |Π| then x1 := maxts(H);
      | if mult(H) ∧ |H| > 1/2 · |Π| then x1 := maxts(H);
521   send x1 ls
      | if uni(H) then x2 := inp := smor(H);
   send x2 every
      | if uni(H) ∧ |H| > 1/2 · |Π| then dec := smor(H);
   Communication predicate: F(ψ1) where ψ1 := (φ1/2, φ1s, φ1/2)

```

■ **Algorithm 3** Three round Paxos algorithm

522 The three round Paxos presented above is also correct by Theorem 21. Once again it is
523 possible to change constants in the first round to 1/3 and in the last round to 2/3 (both in
524 the algorithm and in the communication predicate).

525 One can also wonder about algorithms with coordinators but without timestamps. Here
526 is a possibility that resembles three round Paxos:

```

   send (inp) lr
      | if uni(H) ∧ |H| > 2/3 · |Π| then x1 := smor(H);
      | if mult(H) ∧ |H| > 2/3 · |Π| then x1 := smor(H);
   send x1 ls
527   | if uni(H) then x2 := inp := smor(H);
   send x2 every
      | if uni(H) ∧ |H| > 2/3 · |Π| then dec := smor(H);
   Communication predicate: F(ψ) where ψ := (φ2/3, φ1s, φ2/3)

```

■ **Algorithm 4** Three round coordinator algorithm

528 The algorithm solves consensus by Theorem 19. The constants are bigger than in Paxos
529 because we do not have timestamps: the constraints on constants come from Definition 9,
530 and not from Definition 13. The advantage is that we do not need time-stamps, while keeping
531 the same structure as for three-round Paxos.

532 It is possible to introduce more parameters in these algorithms to analyze for which
533 choices of parameters they solve consensus.

534 8 Conclusions

535 We have characterized all algorithms solving consensus in a fragment of the Heard-Of model.
536 We have aimed at a fragment that can express most important algorithms while trying to
537 avoid ad hoc restrictions (c.f. proviso on page 8). The fragment covers algorithms considered
538 in the context of verification [28, 10] with a notable exception of algorithms sending more than
539 one variable. In this work we have considered only single phase algorithms while originally
540 the model permits also to have initial phases. We believe that it is possible to extend
541 the characterization to incorporate the initial phases, but this would further complicate
542 the results and there are no well-know algorithms that use such phases. More severe and
543 technically important restriction is that we allow to use only one variable at a time. In
544 particular, it is not possible to send pairs of variables.

545 One curious direction of further research would be to list all “best” consensus algorithms
546 under some external constraints; for example the constraints can come from some properties
547 of an execution platform external to the Heard-Of model. This problem assumes that there
548 is some way to compare two algorithms. One guiding principle for such a measure could

549 be efficient use of knowledge [30, 29]: at every step the algorithm does maximum it can do,
 550 given its knowledge of the state of the system.

551 This research is on the borderline between distributed computing and verification. From
 552 a distributed computing side it considers quite a simple model, but gives a characterization
 553 result. From a verification side, the systems are complicated because the number of processes
 554 is unbounded, there are timestamps, and interactions are based on a fraction of processes
 555 having a particular value. We do not advance on verification methods for such a setting.
 556 Instead, we observe that in the context considered here verification may be avoided. We
 557 believe that a similar phenomenon can appear also for other problems than consensus. It is
 558 also an intriguing question to explore how much we can enrich the current model and still
 559 get a characterization. We conjecture that a characterization is possible for an extension
 560 with randomness covering at least the Ben-Or algorithm. Of course, formalization of proofs,
 561 either in Coq or Isabelle, for such extensions would be very helpful.

562 ——— References ———

- 563 1 Marcos K. Aguilera, Carole Delporte-Gallet, Hugues Fauconnier, and Sam Toueg. Partial
 564 synchrony based on set timeliness. *Distributed Computing*, 25(3):249–260, 2012. doi:10.1007/
 565 s00446-012-0158-8.
- 566 2 Benjamin Aminof, Sasha Rubin, Iina Stoilkovska, Josef Widder, and Florian Zuleger. Param-
 567 eterized model checking of synchronous distributed algorithms by abstraction. In Isil Dillig
 568 and Jens Palsberg, editors, *Verification, Model Checking, and Abstract Interpretation - 19th*
 569 *International Conference, VMCAI 2018*, volume 10747 of *Lecture Notes in Computer Science*,
 570 pages 1–24. Springer, 2018. doi:10.1007/978-3-319-73721-8_1.
- 571 3 A.R. Balasubramanian and Igor Walukiewicz. Characterizing consensus in the heard-of model.
