



HAL
open science

Testing consensus implementations using communication closure

Cezara Dragoi, Constantin Enea, Burcu Kulahcioglu Ozkan, Rupak Majumdar, Filip Nksic

► **To cite this version:**

Cezara Dragoi, Constantin Enea, Burcu Kulahcioglu Ozkan, Rupak Majumdar, Filip Nksic. Testing consensus implementations using communication closure. SPLASH 2020 : ACM SIGPLAN conference on Systems, Programming, Languages, and Applications: Software for Humanity, Oct 2021, Chicago / Virtual, United States. 10.1145/3428278 . hal-03134294

HAL Id: hal-03134294

<https://inria.hal.science/hal-03134294>

Submitted on 8 Feb 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Testing Consensus Implementations Using Communication Closure

CEZARA DRĂGOI, INRIA, France and Informal Systems

CONSTANTIN ENEA, IRIF, France

BURCU KULAHCIOGLU OZKAN, MPI-SWS, Germany

RUPAK MAJUMDAR, MPI-SWS, Germany

FILIP NIKSIC, University of Pennsylvania, USA

Large scale production distributed systems are difficult to design and test. Correctness must be ensured when processes run asynchronously, at arbitrary rates relative to each other, and in the presence of failures, e.g., process crashes or message losses. These conditions create a huge space of executions that is difficult to explore in a principled way. Current testing techniques focus on systematic or randomized exploration of all executions of an implementation while treating the implemented algorithms as black boxes. On the other hand, proofs of correctness of many of the underlying algorithms often exploit semantic properties that reduce reasoning about correctness to a subset of behaviors. For example, the *communication-closure* property, used in many proofs of distributed consensus algorithms, shows that every asynchronous execution of the algorithm is equivalent to a *lossy synchronous* execution, thus reducing the burden of proof to only that subset. In a lossy synchronous execution, processes execute in lock-step rounds, and messages are either received in the same round or lost forever—such executions form a small subset of all asynchronous ones.

We formulate the *communication-closure hypothesis*, which states that bugs in implementations of distributed consensus algorithms will already manifest in lossy synchronous executions and present a testing algorithm based on this hypothesis. We prioritize the search space based on a bound on the number of failures in the execution and the rate at which these failures are recovered. We show that a random testing algorithm based on sampling lossy synchronous executions can empirically find a number of bugs—including previously unknown ones—in production distributed systems such as Zookeeper, Cassandra, and Ratis, and also produce more understandable bug traces.

1 INTRODUCTION

Large-scale, fault-tolerant, distributed systems are the backbone for many critical software services. Since they must execute correctly and efficiently in the presence of concurrent and asynchronous message exchanges as well as benign (message loss, process crash) or Byzantine failures (message corruption), the underlying algorithms are intricate. Moreover, even when the algorithms are proven correct, testing production *implementations* of these algorithms remains a significant challenge, precisely because of the enormous number of exceptional conditions that may arise in production.

Testing such distributed systems raises several important challenges:

- (C0) *Test oracle*: Formulating a correctness specification that should hold for the system and a checker for the property on a given execution.
- (C1) *Test harness discovery*: Devising a suitable set of test harnesses (combinations of user requests) that are more likely to expose vulnerabilities, e.g., sets of transactions that access a common set of data fields in the case of a distributed database.
- (C2) *Enumerating executions*: Even if the test harness contains few user requests, the number of possible executions can still be enormous because of a large number of internal steps that can interleave in arbitrary ways (the number of executions can be infinite if failures occur

Authors' addresses: Cezara Drăgoi, INRIA, France, Informal Systems; Constantin Enea, IRIF, France; Burcu Kulahcioglu Ozkan, MPI-SWS, Germany; Rupak Majumdar, MPI-SWS, Germany; Filip Nksic, University of Pennsylvania, USA.

2020. 2475-1421/2020/1-ART1 \$15.00

<https://doi.org/>

frequently and infinitely-often). An important challenge is to define efficient strategies for enumerating the execution space that maximizes the probability of exposing vulnerabilities.

(C3) *Improving interpretability*: Since a vulnerability can be exposed in many different ways, it is desirable to prioritize showing the user executions that are “easily” interpretable and that simplify the task of extracting the root cause and a possible repair.

Specifications for distributed systems is a well-studied topic, e.g., [Lynch 1996], and our paper assumes that a correctness specification is provided. We shall focus on the challenges C1–C3.

Challenge C1 is usually addressed using an exhaustive enumeration of test harnesses with few user requests. Empirically, these harnesses seem to be enough for exposing most vulnerabilities (an instance of the so-called “small scope” hypothesis). Therefore, testing techniques today focus on addressing the challenge C2 and explore message orderings, systematically or randomly, a major concern being to prioritize the search order [Desai et al. 2015; Izrailevsky and Tseitlin 2011; Killian et al. 2007; Kingsbury 2018; Leesatapornwongsa et al. 2014a; Lukman et al. 2019a; Ozkan et al. 2018]. In most existing testing approaches, the underlying distributed protocols are treated as black boxes: tests explore possible schedules of messages and faults in the implementation without considering properties of the underlying algorithms. Up to our knowledge, none of the existing techniques address the challenge C3 of improving interpretability.

In this paper, we describe a testing strategy that addresses both C2 and C3. We pick a subset of executions of a distributed system that, under some reasonable and frequently occurring assumption on the underlying algorithms, represents every other possible execution (any other execution is equivalent to one in this subset). The subset of executions is chosen to follow a symmetric and regular scheduling policy, e.g., synchronizing message exchanges between different processes. Our testing strategy explores only this subset of executions, and it is complete in the limit under the hypothesis that the semantic reduction holds. Since it explores concrete executions of the system it is clearly sound, in the sense that all reported bugs are genuine. The restriction to a subset of executions improves the likelihood that a bounded enumeration is able to expose vulnerabilities (challenge C2) while restricting the scheduling policy improves interpretability (challenge C3). This semantic reduction is mainly based on a property called *communication-closure* [Elrad and Francez 1982], which has been used extensively in designing or proving distributed protocols like Paxos [Chou and Gafni 1988; Damian et al. 2019; Dragoi et al. 2016; Moses and Rajsbaum 2002; von Gleissenthall et al. 2019].

A Semantic Reduction Based on Communication-Closure We model a fault-tolerant distributed system as a set of processes communicating through message passing. Each process maintains local state and executes a sequence of send, receive, and state update actions. Under the standard asynchronous semantics, processes may execute at arbitrarily different speeds and messages can be arbitrarily delayed or lost (process crashes can be modeled as losing all messages sent by or to a process). The space of possible executions is enormous since it is defined by all the interleavings between process actions and all possible ways of introducing message delays or losses.

As stated above, we consider a semantic reduction for such systems which is based on *communication closure*. This property relies on a restricted semantics, that we call *lossy synchronous* [Charron-Bost and Schiper 2009; Gafni 1998; Santoro and Widmayer 1989], and ensures that every asynchronous execution is *indistinguishable* from a lossy synchronous execution. Indistinguishability means that processes go through the same sequence of local states, modulo stuttering, in the two executions. Assuming that the system specification cannot make the difference between indistinguishable executions, which is the case in practice for many specifications of interest, communication-closure ensures that exploring only lossy synchronous executions is complete.

99 While our method is not complete for systems that violate communication closure, it is sound (any
100 reported bug is a true bug).

101 To define the lossy synchronous semantics, we consider that the behavior of each process is
102 structured as a sequence of *rounds*: sequences of send-receive-update actions (this decomposition
103 can be assumed without loss of generality modulo introducing fictitious actions for sending/receiv-
104 ing an empty set of messages and update actions leaving the state unchanged). For example, in a
105 distributed consensus protocol, rounds correspond to preparing a new ballot/view/term, sending
106 and receiving acknowledgments, proposing values, and communicating promises. The lossy syn-
107 chronous semantics imposes that processes execute rounds synchronously and in lock-step, but
108 messages can be lost. Any two processes are in the same round at each point during the execution
109 and all messages sent in a round are either received in the same round or lost forever (messages
110 exchanged in one round may be lost while the ones exchanged in the next round delivered without
111 failure). In contrast, under the asynchronous semantics, processes may be executing different
112 rounds at a point of time and be ready to receive messages from any round in the past or future.

113 We reduce the execution space even further for “leader-based” protocols, a widely used technique
114 for implementing state machine replication. In a leader-based protocol, the communication in each
115 round goes from one process, called *leader*, to all the other processes, or from all processes to the
116 leader. We introduce a restriction of the lossy synchronous semantics, which restricts the way
117 messages are lost in a given round. We define a *uniform* lossy synchronous semantics where the
118 messages that are lost in a given round are precisely those sent or received by a set of processes.
119 Intuitively, this corresponds to isolating each such process from all the other processes in the
120 network. This is a restriction of the lossy synchronous semantics. For instance, in the presence
121 of three processes p_1, p_2, p_3 , the uniform semantics does not allow that a message from p_1 to p_2 is
122 lost while a message from p_1 to p_3 is delivered, or it does not allow that a message from p_2 to p_1 is
123 lost while a message from p_3 to p_1 is delivered (p_1 is not isolated from all the other processes, but
124 only from p_2). It is rather easy to see that the uniform lossy synchronous semantics is complete for
125 leader-based protocols (we show in Section 5 that it is complete for a larger class of protocols).

126 Our testing algorithm enumerates only executions under the uniform lossy synchronous se-
127 mantics. While proving the validity of the reduction to such a semantics (i.e., that our testing
128 algorithm is complete in the limit) is very difficult for production systems (the kind we con-
129 sider in the experimental evaluation), the goal of our work is investigating the following *uniform*
130 *communication-closure hypothesis*: bugs in many distributed systems manifest already at the level
131 of uniform lossy synchronous executions. The validity of this hypothesis leads to a solution for
132 challenge C2 since the space of uniform lossy synchronous executions is much smaller than the
133 whole set of asynchronous executions (see Section 2 for an example) and challenge C3 because the
134 exchange of messages in such executions is quite regular and easy to interpret in comparison to an
135 arbitrary asynchronous execution. While it is hard to evaluate the degree of interpretability in an
136 objective manner, we believe through our own experience that the simple communication patterns
137 in uniform lossy synchronous executions, the lock-step exchange of messages in particular, are
138 definitely easier to debug than an arbitrary schedule of such actions.