 572 <http://arxiv.org/abs/2004.09621>.
- 573 4 Nathalie Bertrand, Igor Konnov, Marijana Lazic, and Josef Widder. Verification of randomized
 574 consensus algorithms under round-rigid adversaries. In Wan Fokkink and Rob van Glabbeek,
 575 editors, *30th International Conference on Concurrency Theory*, volume 140 of *LIPICs*, pages
 576 33:1–33:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.
 577 CONCUR.2019.33.
- 578 5 Martin Biely, Josef Widder, Bernadette Charron-Bost, Antoine Gaillard, Martin Hutle, and
 579 André Schiper. Tolerating corrupted communication. In Indranil Gupta and Roger Wattenhofer,
 580 editors, *Proceedings of the Twenty-Sixth Annual ACM Symposium on Principles of Distributed*
 581 *Computing, PODC 2007*, pages 244–253. ACM, 2007. doi:10.1145/1281100.1281136.
- 582 6 Tushar Deepak Chandra, Vassos Hadzilacos, and Sam Toueg. The weakest failure detector for
 583 solving consensus. *J. ACM*, 43(4):685–722, 1996. doi:10.1145/234533.234549.
- 584 7 Tushar Deepak Chandra and Sam Toueg. Unreliable failure detectors for reliable distributed
 585 systems. *J. ACM*, 43(2):225–267, 1996. doi:10.1145/226643.226647.
- 586 8 Mouna Chaouch-Saad, Bernadette Charron-Bost, and Stephan Merz. A reduction theorem for
 587 the verification of round-based distributed algorithms. In Olivier Bournez and Igor Potapov,
 588 editors, *Reachability Problems, 3rd International Workshop, RP 2009*, volume 5797 of *Lecture*
 589 *Notes in Computer Science*, pages 93–106. Springer, 2009. doi:10.1007/978-3-642-04420-5_10.
- 590 9 Bernadette Charron-Bost, Henri Debrat, and Stephan Merz. Formal verification of consensus
 591 algorithms tolerating malicious faults. In Xavier Défago, Franck Petit, and Vincent Villain,
 592 editors, *Stabilization, Safety, and Security of Distributed Systems - 13th International Sympo-*
 593 *sium, SSS 2011*, volume 6976 of *Lecture Notes in Computer Science*, pages 120–134. Springer,
 594 2011. doi:10.1007/978-3-642-24550-3_11.
- 595 10 Bernadette Charron-Bost and Stephan Merz. Formal verification of a consensus algorithm
 596 in the heard-of model. *Int. J. Software and Informatics*, 3(2-3):273–303, 2009. URL: [http:
 597 //www.ijsi.org/ch/reader/view_abstract.aspx?file_no=273&flag=1](http://www.ijsi.org/ch/reader/view_abstract.aspx?file_no=273&flag=1).

- 599 **11** Bernadette Charron-Bost and André Schiper. The heard-of model: computing in distributed
600 systems with benign faults. *Distributed Computing*, 22(1):49–71, 2009.
- 601 **12** Flaviu Cristian and Christof Fetzer. The timed asynchronous distributed system model. *IEEE*
602 *Trans. Parallel Distrib. Syst.*, 10(6):642–657, 1999. doi:10.1109/71.774912.
- 603 **13** Andrei Damian, Cezara Dragoi, Alexandru Militaru, and Josef Widder. Communication-
604 closed asynchronous protocols. In *CAV (2)*, volume 11562 of *Lecture Notes in Computer*
605 *Science*, pages 344–363. Springer, 2019. URL: <https://dblp.org/rec/conf/cav/DamianDMW19>,
606 doi:10.1007/978-3-030-25543-5_20.
- 607 **14** Henri Debrat and Stephan Merz. Verifying fault-tolerant distributed algorithms in the heard-of
608 model. *Archive of Formal Proofs*, 2012, 2012. URL: [https://www.isa-afp.org/entries/
609 Heard_Of.shtml](https://www.isa-afp.org/entries/Heard_Of.shtml).
- 610 **15** Cezara Dragoi, Thomas A. Henzinger, Helmut Veith, Josef Widder, and Damien Zufferey.
611 A logic-based framework for verifying consensus algorithms. In Kenneth L. McMillan and
612 Xavier Rival, editors, *Verification, Model Checking, and Abstract Interpretation VMCAI*
613 *2014*, volume 8318 of *Lecture Notes in Computer Science*, pages 161–181. Springer, 2014.
614 doi:10.1007/978-3-642-54013-4_10.
- 615 **16** Cezara Dragoi, Thomas A. Henzinger, and Damien Zufferey. Psync: a partially synchronous
616 language for fault-tolerant distributed algorithms. In Rastislav Bodík and Rupak Majumdar,
617 editors, *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles*
618 *of Programming Languages, POPL 2016*, pages 400–415. ACM, 2016. doi:10.1145/2837614.