139
140 **Testing Algorithm** We define a randomized testing algorithm which samples uniform lossy
141 synchronous executions. The algorithm takes as input a harness consisting of n processes running
142 for a maximum of r rounds. Our algorithm limits the sampling space according to several parameters
143 that bound the choice of isolated processes in each round. Note that process isolation is the only
144 source of non-determinism in the uniform lossy synchronous semantics since processes execute
145 rounds in lock-step (the interleaving between actions of different processes is fixed modulo actions
146 which commute trivially like sends done in parallel by two different processes). The first parameter
147

148 is a bound d on the number of isolated processes across all the rounds in the execution while the
149 second parameter k sets the frequency at which isolated processes re-join the network.

150 While the choice of the parameter d is motivated by an empirical “small scope” observation that
151 many bugs in implementations already occur under a rather small number of isolated processes
152 (transient faults), the second parameter k is motivated by the structure of standard distributed
153 algorithms, e.g., state machine replication algorithms. Typically, the sequence of rounds in a
154 process is further decomposed into a sequence of *phases* (a phase is a sequence of rounds) with
155 *successful phases*, when the system makes progress towards its specification, and *unsuccessful*
156 *phases*, when progress is not possible because of failures (e.g., message loss), but some computation
157 needs to be performed to ensure that the system remains safe. For example, in a state machine
158 replication algorithm, a successful phase corresponds to committing a single command (transition)
159 of the machine, provided that enough messages are delivered in each of its rounds. In more faulty
160 scenarios, i.e., when the network is temporarily partitioned such that there is no majority that
161 can communicate reliably, the system will execute several unsuccessful phases until the network
162 delivers sufficiently many messages in a phase to commit a client request. The desirable choice for
163 the rate k at which processes re-join the network equals the length of a phase in the system under
164 test. The testing algorithm uses k and d to generate executions that alternate successful phases
165 (having few to no processes isolated) and unsuccessful phases (having sufficiently many processes
166 isolated to prevent progress). However, the user is not required to have protocol specific insights
167 about the length of a phase. The testing algorithm drives the exploration through executions where
168 the set of isolated processes changes at every k rounds. The sampling space grows as d is increased
169 and k is decreased, covering the whole space of uniform lossy synchronous executions when d
170 grows to infinity and $k = 1$.

171 Our algorithm samples executions of the harness satisfying the bounds d and k , and guarantees
172 that each execution is picked with a certain minimum probability. This leads to precise probabilistic
173 guarantees about hitting a specific execution. This algorithm is sound, i.e., the reported bugs are
174 not spurious, and complete in the limit when the reduction to the uniform lossy semantics is valid.
175

176 **Evaluation** We evaluated the effectiveness of our testing algorithm on large scale distributed
177 systems such as Cassandra, Ratis, and Zookeeper. Our evaluation focuses on detecting consistency
178 violations, a major source of bugs in distributed systems. We experimentally show that our testing
179 algorithm (1) compares favorably with testing based on random search: it detects several known
180 and novel bugs by sampling from a much smaller subset of executions (showing that uniform lossy
181 executions already cover many bugs), and (2) enables exploration even with little instrumentation of
182 the source code. In particular, our testing tool was able to detect several previously unknown bugs
183 in recent versions of Zookeeper and Ratis. Moreover, the buggy traces produced by our algorithm
184 are informative. The synchronous traces are more understandable when compared with the usual
185 asynchronous ones produced by other state-of-the-art techniques.

186 The generality of our method goes beyond the evaluated benchmarks. Distributed systems
187 are all about coordination in the absence of a global clock. Communication-closure highlights
188 rounds, an encoding of a local notion of time which is used by processes to coordinate and
189 accomplish collective tasks. Rounds are a good abstraction of timestamps, vector clocks, or any
190 other synchronizations mechanism that must be implemented by a distributed protocol. The
191 communication-closed executions of a systems are the core of any protocol (even if the protocol
192 has not been shown communication-closed), because they include the executions for which local
193 time can be mapped on a global notion of time. Therefore, even for systems where our testing is
194 not complete, prioritizing communication-closed executions is an important heuristic.
195

197 **Contributions and Outline** In this paper we propose a framework for reducing the search space
198 in testing based on communication-closure, a well established design and reasoning principle for
199 fault-tolerant distributed systems.

200 Our testing framework complements theoretical concepts from the distributed computing com-
201 munity (communication closure) with novel search prioritization and randomization techniques
202 (which are specific to the use of communication closure and the systems under study). Despite
203 the fact that communication closure is a rather established and well-studied concept in theoretical
204 terms, it has never been proposed as a way of building better testing tools. Our work transfers the
205 theoretical insight to testing tools that find bugs in real-world, deployed, applications.

206 Our contributions and outline are summarized as follows:

- 207 • we develop a theoretical framework for stating and using the communication-closure hy-
208 pothesis in testing (§3 and §4),
- 209 • we define the uniform restriction of the lossy synchronous semantics prescribed by communication-
210 closure which limits message losses to isolating a set of processes and which is complete for
211 a large class of practical distributed algorithms (§5),
- 212 • we define a randomized testing algorithm with precise probabilistic guarantees that samples,
213 uniform lossy synchronous executions under certain bounds on the occurrence of network
214 link failures (§6)
- 215 • we conduct an empirical evaluation on production distributed systems (§7).

217 2 OVERVIEW

218 We demonstrate our testing framework on the distributed protocol listed in Fig. 1, where a set of
219 processes must agree on a total order between a set of commands. These commands are inputted
220 one by one while the protocol is running, and possibly concurrently, at different processes at the
221 same time. This is a simplified version of state machine replication (based on Paxos [Lamport 2005])
222 in which we omit how commands are communicated to the protocol and assume that they are
223 created by invoking a `new_command` function (see line 28). Each process maintains a sequence of
224 commands (in a local variable `log`) which is outputted when certain conditions are fulfilled (see
225 line 42). The intended specification is that any two such outputs, possibly from different processes
226 or from the same process but at different points in time, must be comparable with respect to the
227 standard prefix order between sequences. The protocol would be incorrect if for instance, two
228 processes would output a and b , respectively (since neither is a prefix of the other one).

229 The pseudocode in Fig. 1 is executed an arbitrary number of times by each process participating
230 in the protocol, and an execution of the protocol is a standard interleaving of steps from different
231 processes. Like in many other distributed protocols, each process executes a sequence of *rounds*.
232 Each round consists of a sequence of message sends, receives, and state updates, in this order. The
233 protocol periodically tries to extend the sequence of commands on which processes agree with a
234 new command by executing a sequence of four rounds, called *phase*,¹ in each process. In each phase,
235 a process, called the leader, gets a new command and tries to store into the log of a quorum formed
236 of more than half of the processes. The quorum is essential for fault tolerance. If all processes
237 execute synchronously (in lockstep) and all messages are delivered, then each process ends up
238 extending their local sequence `log` with the new command. Such an execution is given in Fig. 2(a).
239 Each process in this execution executes two phases: the first phase appends command a while the
240 second appends command b . The possible outputs are related by prefix order as expected. If too
241 many messages are lost (the cardinality constraints at lines 26 and 41 that check for a quorum are
242 not satisfied) while a process is executing a phase to process a new command, then its `log` remains
243

244 ¹In other works, a phase may be called *ballot* or *view*.


```

246 1 //Local variables
247 2 int last = phase = 0
248 3 var log = ε
249 4 var my_id, leader
250 5 var step
251 6
252 7 //@Round Prepare
253 8 //@Snd:
254 9 if (getLeader(phase) == my_id)
255 10     send to all ("Prepare", phase+1, my_id)
256 11     receive_messages()
257 12 //@Upd:
258 13 if received m=("Prepare", m.phase, m.sender)
259 14     with m.phase >= phase
260 15         last = phase //@bugfix remove
261 16         phase = m.phase
262 17         leader = m.sender;
263 18         step = "Ack"
264 19
265 20 //@Round Ack
266 21 //@Snd:
267 22 if(step=="Ack") send to leader
268 23     ("Ack", phase, (last, log))
269 24     receive_messages()
270 25 //@Upd:
271 26 if (step=="Ack") && received > n/2 messages ("Ack",phase,_)
272 27     log = select_log_from_received_messages()
273 28     log = log @ new_command()
274 29     step = "Propose"
275 30 if(my_id != leader) step = "Propose"
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295 25 //@Round Propose
296 26 //@Snd
297 27 if (leader == my_id && step == "Propose")
298 28     send to all ("Propose", phase, log)
299 29     receive_messages()
300 30 //@Upd:
301 31 if received from leader a message
302 32     m=("Propose",phase, m.log)
303 33     log = m.log
304 34     step = "Promise"
305 35 //@bugfix add last = phase
306 36
307 36 //@Round Promise
308 37 //@Snd:
309 38 if(step == "Promise") send to all
310 39     ("Promise", phase, log)
311 40     receive_messages()
312 41 //@Upd:
313 42 if received more than n/2 messages
314 43     ("Promise", phase, log)
315 44     output log

```

Fig. 1. A Paxos-like state machine replication protocol. The number of processes participating in the protocol is denoted by n . Each process has a number of local variables, listed at lines 2–5: `my_id` is a constant storing the id of the process, and `log` stores the sequence of commands to be outputted (`@` denotes sequence concatenation at line 28). The code represents a *phase* defined as a sequence of four *rounds*. Each round consists of message sends (annotation `@Snd`), receives, and state updates (annotation `@Upd`). Each phase has a designated *leader* which is set by the call to the deterministic `getLeader` function.

unchanged and it begins a new phase (the same happens if the process executes much faster than many other processes). Although the protocol should tolerate any such exceptional conditions and satisfy the intended specification, this is *not* actually true.

Fig. 3 shows an execution that violates this specification where the processes output sequences a and b , which are incomparable w.r.t. prefix order. This execution contains four phases: during the first phase, enough messages are delivered so that two processes can output a ; many messages are lost in the next two phases, so no process can extend their log; enough messages are delivered during the fourth phase, but processes end up “forgetting” about command a , and output the singleton sequence b .

In order to understand the details of the bug in Figure 3, we take a closer look at the implementation. Each process keeps track of the current phase it executes (using the local variable `phase`). Due to faults processes may be in different phases. In each phase a processes executed the four rounds in Fig. 1; the rounds are named in comments (lines 7, 20, 25, and 36). In the first round the leader looks for a quorum of processes to learn the most up-to-date log stored by its peers (the leader might have a stale local log due to faults). To this, the leader broadcasts a Prepare message that contains the leader’s phase (line 10). The processes that receive the leader’s message join the leader’s phase by updating the local phase variable to the leader’s phase, unless they are already in

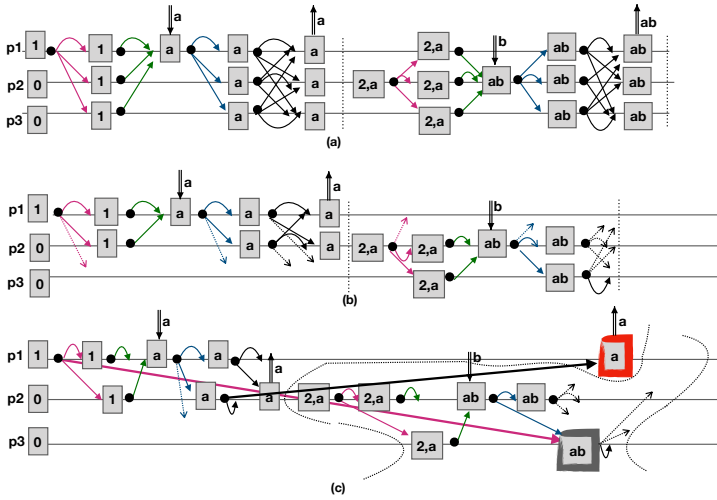


Fig. 2. Three executions of the protocol in Fig. 1 that involve three processes p1, p2, and p3: (a) a synchronous execution where no message is lost, (b) a lossy synchronous execution, (c) an asynchronous execution indistinguishable from the lossy synchronous execution in (b). Each horizontal line shows time progressing for each process. Boxes contain fragments of local state: the numbers represent the value of the phase variable while the strings represent the value of log. Colored arrows between the horizontal lines show the messages exchanged. Dotted arrows in (b) and (c) indicate dropped messages. Double arrows \Downarrow denote input commands (values returned by `new_command`) while \Uparrow denote output command sequences. Each phase ends with a vertical dotted line in the figures.

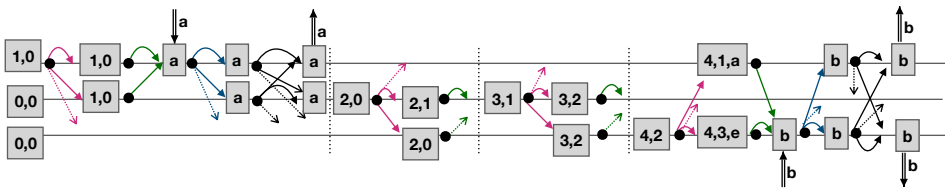


Fig. 3. An incorrect execution of the protocol in Fig. 1 (with three processes). We use the same conventions as in Fig. 2. The pairs of integers, e.g., (1,0) and (2,0), represent the values of the local variables (phase, last). The log values are e, a and b, where e denotes the empty log.

a higher one (line 16). The processes that join the leader’s phase, called followers, store in `last` (line 15) the number of the last phase they participated in. In the second round, each follower sends an Ack message to the leader (line 23), including the values of its local sequence log and the last phase the follower participated in, which is used to date the sent log. If the leader receives more than $n/2$ Ack messages, it has a quorum and sufficient information to compute the most recent value of the log: it selects the log coming from the process that participated in the most recent phase, i.e., the one with the highest value of `last` (line 27). In the third round, the leader sends the most recent log extended with a new command to all processes in a Propose message (line 28). The processes that receive the leader’s message update their log accordingly. In the last round, processes exchange their current log and phase, by sending Promise messages to all the processes (line 38). A process that receives $n/2$ Promise messages for the same phase and the same value of the log, outputs this log value.

344 The value of the last phase a process participated in is sent along with the value of the log
345 in the round Ack. This is crucial for correctness because it prevents losing requests. The bug in
346 Fig. 3 is caused by an incorrect computation of the most recent log after two phases when too
347 many messages were lost. In the fourth phase the leader receives two log values in round “Ack”:
348 a from p1 and the empty log from p3. The leader picks the empty log of p3 as being the most
349 recent one, because the last phase number accompanying it is higher than the last phase number
350 accompanying the log of p1 containing a. The bug happens because of a misinterpretation of what
351 “participating” in a ballot means. Processes should recall the last phase when they received a new log
352 from the leader, not the last phase they joined. Process p3 joins phases 2 and 3 but its algorithmically
353 meaningful state, i.e., the log, does not change in these phases. Therefore by updating the value of
354 last when receives a Prepare message from the leader, p3 incorrectly makes its log more recent
355 than it is. A correct implementation requires removing the update of last from the round Prepare
356 (line 15) and adding an update of last to phase in round Propose when the process receives the
357 leader’s new log proposal, in line 34.

358
359 **Asynchronous vs. Lossy Synchronous Semantics** The standard asynchronous semantics of this
360 protocol allows arbitrary interleavings of steps from different processes under a non-deterministic
361 network that can drop arbitrarily many messages. Different processes may execute different rounds
362 at the same time and they may receive arbitrarily delayed messages. For example, Fig. 2(c) shows
363 an asynchronous execution where some messages are lost and others are delayed. In this execution,
364 processes go through two phases: in the first one, the leader p1 transmits the command a to
365 {p1, p2}, and in the second phase, the leader p2 transmits the second command b to {p2, p3}. The
366 non-determinism in this semantics due to scheduling and message loss leads to an enormous
367 number of executions. Standard exhaustive or random enumerations of this space of executions are
368 very unlikely to be effective in exposing potential vulnerabilities like the one in Fig. 3.

369 A smaller space of executions can be defined by considering a *synchronous* semantics in which
370 each round is executed at the same time by all processes and every message is delivered. An
371 execution fragment where each process executes a round is called a *synchronized round*. Fig. 2(a)
372 shows such an execution with 8 synchronized rounds. This semantics is however too restricted
373 since it cannot exercise the protocol’s capabilities of tolerating faults, e.g., message loss.

374 An intermediate point is the *lossy synchronous* semantics, which is a weakening of the synchron-
375 ous semantics where messages can be dropped arbitrarily. An execution under this semantics is
376 still a sequence of synchronized rounds, but messages can be either delivered in the same synchron-
377 ized round they were sent or dropped and never delivered in the future. Fig 2(b) shows such an
378 execution: the leader p1 of the first phase sends the command a but only p1 and p2 receive it, and
379 in the second phase, the command b sent by the leader p2 is received only by p2 and p3.

380 In general, the lossy synchronous semantics contains a subset of the possible executions (under the
381 asynchronous semantics). However, most distributed protocols are designed to be *communication-*
382 *closed*, i.e., so that the two semantics are equivalent (every asynchronous execution is equivalent
383 to a lossy synchronous one) [Dragoi et al. 2016; Elrad and Francez 1982]. The protocol in Fig. 1 is
384 communication-closed. For example, the asynchronous execution in Fig. 2(c) is equivalent to the
385 one in Fig. 2(b), in the sense that each process passes through the same sequence of local states in
386 both executions.

387 The key observation in our testing algorithm is that, when testing for a given specification, we
388 can restrict attention only to lossy synchronous executions, instead of the much larger class of
389 asynchronous executions. Note that the bug in Fig. 3 is an incorrect lossy synchronous execution.
390 This execution represents a large class of equivalent asynchronous executions (all bugs). When the
391 underlying protocol is communication-closed, there is no loss of generality.

392

393 **Uniform Lossy Synchronous Semantics** In fact, our testing algorithm considers a further re-
 394 striction of the lossy synchronous semantics, called *uniform*, which limits the choice of messages
 395 to be dropped in a synchronized round. Consider for instance the first synchronized round in
 396 Fig. 3. Choosing to drop the message from the first to the third process is the same as choosing to
 397 isolate the third process from the rest of the processes (meaning that all the messages sent by or
 398 to the third process are dropped). Equating dropping messages with isolating a set of processes
 399 is valid for synchronized rounds where a single process sends messages or when all messages
 400 are sent to the same process. In practice, to minimize the number of exchanged messages, many
 401 distributed protocols are defined in such a way, each round (phase) having a designated *leader*
 402 (like in the first three rounds in Fig. 1). For synchronized rounds with a different communication
 403 structure, e.g., the last round in our protocol, choosing only to isolate a set of processes instead of
 404 dropping a specific set of messages *may* be a restriction that leads to incompleteness. For instance,
 405 the last synchronized round in Fig. 3 corresponds to isolating the second process. Dropping another
 406 message, say from p_1 to p_3 , could not be simulated as a set of isolated processes. As we show in
 407 Section 5 this restriction is actually complete even for this protocol.

408 **Our Testing Algorithm.** Our testing algorithm randomly samples uniform lossy synchronous
 409 executions where the number of isolated processes in the run is at most d and where every isolated
 410 process can reconnect to the network every k -th round. The values of d and k are inputs to the
 411 algorithm. Intuitively, d is a bound on the number of messages that can be dropped during an
 412 execution while k should ideally correspond to the number of rounds in a phase of the protocol
 413 (this is however not a requirement and k can be arbitrary). The latter is motivated by the fact that
 414 in many algorithms, once a process becomes isolated in a phase, it cannot make progress (change
 415 its local state) during the same phase even if it reconnects later. For the protocol in Fig. 1, if a
 416 process does not receive a “Prepare” message in the first round, it cannot change its state because
 417 of messages received in the later rounds of the same phase. As d increases and k decreases, the
 418 algorithm covers more and more of the execution space. For each execution, the algorithm applies
 419 a user-provided procedure for checking the intended specification.