619 2837650.
- 620 **17** Cynthia Dwork, Nancy A. Lynch, and Larry J. Stockmeyer. Consensus in the presence of partial
621 synchrony. *J. ACM*, 35(2):288–323, 1988. URL: <http://doi.acm.org/10.1145/42282.42283>,
622 doi:10.1145/42282.42283.
- 623 **18** Michael J. Fischer, Nancy A. Lynch, and Mike Paterson. Impossibility of distributed consensus
624 with one faulty process. *J. ACM*, 32(2):374–382, 1985. doi:10.1145/3149.214121.
- 625 **19** Felix C. Freiling, Rachid Guerraoui, and Petr Kuznetsov. The failure detector abstraction.
626 *ACM Comput. Surv.*, 43(2):9:1–9:40, 2011. doi:10.1145/1883612.1883616.
- 627 **20** Eli Gafni. Round-by-round fault detectors: Unifying synchrony and asynchrony (extended
628 abstract). In Brian A. Coan and Yehuda Afek, editors, *Proceedings of the Seventeenth Annual*
629 *ACM Symposium on Principles of Distributed Computing, PODC '98*, pages 143–152. ACM,
630 1998. doi:10.1145/277697.277724.
- 631 **21** Rachid Guerraoui and Michel Raynal. The alpha of indulgent consensus. *Comput. J.*, 50(1):53–
632 67, 2007. doi:10.1093/comjnl/bx1046.
- 633 **22** Michel Hurfin, Achour Mostéfaoui, and Michel Raynal. A versatile family of consensus protocols
634 based on chandra-toueg’s unreliable failure detectors. *IEEE Trans. Computers*, 51(4):395–408,
635 2002. doi:10.1109/12.995450.
- 636 **23** Igor Konnov, Helmut Veith, and Josef Widder. SMT and POR beat counter abstraction:
637 Parameterized model checking of threshold-based distributed algorithms. In Daniel Kroening
638 and Corina S. Pasareanu, editors, *Computer Aided Verification - 27th International Conference,*
639 *CAV 2015*, volume 9206 of *Lecture Notes in Computer Science*, pages 85–102. Springer, 2015.
640 doi:10.1007/978-3-319-21690-4_6.
- 641 **24** Igor V. Konnov, Marijana Lazic, Helmut Veith, and Josef Widder. A short counterexample
642 property for safety and liveness verification of fault-tolerant distributed algorithms. In Giuseppe
643 Castagna and Andrew D. Gordon, editors, *Proceedings of the 44th ACM SIGPLAN Symposium*
644 *on Principles of Programming Languages, POPL 2017*, pages 719–734. ACM, 2017. URL:
645 <http://dl.acm.org/citation.cfm?id=3009860>.
- 646 **25** Igor V. Konnov, Helmut Veith, and Josef Widder. On the completeness of bounded model
647 checking for threshold-based distributed algorithms: Reachability. *Inf. Comput.*, 252:95–109,
648 2017. doi:10.1016/j.ic.2016.03.006.
- 649 **26** Jure Kukovec, Igor Konnov, and Josef Widder. Reachability in parameterized systems: All
650 flavors of threshold automata. In Sven Schewe and Lijun Zhang, editors, *29th International*

- 651 *Conference on Concurrency Theory, CONCUR 2018*, volume 118 of *LIPIcs*, pages 19:1–19:17.
652 Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018. doi:10.4230/LIPIcs.CONCUR.2018.
653 19.
- 654 27 Ognjen Maric. *Formal Verification of Fault-Tolerant Systems*. PhD thesis, ETH Zurich, 2017.
- 655 28 Ognjen Maric, Christoph Sprenger, and David A. Basin. Cutoff bounds for consensus algo-
656 rithms. In Rupak Majumdar and Viktor Kuncak, editors, *Computer Aided Verification - 29th*
657 *International Conference, CAV 2017*, volume 10427 of *Lecture Notes in Computer Science*,
658 pages 217–237. Springer, 2017. doi:10.1007/978-3-319-63390-9_12.
- 659 29 Yoram Moses. Knowledge in distributed systems. In *Encyclopedia of Algorithms*, pages
660 1051–1055. Springer, 2016. doi:10.1007/978-1-4939-2864-4_606.
- 661 30 Yoram Moses and Sergio Rajsbaum. A layered analysis of consensus. *SIAM J. Comput.*,
662 31(4):989–1021, 2002. doi:10.1137/S0097539799364006.