420 **Advantage of Our Algorithm: Smaller Sample Set of Executions.** Sampling from uniform
 421 lossy synchronous executions reduces the size of the sample set of executions significantly. Consider
 422 the protocol execution in Fig. 1 with 3 processes, 16 synchronized rounds, and $k = 4$ (these
 423 constraints are those satisfied by the buggy execution in Fig. 3; the picture omits the last two rounds
 424 from the second and third phase because no process sends any message). The number of uniform
 425 lossy synchronous executions of the protocol is about 10^7 (each one of 3 processes can be isolated
 426 at one of $k = 4$ rounds in $16/4 = 4$ phases). In comparison, the number of lossy synchronous
 427 executions which are not necessarily uniform is about 10^{43} (any subset of the 9 communication
 428 links between the processes can be lossy in a round).
 429

430 3 DISTRIBUTED PROTOCOLS

431 We describe the theoretical foundation of our work in the context of an abstract notion of protocols
 432 that abstracts away from a particular syntax. We define the standard asynchronous semantics for
 433 such protocols, which allows arbitrary interleavings of steps from different processes and arbitrary
 434 loss of messages.
 435

436 **Protocols.** We fix a set \mathbb{P} of process identifiers and an arbitrary set \mathbb{V} of message payloads. A
 437 message is a triple $(p, q, v) \in \mathbb{P} \times \mathbb{P} \times \mathbb{V}$ where p represents the source of the message, q its destination,
 438 and v the payload. The set of all messages is denoted by \mathbb{M} . A process with identifier p is a tuple
 439 $A = (\Sigma, s_0, Snd, Upd)$ where:

- 440 • Σ is a set of process local states, and s_0 is the initial state of the process,

$$\begin{array}{c}
442 \\
443 \\
444 \\
445 \\
446 \\
447 \\
448 \\
449 \\
450 \\
451 \\
452 \\
453 \\
454 \\
455 \\
456 \\
457 \\
458 \\
459 \\
460 \\
461 \\
462 \\
463 \\
464 \\
465 \\
466 \\
467 \\
468 \\
469 \\
470 \\
471 \\
472 \\
473 \\
474 \\
475 \\
476 \\
477 \\
478 \\
479 \\
480 \\
481 \\
482 \\
483 \\
484 \\
485 \\
486 \\
487 \\
488 \\
489 \\
490
\end{array}$$

$$\begin{array}{c}
\text{SEND} \quad \frac{\mathcal{P}(p).Snd(ls(p)) = M}{(pool, ls) \xrightarrow{\text{send}(p)} (pool \cup M, ls)} \quad \text{ENVIRONMENT} \quad \frac{M \subseteq pool}{(pool, ls) \xrightarrow{\text{env}(M)} (M, ls)} \\
\text{A-UPDATE} \quad \frac{M \subseteq pool \cap (\mathbb{P} \times \{p\} \times \mathbb{V}) \text{ and } \mathcal{P}(p).Upd(ls(p), M) = s}{(pool, ls) \xrightarrow{\text{a-update}(p)} (pool \setminus M, ls[p \mapsto s])} \\
\text{S-UPDATE} \quad \frac{M = pool \cap (\mathbb{P} \times \{p\} \times \mathbb{V}) \text{ and } \mathcal{P}(p).Upd(ls(p), M) = s}{(pool, ls) \xrightarrow{\text{s-update}(p)} (pool \setminus M, ls[p \mapsto s])}
\end{array}$$

Fig. 4. Transition rules for protocol semantics.

- $Snd : \Sigma \rightarrow 2^{\mathbb{M}}$ is the message sending function: $Snd(s) = M$ denotes the fact that p sends the set of messages M when in local state s . As expected, we assume that p is the source of all the messages in M .
- $Upd : \Sigma \times 2^{\mathbb{M}} \rightarrow \Sigma$ is the state-update function: $Upd(s, M)$ is the next state of the process p given its current state s and that it received the set of messages M (we assume that p is the destination of all the messages in M).

Given a process A , we refer to components of A using $A.\Sigma$, $A.s_0$, and so on.

A *protocol* \mathcal{P} maps each process identifier $p \in \mathbb{P}$ to a process $\mathcal{P}(p)$ with identifier p .

Example 3.1. Consider the protocol in Fig. 1. A state is a valuation of the process local variables (declared in the protocol) including a variable representing the control location. The initial state s_0 of any process, has an the empty log of requests, $s_0(\text{log_val}) = \epsilon$, the ballot counter is zero, $s_0(\text{ballot}) = 0$, $s_0(\text{step}) = \text{Prepare}$, and $s_0(\text{last}) = 0$.

The functions Snd and Upd are based on the code snippets that send messages, respectively update the local state (highlighted in the figure with matching labels). For example, for any process p , given a state $s \in \Sigma$ with the program counter at lines 10 (the send of the round Prepare),

$$Snd(s) = \begin{cases} \{(p, q, (\text{"Prepare"}, s(\text{ballot}))) \mid q \in \mathbb{P}\} & \text{if } \text{get_leader}() = p, \\ \emptyset & \text{otherwise.} \end{cases}$$

For any process in some state s , if the program counter is at line 14 (the update of the round Prepare) then $Upd(s, M) = s'$ if there is $m \in M$ s.t. $m.\text{ballot} > s(\text{ballot})$ and $Upd(s, M) = s$ otherwise, where M is the current set of received messages and s' differs from s on the following variables: $s'(\text{last}) = s(\text{ballot})$, $s'(\text{ballot}) = s(m.\text{ballot})$, $s'(\text{step}) = \text{Ack}$, $s'(\text{leader}) = m.\text{sender}$.

A *configuration* of a protocol \mathcal{P} is a tuple $(pool, ls)$ where $pool$ is a set of messages in transit and ls maps each process identifier $p \in \mathbb{P}$ to a process local state in $\mathcal{P}(p).\Sigma$. Given a configuration $c = (pool, ls)$ we use $c.pool$ and $c.ls$ to refer to its components.

Asynchronous Semantics. The asynchronous semantics of a protocol \mathcal{P} is defined using a set of transition rules given in Figure 4. The rule SEND represents a transition in which a given process p sends all messages prescribed by its message sending function Snd in a given state. These messages are added to the pool of messages in transit and the process local states remain unchanged. The rule A-UPDATE represents a transition in which a set of messages M is delivered to a process p and p updates its local state according to its state-update function Upd . The set of messages M is chosen non-deterministically from the set $pool$ of messages in transit with destination p . This models

adversarial networks in which messages can be delayed arbitrarily. The rule ENVIRONMENT is used to model networks that can also drop messages arbitrarily. It defines a set of transitions that can delete an arbitrary set of messages from the pool of messages in transit. These transitions are labeled by $\text{send}(p)$, $\text{a-update}(p)$, and $\text{env}(M)$ where M is the set of messages kept by an ENVIRONMENT transition, respectively.

An *asynchronous execution* of a protocol \mathcal{P} is a sequence of transitions between configurations $c_0 \xrightarrow{\ell_0} c_1 \xrightarrow{\ell_1} \dots \xrightarrow{\ell_{m-1}} c_m$ where each $\ell_i \in \{\text{send}(p), \text{a-update}(p), \text{env}(M) : p \in \mathbb{P}, M \subseteq \mathbb{M}\}$, for all $0 \leq i \leq m - 1$. The set of asynchronous executions of a protocol \mathcal{P} is denoted by $\text{AsyncEx}(\mathcal{P})$.

Example 3.2. Fig. 2(c) shows an asynchronous execution of the protocol in Fig. 1. Each square represents a state update transition, each filled circle represents a send transition, and each edge represents a message produced by the sending function. The initial states are given by circles labeled with the initial ballot number. The execution omits send transitions that produce an empty set of messages, e.g. the first and third send actions of process p_2 . The interleaving of transitions performed by different processes is represented by the order between squares and filled circles.

Each solid edge represents a message produced during the send transition where it starts and delivered during the state update transition where it ends. Each dotted edge represents a message dropped by environment transitions. Note that some messages are delayed, i.e., they are delivered during a state update transition that occurs later in the execution and not immediately after the send transition that generated them. For instance, the message $(p_2, (\text{“Promise”}, 1, \text{a}), p_1)$ represented by the bold edge arrives with a long delay to process p_1 . The fact that the asynchronous executions can interleave send and update transitions arbitrarily is essential for modeling such delays.

4 COMMUNICATION-CLOSED PROTOCOLS

In this section we define the lossy synchronous semantics exploited by the testing algorithm, and the communication-closure property stating that this semantics is indistinguishable from the standard asynchronous semantics.

Lossy Synchronous Semantics. We consider a lossy synchronous semantics where executions are sequences of *synchronized rounds* in which all processes start by sending the set of messages determined by their local state before updating their local state using a non-deterministically chosen set of messages to receive. These rounds are communication-closed in the sense that the messages which are sent but not received within one round are lost. There is no fixed relation between the messages lost in different rounds.

Formally, a *synchronized round* between two configurations c_0 and $c_{2 \cdot n+1}$ with a set of processes $\mathbb{P} = \{p_0, \dots, p_{n-1}\}$ is a sequence of transitions

$$c_0 \xrightarrow{\text{send}(p_0)} c_1 \dots \xrightarrow{\text{send}(p_{n-1})} c_n \xrightarrow{\text{env}(M)} c_{n+1} \xrightarrow{\text{s-update}(p_0)} c_{n+2} \dots \xrightarrow{\text{s-update}(p_{n-1})} c_{2 \cdot n+1}$$

where the $\text{s-update}(\cdot)$ transitions are defined by the rule S-UPDATE in Figure 4. These transitions represent a variation of the update transitions from the asynchronous semantics where *all* messages which are still in transit are received and used to update the state of a process. A process may still receive a subset of the sent messages because of the $\text{env}(\cdot)$ transition scheduled before all update transitions.

We use $c_0 \xrightarrow{\text{round}(M)} c_{2 \cdot n+1}$ to denote the sequence of transitions in a synchronized round. A *lossy synchronous execution* is a sequence of synchronized rounds $c_0 \xrightarrow{\text{round}(M_0)} c_1 \dots \xrightarrow{\text{round}(M_{m-1})} c_m$. The set of lossy synchronous executions of a protocol \mathcal{P} is denoted by $\text{SyncEx}(\mathcal{P})$. All synchronous executions we consider are lossy synchronous.

540 *Example 4.1.* Fig. 2(a) and Fig. 2(b) show two lossy synchronous executions of the protocol in
 541 Fig. 1. The conventions for representing send and update transitions, and messages are the same as
 542 in Example 3.2. The transitions that are aligned vertically are ordered from top to bottom.

543 For the execution in Fig. 2(a), it is assumed that the environment transitions preserve the content
 544 of the pool of messages in transit (no messages are dropped). Under the synchronous semantics
 545 no messages are delayed and all send and update transitions are executed in lock-step: the k^{th}
 546 send (resp., update) is executed simultaneously on all processes. This execution goes through eight
 547 rounds, each process iterating twice over the code in Fig. 1.

548 For the execution in Fig. 2(b), the environment transitions drop the messages represented by
 549 dotted edges.

550 **Communication-Closed Protocols.** The *behavior* of a process p in a (synchronous or asynchro-
 551 nous) execution $\eta = c_0 \xrightarrow{\ell_0} c_1 \xrightarrow{\ell_1} \dots \xrightarrow{\ell_{m-1}} c_m$, denoted by $\eta \downarrow p$, is the sequence of states of p in the
 552 configurations c_0, \dots, c_m , i.e., $\eta \downarrow p = c_0.ls(p) \dots c_m.ls(p)$. Two sequences of local states σ and σ' are
 553 called *equivalent up to stuttering*, denoted $\sigma \equiv \sigma'$, when they coincide modulo removing consecutive
 554 repetitions of the same state. An execution η_1 is *indistinguishable* from another execution η_2 , which
 555 is denoted by $\eta_1 \equiv \eta_2$, if $\eta_1 \downarrow p \equiv \eta_2 \downarrow p$ for each $p \in \mathbb{P}$.

556 *Example 4.2.* The executions in Fig. 2(b) and Fig. 2(c) are indistinguishable. The executions show
 557 only (the modification of) the values of the variables `ballot` and `log_val`. The values of the other
 558 variables are also equal modulo stuttering. For example, p_1 goes through the states s_0, s_1, s_2, s_3, s_4 in
 559 both executions where s_0 is the initial state, $s_1(\text{ballot}) = 1$, $s_1(\text{log_val}) = \epsilon$, $s_1(\text{step}) = \text{"Prepare"}$
 560 $s_2(\text{ballot}) = 1$, $s_2(\text{log_val}) = a$ and $s_2(\text{step}) = \text{"Propose"}$. The states s_3 and s_4 differ from s_2 only
 561 in the value of the variable `step`, i.e. $s_3(\text{step}) = \text{"Promise"}$ and $s_4(\text{step}) = \text{"Prepare"}$.

562 *Definition 4.3.* A protocol \mathcal{P} is called *communication-closed* when for each asynchronous execu-
 563 tion $\eta_1 \in \text{AsyncEx}(\mathcal{P})$ there is a lossy synchronous execution $\eta_2 \in \text{SyncEx}(\mathcal{P})$ such that $\eta_1 \equiv \eta_2$.

564 Communication-closure is a property which is met by all the replicated state machine or consen-
 565 sus protocols we are aware of, e.g., Paxos [Lamport 2005], Multi-Paxos [Chandra et al. 2007],
 566 EPaxos [Moraru et al. 2013], ViewStamped [Oki and Liskov 1988]. Intuitively, this property is
 567 achieved using the following principles: (1) each process uses a set of variables to encode a local
 568 notion of time, called round number, which is monotonically increasing, (2) every message carries
 569 some metadata that associates it with some unique round number, and (3) a process updates its state
 570 using only messages whose round number equals the process's local round number. Assuming these
 571 constraints, any asynchronous execution can be rewritten to an indistinguishable synchronous
 572 execution by essentially, reordering commutative transitions [Damian et al. 2019; Elrad and Francez
 573 1982; Moses and Rajsbaum 2002].

574 For example, the round number of the protocol in Fig. 1 is defined by the values of the pair of
 575 variables (`ballot`, `step`). We consider the lexicographic order over the values of (`ballot`, `step`)
 576 where the four values of the variable `step` are ordered as "Prepare" < "Ack" < "Propose" < "Promise"
 577 (`ballot` is an integer variable and its values are ordered as usual), and define the round number of a
 578 process in state s as the position in the lexicographic order of the values of (`ballot`, `step`) in s . Then,
 579 every sent message m has two fields $m.\text{ballot}$ and $m.\text{step}$ that represent its round number (in the
 580 same way as the pair of local variables (`ballot`, `step`) represents the process's local round number).
 581 The third condition relates message round numbers with process round numbers. Before using the
 582 payload of a received message to update the local state, e.g., before reading $m.\text{sender}$ at line 18
 583 or $m.\text{log_val}$ at line 27 and storing their values in some local variable, the code ensures that the
 584 round number of the message equals the process's local round number, i.e., $m.\text{ballot} == \text{ballot}$
 585
 586
 587
 588

and $m.\text{step} == \text{step}$. If this is not the case, the message is either not used to update the local state or the round number of the process is first increased to match the message's round number at line 16 before using the message's content to update the state at line 18.

When systems are not known to be communication-closed, one can identify the subset of communication-closed executions. In this case, the lossy synchronous executions represent a subset of the set of executions of the distributed system.

5 UNIFORM EXECUTIONS

In this section, we present a restriction of the lossy synchronous semantics in which the faults (message losses) modeled by the environment transitions are uniform, that is the messages sent by a subset of the process are received. This restriction is complete for standard state machine replication and consensus protocols, up to indistinguishability.

A synchronized round $c \xrightarrow{\text{round}(M)} c'$ is called *uniform* if there exists a set of processes Π such that the set of messages received in the round (by some process) is *exactly* the set of messages sent by a process from Π to a process in Π , i.e.,

$$((p, q, v) \in \mathcal{P}(p).\text{Snd}(c.\text{ls}(p)) \wedge \{p, q\} \subseteq \Pi) \Leftrightarrow (p, q, v) \in M,$$

for every p, q, v . The set of processes Π is called the *kernel* of the round. A lossy synchronous execution is called *uniform* when it is a sequence of uniform rounds.

Example 5.1. The synchronous executions in Fig. 2(a), Fig. 2(b) (described also in Example 4.1), and Fig 3 are uniform. In Fig. 2(a), the kernel of each synchronized round is the set of all processes. For the execution in Fig. 2(b), $\Pi_1 = \{p1, p2\}$ is the kernel of the first four synchronized rounds (the first phase), $\Pi_2 = \{p2, p3\}$ is the kernel of the next three synchronized rounds (the first three rounds of the second phase), and $\Pi_3 = \{p3\}$ is the kernel of the last synchronized round.

Figure 5(a) shows a non-uniform execution, where in the last synchronized round, process $p1$ receives messages from $\{p1, p2\}$, the message from $p3$ being lost, and $p2$ receives messages from $\{p2, p3\}$, the message from $p1$ being lost.

A synchronized round $c \xrightarrow{\text{round}(M)} c'$ is *one-to-all* if there exists at most one process p sending messages in this round, i.e., $\mathcal{P}(q).\text{Snd}(c.\text{ls}(q)) = \emptyset$ for every $q \neq p$, and *all-to-one* if all processes send messages to a single process p , i.e., for every q , if $\mathcal{P}(q).\text{Snd}(c.\text{ls}(q)) \neq \emptyset$, then there exists $v \in \mathbb{V}$ such that $\mathcal{P}(q).\text{Snd}(c.\text{ls}(q)) = \{(q, p, v)\}$. A protocol \mathcal{P} is *leader-based* iff all its synchronous executions are sequences of one-to-all or all-to-one synchronized rounds.

Example 5.2. For the protocol in Example 3.1 (Figure 1), a synchronized round where the leader sends a "Prepare" message to all processes (the first round a phase) is *one-to-all* while a synchronized round where processes send an "Ack" message to the leader is an *all-to-one* round (the second round a phase). Note that *all* refers to the maximum number of processes that can receive, resp., send messages, in a synchronized round (messages can be dropped during an environment transition).

The following theorem implies that the restriction to uniform lossy synchronous executions is complete for leader-based protocols which are also communication closed.

THEOREM 5.3. *Every lossy synchronous execution of a leader-based protocol is uniform.*

All benign consensus and replicated state machine implementations [Junqueira et al. 2011; Lakshman and Malik 2010a; Moraru et al. 2013] are leader-based, and hence satisfy Th. 5.3. However, our running example in Fig. 1 is not leader-based since the last round uses an *all-to-all* communication. In the Promise round, all processes that received the leader's proposed log (in the previous round)

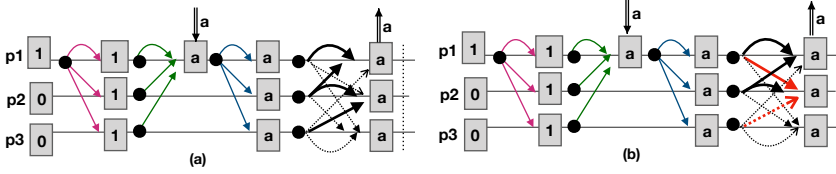


Fig. 5. Two synchronous executions of the protocol in Figure 1.

Table 1. The comparison of sample set sizes of different algorithms.

Set of executions with:	Upper bound on the size
Arbitrary message losses	$2^{n^2 r}$
k -periodic losses	$k^{n^2 r/k}$
d -bounded k -periodic losses	$\leq C(n^2 r, d) \leq (n^2 r)^d$
Arbitrary uniform executions	2^{nr}
k -periodic uniform executions	$k^{nr/k}$
d -bounded k -periodic uniform executions	$\leq C(nr, d) \leq (nr)^d$
Arbitrary reorderings	$\leq (n^2 r)!$
d -bounded reorderings (PCT)	$\leq w \cdot C(nr, (d-1)) \cdot (d-1)!$

broadcast this proposal to all the processes in the network. A process that receives more than $n/2$ messages, having as payload the same log value, transmits this log to the client. The uniform executions with all-to-all communication are a strict subset of the lossy synchronous executions.