- 663 31 Achour Mostéfaoui and Michel Raynal. Solving consensus using chandra-toueg’s unreliable
664 failure detectors: A general quorum-based approach. In Prasad Jayanti, editor, *Distributed*
665 *Computing, 13th International Symposium*, volume 1693 of *Lecture Notes in Computer Science*,
666 pages 49–63. Springer, 1999. doi:10.1007/3-540-48169-9_4.
- 667 32 Michel Raynal and Julien Stainer. Synchrony weakened by message adversaries vs asynchrony
668 restricted by failure detectors. In Panagiota Fatourou and Gadi Taubenfeld, editors, *ACM*
669 *Symposium on Principles of Distributed Computing, PODC ’13*, pages 166–175. ACM, 2013.
670 doi:10.1145/2484239.2484249.
- 671 33 Olivier Rützi, Zarko Milosevic, and André Schiper. Generic construction of consensus algorithms
672 for benign and byzantine faults. In *Proceedings of the 2010 IEEE/IFIP International Conference*
673 *on Dependable Systems and Networks, DSN 2010*, pages 343–352. IEEE Computer Society,
674 2010. doi:10.1109/DSN.2010.5544299.
- 675 34 Yee Jiun Song, Robbert van Renesse, Fred B. Schneider, and Danny Dolev. The building
676 blocks of consensus. In Shrishia Rao, Mainak Chatterjee, Prasad Jayanti, C. Siva Ram Murthy,
677 and Sanjoy Kumar Saha, editors, *Distributed Computing and Networking, 9th International*
678 *Conference, ICDCN 2008*, volume 4904 of *Lecture Notes in Computer Science*, pages 54–72.
679 Springer, 2008. doi:10.1007/978-3-540-77444-0_5.
- 680 35 Ilina Stoilkovska, Igor Konnov, Josef Widder, and Florian Zuleger. Verifying safety of
681 synchronous fault-tolerant algorithms by bounded model checking. In Tomás Vojnar and Lijun
682 Zhang, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 25th*
683 *International Conference, TACAS 2019*, volume 11428 of *Lecture Notes in Computer Science*,
684 pages 357–374. Springer, 2019. doi:10.1007/978-3-030-17465-1_20.
- 685 36 Marcelo Taube, Giuliano Losa, Kenneth L. McMillan, Oded Padon, Mooly Sagiv, Sharon
686 Shoham, James R. Wilcox, and Doug Woos. Modularity for decidability of deductive verification
687 with applications to distributed systems. In Jeffrey S. Foster and Dan Grossman, editors,
688 *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and*
689 *Implementation, PLDI 2018*, pages 662–677. ACM, 2018. doi:10.1145/3192366.3192414.
- 690 37 Tatsuhiro Tsuchiya and André Schiper. Verification of consensus algorithms using satisfiability
691 solving. *Distributed Computing*, 23(5-6):341–358, 2011. doi:10.1007/s00446-010-0123-3.
- 692 38 Robbert van Renesse, Nicolas Schiper, and Fred B. Schneider. Vive la différence: Paxos vs.
693 viewstamped replication vs. zab. *IEEE Trans. Dependable Sec. Comput.*, 12(4):472–484, 2015.
694 doi:10.1109/TDSC.2014.2355848.
- 695 39 Klaus von Gleissenthall, Nikolaj Bjørner, and Andrey Rybalchenko. Cardinalities and universal
696 quantifiers for verifying parameterized systems. In Chandra Krintz and Emery Berger, editors,
697 *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and*
698 *Implementation, PLDI 2016*, pages 599–613. ACM, 2016. doi:10.1145/2908080.2908129.
- 699 40 Josef Widder and Ulrich Schmid. The theta-model: achieving synchrony without clocks.
700 *Distributed Computing*, 22(1):29–47, 2009. doi:10.1007/s00446-009-0080-x.
- 701 41 James R. Wilcox, Doug Woos, Pavel Panchekha, Zachary Tatlock, Xi Wang, Michael D.
702 Ernst, and Thomas E. Anderson. Verdi: a framework for implementing and formally verifying

- 703 distributed systems. In David Grove and Steve Blackburn, editors, *Proceedings of the 36th*
704 *ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages
705 357–368. ACM, 2015. doi:10.1145/2737924.2737958.
- 706 **42** Doug Woos, James R. Wilcox, Steve Anton, Zachary Tatlock, Michael D. Ernst, and Thomas E.
707 Anderson. Planning for change in a formal verification of the raft consensus protocol. In Jeremy
708 Avigad and Adam Chlipala, editors, *Proceedings of the 5th ACM SIGPLAN Conference on*
709 *Certified Programs and Proofs*, pages 154–165. ACM, 2016. doi:10.1145/2854065.2854081.