The protocol in Fig. 1 confirms, beyond leader-based algorithms, the hypothesis that bugs manifest in uniform executions, as the incorrect execution from Fig. 3 is uniform. The underlying principle is that for any non-uniform execution of the protocol in Fig. 1 either there exists an indistinguishable uniform execution or there exists a uniform execution that exposes the same log values to the client. Fig. 5(b) shows a uniform execution that is indistinguishable from the execution in Fig. 5(a). The equivalence relation between non-uniform and uniform executions (w.r.t. the client's observations) is proved using a key insight from consensus proofs: a process communicates with the client only when the system is in a univalent (global) state, i.e., $|\{p \mid \log(p) = val \wedge \text{last}(p) \geq b\}| > n/2$ for some integer b , which means that val is a stable prefix of the log.

Finally, note that the protocol in Fig. 1 continues to solve state machine replication if we replace the last all-to-all round with an all-to-one round. In the modified Promise round processes send a Promise message only to the leader (instead of broadcasting it) acknowledging its proposal and only the leader transmits the log to the client, in case it received $n/2$ Promise messages.

6 RANDOM SAMPLING FROM UNIFORM EXECUTIONS

We now present our testing algorithm. Theoretically, the effectiveness of the testing algorithm is based on the fact that it samples from a relatively small (yet, complete in the limit) set of executions, rather than from all possible executions of a protocol.

6.1 The Space of Executions

Before giving the sampling procedure, let us consider the size of the space of executions, and compare the space of executions to other techniques. For the comparison, we consider a test harness consisting of n processes running a total of r rounds. In addition to these parameters, we consider two additional parameters to prioritize the search: the *periodicity* k and the number of

687 isolated processes d . Given a lossy synchronous execution τ , we say that a process p starts at round
 688 i in τ if p is included in the kernel of the i -th round in τ but it is not included in the kernel of the
 689 previous round (round $i - 1$). Then, a uniform execution τ is k -periodic if a process can start only at
 690 a round which is a multiple of k .

691 Consider the uniform execution in Figure 3. It has 4 phases and 4 rounds in each phase. The
 692 figure omits the last two “empty” rounds in the second and third phase, where no messages are
 693 sent. The 4-periodic execution of this example recovers isolated processes after every $k = 4$ rounds,
 694 that is in the beginning the second phase with ballot 2, the third phase with ballot 3, and the
 695 fourth phase with ballot 4. The k -periodic executions take the empty rounds into account.

696 *k*-periodic uniformity. Consider an execution with n processes running a protocol with r rounds.
 697 In a non-uniform execution, any subset of the n^2 communication links can have a message loss
 698 in each round, resulting in $2^{n^2 r}$ possible executions. In a uniform execution, the corresponding
 699 number is 2^{nr} . In a k -periodic non-uniform execution, each of the links can be broken at any k
 700 rounds in all r/k phases, resulting in $k^{n^2 r/k}$ executions. In a k -periodic uniform execution, by a
 701 similar argument, the number of executions is $k^{nr/k}$. For the example in Figure 3 with 3 processes,
 702 4 rounds and 4 phases, the sample set of executions is around 10^{43} for non-uniform executions and
 703 only around 10^7 for 4-periodic uniform executions.

704 *d*-bounding. While k -periodic uniformity already reduces the size of the execution space, bound-
 705 ing the set to d -bounded k -periodic uniform executions, i.e., executions with d isolated processes
 706 over all rounds, further reduces it. This bound reduces the asymptotic size of the space of executions
 707 so that it is exponential only in the bounding parameter d but polynomial in the number of rounds
 708 and processes. The bounded version of the non-uniform case has an upper bound of $(n^2 r)^d$. When
 709 we further restrict to the uniform case, we get an upper bound of $(nr)^d$. The actual sample set is
 710 smaller for $d > n$ since we cannot isolate more than n processes into a period of k rounds.²

711 Table 1 summarizes upper bounds on the number of executions for various choices (arbitrary
 712 message losses vs. uniform executions, k -periodic, and d -bounded k -periodic). Additionally, it shows
 713 the number of executions explored by a state-of-the-art sampling algorithm (PCT [Kulahcioglu
 714 Ozkan et al. 2018]) that is oblivious to rounds.

715 The size of the set of d -bounded k -periodic uniform executions is asymptotically smaller than
 716 the others on Table 1. Moreover, the characterization of the bounding parameter for k -periodic
 717 uniform executions requires a smaller value of d to reproduce an execution.

719 6.2 The Sampling Algorithm

720 Our testing algorithm (Algorithm 1) takes a test harness consisting of a set \mathbb{P} of n processes running
 721 at most r rounds, and randomly samples from the set of k -periodic uniform executions with at
 722 most d isolated processes, i.e., from a sample space of size at most $(nr)^d$. The algorithm ensures
 723 that each execution is picked with probability at least $1/(nr)^d$.

724 Given the set of processes \mathbb{P} , upper bound on rounds r , and the parameters k and d , the algorithm
 725 distributes the d failures into r/k phases (line 1). For each phase, in its first round (line 5), the
 726 algorithm selects a set of d_{phase} processes to isolate in the current phase (line 6). For each of the d_{phase}
 727 selected processes, the algorithm chooses the first round in which the process is isolated (line 7). The
 728 algorithm isolates these processes by simply dropping them from the kernel of the corresponding
 729 rounds (line 8). We write $f^{-1}([0, n])$ to denote $\bigcup_{0 \leq i \leq n} f^{-1}(i)$ and use this to propagate process
 730 isolation in a phase until the end of that phase. The algorithm simulates re-establishment of faulty
 731 links by resetting the isolated set of processes in every k rounds.

733 ² The size of d -bounded k -periodic set of executions can be more precisely characterized by inclusion-exclusion principle
 734 [Charalambides 2018] or using q-binomial coefficients [Kac and Cheung 2001].

736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784

```

Input: A test harness with a set  $\mathbb{P}$  of  $n$  processes and at most  $r$  rounds
Input Parameters: A period  $k$  and a bound  $d$  on the number of isolated processes
1 distribute  $d$  into  $d_0, \dots, d_{(r/k-1)}$  s.t.  $\sum_{0 \leq i < r/k} d_i = d$  and  $d_{0 \leq i < r/k} \leq |\mathbb{P}|$ ;
2 for  $i := 0$  to  $r - 1$  do
3    $phase := i / k$ ;
4    $roundInPhase := i \% k$ ;
5   if  $roundInPhase = 0$  then
6     choose u.a.r.  $d_{phase}$  processes from  $\mathbb{P}$  as  $\mathbb{P}_{phase}$  ;
7     choose u.a.r.  $f : \mathbb{P}_{phase} \rightarrow [0, k - 1]$ ;
8     schedule round with kernel  $\mathbb{P} \setminus f^{-1}([0, roundInPhase])$ ;
9 check specification on execution trace

```

Algorithm 1: Randomized sampling from k -periodic uniform executions with bound d .

The algorithm can be modified to sample executions with an unbounded number of isolated processes. For this, we omit the parameter d together with the lines 1 and 6 in the algorithm. On line 7, we isolate any process at any round.

PROPOSITION 6.1 (SOUNDNESS AND RELATIVE COMPLETENESS). (1) *Algorithm 1 samples each synchronous uniform executions of periodicity k and up to d isolated processes with probability at least $1/(nr)^d$.* (2) *Let \mathcal{P} be a leader-based communication-closed distributed protocol. For any asynchronous execution of \mathcal{P} , there is a test harness and parameters d and k such that Algorithm 1 run on the harness with (d, k) samples an indistinguishable execution with positive probability.*

The bugs reported by the testing algorithm are not spurious as the testings enumerates actual executions of the system under test. The applicability does not depend on whether the system under test is indeed communication-closed, that is if all asynchronous executions have a synchronous indistinguishable counter-part. If the system is not communication closed the algorithm will cover an important sub-set of executions.

7 EXPERIMENTAL EVALUATION

We present an empirical evaluation of our approach on production implementations of three fault-tolerant protocols: Cassandra’s Paxos [Lakshman and Malik 2010b], Zookeeper’s atomic broadcast (ZAB) [Hunt et al. 2010], and the Raft [Ongaro and Ousterhout 2014] implementation in Ratis. This evaluation addresses the following research questions:

RQ1 Is our testing algorithm effective at detecting fault tolerance bugs in large scale systems?

RQ2 How do the algorithm parameters affect the efficacy in detecting bugs?

RQ3 How do different implementations of our algorithm affect the effectiveness at detecting bugs?

To address **RQ1** we show that our framework is indeed able to discover bugs in these implementations, some of them being unknown before our work. We also compare its effectiveness with a baseline approach that explores arbitrary asynchronous executions with arbitrary message losses.

For **RQ2**, we tested each system under varying bounds for the number of isolated processes. For Cassandra, we also evaluated the effect of varying the periodicity of isolation recovery.

For **RQ3**, we experimented with three implementations of Algorithm 1, that provide different approximations of the lossy synchronous semantics. These implementations differ in the instrumentation effort and required information about the internals of the system under test.

785 *Heavy system instrumentation.* This is a precise implementation of Algorithm 1 that instruments the
786 system in order to enforce the lossy synchronous semantics and to control the isolation of processes
787 precisely. This requires identifying the messages sent in a certain round and controlling their
788 delivery so that they are delivered only in the context of the same synchronized round they were
789 sent (or dropped). The round of a message is identified by looking at the metadata stored in that
790 message. The presence of such metadata is actually a common design principle for fault-tolerant
791 systems [Fekete and Lynch 1990]. To control the delivery of messages, the instrumentation adds a
792 layer on top of the network which collects the messages in flight, and enforces their delivery to be
793 synchronous. We used this implementation to test Cassandra.

794 *Lightweight system instrumentation.* This implementation looks at the metadata stored in the
795 messages to identify those that should be dropped according to Algorithm 1, but it only approximates
796 the lossy synchronous semantics. In this approximation, processes execute a *phase* in lockstep,
797 but they may run the rounds inside the same phase asynchronously. The lockstep execution of
798 phases is enforced using high-enough timeouts, which ensure that each process terminates a phase
799 before advancing in the execution (a phase usually corresponds to handling one client request). We
800 implemented this approach for testing Ratis.

801 *No system instrumentation.* A coarse version of Algorithm 1 can be implemented using only the
802 API methods of the system under test (treating the system as a black-box). The tester uses timeouts
803 to enforce a lockstep execution of phases, but does not look inside messages to decide which ones
804 should be dropped. Instead, it uses API methods for stopping or starting a process at the beginning
805 of a phase as an approximation for isolating/deisolating a process during a phase. We used this
806 approach for testing Ratis and Zookeeper.

808 7.1 Cassandra

809 Cassandra ensures serializability of transactions using an implementation of Paxos. This protocol is
810 used to make different processes (replicas) agree on an order in which to execute the transactions
811 submitted by the client. Each phase consists of six “one-to-all” or “all-to-one” rounds similar to those
812 in Fig. 1: Prepare/Promise, Propose/Accept, and Commit/Ack (therefore all its lossy synchronous
813 executions are uniform).

814 We test Cassandra using a harness with three processes and three transactions, two of which
815 update the same key. At the end of the tests, we read the values of the keys and check for the
816 serializability of the processed transactions. This harness admits a difficult to detect buggy behavior
817 in Cassandra 2.0.0 when messages are lost at subtle points of execution [Apache 2013]: one of the
818 processes does not receive the messages sent during the rounds processing the first two transactions,
819 and when this process becomes a leader instead of trying to process a third transaction, it recommits
820 the first one that was already executed, violating serializability.

821 We tested Cassandra using a precise implementation of Algorithm 1, that controls the messages
822 to be dropped or their delivery (the “heavy system instrumentation” described above). We bounded
823 the length of the executions to at most 24 rounds.³

824 *The effect of varying parameters.* We evaluate the effect of varying the values of the parameters
825 d and k when testing with the harness described above. For each assignment of parameters, we
826 sampled 1000 executions. For each set of tests, we report in Table 2 the average number of rounds
827 and phases that are executed by a quorum of processes⁴ (as $\#rnds\checkmark$ and $\#phs\checkmark$), in addition to the

829 ³Source code at <https://github.com/burcuku/explorer-server>

830 ⁴The parameters d and k affect the distribution of the isolated processes in an execution, which in turn may affect the
831 length of an execution. The processing of the three transactions can finish in 18 rounds if no messages are lost, or more
832 rounds when processes are isolated and quorums cannot be formed.

Table 2. The number of buggy executions detected by sampling from d -bounded k -uniform executions. On the left, we list the results for $d = 8$ and varying k . On the right, we list them for $k = 6$ and varying d .

k -uniform	#rnds	#rnds✓	#phs	#phs✓	#msgs	#buggy	d -bounded	#rnds	#rnds✓	#phs	#phs✓	#msgs	#buggy
$k = 1$	21.67	18.14	3.61	2.87	49.13	0	$d = 3$	19.08	18.08	3.18	3.00	48.47	0
$k = 2$	21.60	18.07	3.60	2.87	48.87	0	$d = 4$	20.11	18.29	3.35	2.99	48.62	0
$k = 4$	22.80	17.53	3.80	2.64	46.76	0	$d = 5$	20.91	18.17	3.48	2.93	47.90	1
$k = 6$	22.86	17.10	3.81	2.63	44.78	2	$d = 6$	21.69	17.98	3.61	2.86	47.13	1
$k = 8$	23.71	6.61	3.95	1.03	20.23	0	$d = 8$	22.86	17.10	3.81	2.63	44.78	2
$k = 10$	23.81	6.36	3.97	0.94	19.60	0	$d = 10$	23.61	15.72	3.93	2.31	41.83	1

average number of rounds (**#rnds**), phases (**#phs**), messages (**#msgs**), and the number of times a buggy execution is sampled (**#buggy**). We mark a round to have a quorum if the kernel of that round consists of a majority of processes. Similarly, we mark a phase to have a quorum if the corresponding user request takes effect (i.e., a written value is committed) on a majority of processes.

The left of Table 2 lists the results when varying $k = \{1, 2, 3, 4, 6, 8\}$ and fixing $d = 8$ (this value of d is high-enough for reproducing the bug). For values of k smaller than the number of rounds in a phase, executions have a higher number of rounds and phases with a quorum. This can be explained by the fact that the isolated processes get a chance to recover from message losses during the execution of the phase. As k increases, fewer rounds have a quorum, resulting in an increase in the total number of rounds. When $k > 6$, links are not re-established at the beginning of a phase and faults propagate to succeeding phases. This causes the protocol to fail to process user requests in later phases. Only about a single phase is successful for $k = 8, 10$ on average.

The data on the right of Table 2 shows that as d increases, the average number of rounds and phases executed by a quorum of processes decreases due to a higher frequency of message losses. Consistently, the average of the total number of rounds and phases in an execution increases due to the repetition of no-quorum phases. In the extreme case with an unbounded number of isolated processes, a minority of rounds are executed by a quorum, failing to process even a single request on average. Tests with a bounded number of isolated processes produce executions with both quorum and no-quorum phases which are more likely produce a buggy behavior. In our experiments, we could reproduce the bug by taking $d \in \{5, 6, 8, 10\}$.

Testing Cassandra with a baseline algorithm. As a baseline for testing fault tolerance of a system against network failures, we consider a naive randomized algorithm. This algorithm samples from the set of executions with arbitrary message losses, by randomly dropping a message with some probability. We tested Cassandra 1000 times using different probabilities $p = 0.125, 0.25, 0.5$. In our evaluation, none of those tests could hit the bug in the system. The infrequency of hitting the bug is not surprising since the bug in Cassandra is known to be a difficult bug and it is reproduced only in few executions in previous works [Kulahcioglu Ozkan et al. 2019; Leesatapornwongsa et al. 2014b].

7.2 Ratis

Ratis [Apache 2020] is an implementation of the Raft protocol [Ongaro and Ousterhout 2014], usable in large-scale systems such as Hadoop Ozone key-value store. Ratis is in early stages of development, currently in version 0.6.0. Raft is a consensus protocol for state machine replication. Similarly to Paxos and our motivating example, operations on the state machine are sent to the leader of the Ratis cluster. The leader appends operations to its log and replicates the operations to other servers. An operation is committed once the leader receives acknowledgements from a majority of servers. Differently from Paxos, a server can become leader only if its log is at least as up-to-date with the other servers. Raft consists of *leader election* or *log replication* rounds. The servers exchange RequestVote/RequestVoteReply messages for leader election, and AppendEntries/

883 AppendEntriesReply messages for log replication and as heartbeat messages. Similarly to other
884 consensus protocols, Raft uses only “one-to-all” and “all-to-one” rounds.

885 We tested Ratis using an implementation of our algorithm based on lightweight instrumentation.⁵
886 A test harness consists of a number of client requests submitted to the Ratis cluster and the maximal
887 number of rounds in an execution, approximated using a timeout. When processed, each request
888 extends the replicated log with some message. During the processing of the requests, we introduced
889 message losses as prescribed by our algorithm. At the end of a test, we ran the system without
890 failures for some time to allow the cluster to recover and synchronize its servers. Finally, we
891 checked whether the system could tolerate the introduced message losses by checking the following
892 properties extracted from [Ongaro and Ousterhout 2014] and the unit tests in Ratis:

893 P1 The servers eventually elect a leader.

894 P2 All servers eventually store all log entries.

895 P3 After sending a request, a client eventually receives a reply.

896 While these specifications are liveness properties, we checked for bounded-liveness variations
897 where they are required to be satisfied within a bounded amount of time. To define the time bounds
898 we use a heuristic similar to [Killian et al. 2007]. We run the system without any message loss
899 (failures) several times to determine the average time required to synchronize the servers. In our
900 tests, we allowed the system to run significantly longer to recover after the message losses.

901 We tested Ratis using $n = 3$ servers, 4 client requests, and a varying number of failures (isolated
902 processes) distributed into $r = 8$ rounds. The number of rounds is counted based on the size of
903 the replicated log (which is observed by the instrumentation). We used a period $k = 2$ to recover
904 isolated processes. At the end of the 8 rounds, we continue running the system without any failures
905 leaving a timeout of 2 seconds to allow the servers synchronize. Ratis has significant amount of
906 support code for the transport layer libraries it uses, namely gRPC and Netty. This can lead to
907 different system behavior when run with different transport options. To cover both behaviors, we
908 tested Ratis using both gRPC and Netty libraries.

909 *Testing Ratis using the lightweight system instrumentation.* We tested an instrumented version of
910 Ratis which enables our algorithm to read the content of in flight messages and be able to drop
911 them. The algorithm uses the information in the messages (more specifically, the size of the sender’s
912 log) to identify the current round of a server. Then, we isolate selected servers in selected rounds
913 by dropping the messages of those rounds from/to the isolated servers.

914 We tested Ratis 1000 times using different values for the bound on the number of isolated
915 processes $d = 1, \dots, 7$. In Table 3, we list the number of violations to the specifications P1, P2 and P3
916 detected in our tests for each value of d . In many test executions with gRPC, we observed violations
917 to P2 or P3. In the failing tests, a follower server has inconsistent entries with the leader, and sends
918 a negative reply to leader’s AppendEntries message. Inconsistency in the servers logs can arise
919 when the leader cannot fully replicate all of the entries in its log, e.g., when it disconnects before
920 sending AppendEntries messages. In the problematic executions, the leader and the follower with
921 inconsistent entries repeatedly send the same messages to each other and fail to synchronize in
922 hundreds of exchanged messages. Our bug report for this problem is currently open.⁶ In our failing
923 tests with Netty, we discovered a liveness bug which causes the violation of P3. In the buggy
924 execution, the leader gets disconnected from the cluster after it receives a client request. Then,
925 the cluster elects a new leader. While the client is successfully redirected to the new leader in the
926 implementation for the gRPC adapter, the implementation for Netty causes the client to indefinitely
927 wait for a reply from the old leader. Our bug report for this problem is already acknowledged by the
928

929 ⁵Source code is available at <https://github.com/burcuku/explorer-server>.

930 ⁶<https://issues.apache.org/jira/projects/RATIS/issues/RATIS-946>

Table 3. The number of violations to properties P1, P2 and P3 in Ratis detected by our algorithm using lightweight system instrumentation.

	d	1	2	3	4	5	6	7
Ratis with gRPC	P1	0	0	0	0	0	0	0
	P2	121	199	242	192	103	65	61
	P3	0	0	2	5	22	64	111
Ratis with Netty	P1	17	291	418	576	917	986	995
	P2	362	592	710	778	958	989	995
	P3	151	285	331	472	888	984	992

Table 4. The number of violations to properties P1, P2 and P3 in Ratis detected by our algorithm *without* system instrumentation. On the left, we list the results for the implementation using server blocking methods in Ratis test API. On the right, we list them for the implementation using server kill/restart methods.

	d	1	2	3	4	5	6	7
Ratis with gRPC	P1	0	0	0	0	0	0	0
	P2	0	0	0	1	0	0	0
	P3	0	1	16	88	182	366	523
Ratis with Netty	P1	0	0	2	0	0	0	0
	P2	0	0	1	1	0	2	9
	P3	0	9	69	159	262	497	620

	d	1	2	3	4	5	6	7
Ratis with gRPC	P1	0	0	0	0	0	0	0
	P2	0	0	0	12	23	47	57
	P3	0	18	110	197	205	276	319
Ratis with Netty	P1	0	0	1	3	1	3	7
	P2	0	0	1	0	0	0	0
	P3	0	16	11	96	93	118	79

Table 5. The number of violations detected in Ratis by using a baseline randomized testing algorithm which drops messages with a given probability. We rely on our instrumentation for selectively dropping messages.

p : probability of dropping a message	0.125	0.25	0.50	p : probability of dropping a message	0.125	0.25	0.50		
Ratis with gRPC	P1	1	2	6	Ratis with Netty	P1	994	971	497
	P2	0	1	25		P2	998	983	462
	P3	0	2	155		P3	999	996	179

Ratis developers.⁷ We also observed high number of tests where the servers cannot elect a leader (failing P1) when some messages are dropped. This violation occurs frequently and it is produced by dropping almost any message in the log synchronization of the servers. Our bug report for this violation is also currently open.⁸

Testing Ratis without additional instrumentation. We also implemented two coarser versions of our algorithm where we only use the methods provided by Ratis test API. In one of the implementations, we isolated the servers by using Ratis test API's server isolation methods which block outgoing/incoming messages from/to servers. In the other one, we used server kill and restart methods to isolate servers for some duration. In our implementations, we distributed d number of process isolations into a number of phases which are approximately determined by some timeouts. At the beginning of each phase, we isolated a randomly sampled subset of processes. If the phase has a majority of processes alive, we wait until the system elects a leader (the Ratis API provides a method for checking the leader of a cluster) and submitted 3 client requests. After that, we isolated some other randomly sampled processes and we wait for 2 seconds for the servers to process the requests. At the end of the phase, we recover the isolated processes for the next phase. We ran the system 1000 times for each value of $d = 1, \dots, 7$.

On the right of Table 4, we list the number of violations to P1, P2 and P3 detected by testing the system using the Ratis API blocking methods. Some tests detects violations of P3, where the executions fail to serve some client requests within timeout. However, the frequency of executions

⁷<https://issues.apache.org/jira/projects/RATIS/issues/RATIS-844>

⁸<https://issues.apache.org/jira/projects/RATIS/issues/RATIS-1048>

981 violating P1 or P2 is very low. A reason for these tests to miss violations might be the behavior
982 of process isolation methods in the Ratis test API. Instead of dropping messages, the isolation
983 methods block messages of a process by sleeping the thread delivering the message until the server
984 is deisolated. This might result in servers to process blocked messages once they are deisolated. In
985 our instrumentation, messages from/to the isolated process are dropped completely.

986 On the left of Table 4, we list the number of violations by testing the system using the Ratis
987 server kill/restart methods. In these tests we can observe violations to all P1, P2 and P3, in smaller
988 numbers than the tests with instrumentation. A reason for that might be blocking processes for
989 some duration is coarse grained and less selective on which particular messages will be dropped.

990 *Testing Ratis with a baseline algorithm.* Table 5 lists the number of violations detected by a naive
991 random algorithm, which samples from the set of executions with arbitrary message losses. We
992 rely on our instrumentation for dropping messages. The algorithm takes a probability value p
993 as input and drops a message with the probability p . For each different value of the probability,
994 $p = 0.125, 0.25, 0.5$ we tested the system with 1000 executions. In Netty, the tests produce executions
995 which violate P1 and therefore P2 due to lack of synchronization in the absence of the leader.
996 However, only a few of the tests could hit an execution with inconsistent servers using gRPC
997 adapter.

998 In conclusion, in the context of Ratis, the implementation of Algorithm 1 based on a lightweight
999 system instrumentation is quite effective and it hits a higher number of problematic executions
1000 in comparison to the coarse-grain implementation (based solely on the Ratis API without any
1001 instrumentation) or testing with a baseline randomized algorithm.

1002 7.3 Zookeeper

1004 We tested Apache Zookeeper, a strongly consistent distributed key-value store that relies on the ZAB
1005 (Zookeeper Atomic Broadcast) protocol, using a coarse-grained implementation of our sampling
1006 algorithm based exclusively on the API of the system, without additional instrumentation.⁹

1007 Our implementation enforces lockstep execution of *abstract phases*, which subsume a sequence
1008 of phases at the algorithmic level, starting from an event that causes the servers to start exchanging
1009 messages to a steady state. The length of an abstract phase is approximated in two ways. First, after
1010 starting a set of servers, a steady state is reached once the client-facing handlers detect that the
1011 servers have been started. During this time, the servers will have executed part of the ZAB protocol
1012 to agree on the most recent log of client requests. Second, after a client request, reaching a steady
1013 state is approximated with a 100ms timeout, empirically sufficient for the servers to commit the
1014 request. We use the system API to approximate points in execution where the system reaches a
1015 steady state and to inject faults (isolate servers) only at these points. This relaxed approach loses
1016 completeness, but it is easier to deploy since it does not require instrumentation. As we demonstrate
1017 in this section, it is sufficient for exposing interesting behaviors and bugs in Zookeeper.

1018 Our tool programmatically starts Zookeeper servers as threads, making them easier to manipulate
1019 than if they were separate processes. Each server is paired with a client-facing handler, which is
1020 also part of the Zookeeper API. The handler is used to detect a change in the server's state (is it up
1021 or down), and to initiate a client request (get or set a key-value pair).

1022 A test is parameterized by the number of servers n , a fault budget d , and a *test harness*. The
1023 test harness is determined by the client requests and the number of abstract phases, which are
1024 organized as a sequence of *steps*. A step can be either an empty step or a request step. An empty
1025 step, denoted as empty, consists of a single abstract phase that involves starting a set of servers
1026 and waiting for them to reach steady state. A request step consists of two abstract phases: the first
1027

1028 ⁹Source code is available at <https://github.com/fniksic/zootester>.

1030 one is like in the empty step, and the second one involves initiating a client request and waiting for
 1031 steady state, this time approximated with a 100ms timeout. We support two kinds of client requests:
 1032 a write request and a conditional write request. A write request for setting key k to value v on
 1033 server s is written as $s : k \leftarrow v$, and a conditional write request for setting key k_2 to value v_2 on
 1034 server s , provided that key k_1 is set to v_1 , is written as $s : k_1 = v_1 ? k_2 \leftarrow v_2$. In our tests we use
 1035 integer values. We identify requests and request steps and use the same notation for both.

1036 A test with n servers, a fault budget d , and a test harness with p steps is executed in the following
 1037 way. First there is an initial step in which all keys appearing in the harness are set to zero. Then
 1038 we use a version of Algorithm 1 to sample a random execution of the harness with d faults: we
 1039 distribute d faults over p steps, and additionally, if a step is a request step consisting of two abstract
 1040 phases, we randomly assign some of the faults to the second abstract phase in the step. At the
 1041 beginning of a step, we randomly choose a kernel of servers to start according to the number
 1042 of faults assigned to the first abstract phase in the step. If there is a second abstract phase, we
 1043 randomly choose servers to stop, again according to the number of faults assigned to the abstract
 1044 phase. At the end of a step, we stop all servers and proceed to the next step. Finally, once all steps
 1045 are executed, we start all servers and check that they are in the same final state, and that the final
 1046 state is allowed under some *sequentially consistent* execution of the requests.

1047 In our first experiment, we focus on exposing bug ZK-2832¹⁰, reported to occur in Zookeeper
 1048 3.4.9. The bug causes the servers to diverge; thus, we will refer to the bug as the *divergence* bug.
 1049 The reporter of the bug provided a test with the exact steps to deterministically reproduce the bug.
 1050 The steps involve three servers handling two client requests in presence of four faults. The client
 1051 requests set new values to two different keys. At the end the servers diverge: two servers disagree
 1052 on the value associated with one of the keys.

1053 Interestingly, the deterministic test provided by the bug's reporter fails to reproduce the bug in
 1054 releases of Zookeeper more recent than 3.4.9. Even though the bug report was still open at the time
 1055 of writing, it may seem that the bug has disappeared. Unfortunately, this is not the case: we were
 1056 able to reproduce the bug in Zookeeper 3.5.8, released in May 2020.

1057 Using our tool, we can represent the steps from the deterministic test as the following harness
 1058 involving servers s_0, s_1, s_2 and keys k_0, k_1 : $H_{\text{div}} = [s_1 : k_0 \leftarrow 101; \text{empty}; s_2 : k_1 \leftarrow 302]$. The exact
 1059 values assigned to the keys in the harness are not important, as long as they are distinct.

1060 We ran the harness with different values of the fault budget d . For each d from 0 to 9 we ran
 1061 1,000 executions and observed divergence in 0 to 5 executions per test. As a comparison, we ran a
 1062 baseline test in which we execute harness steps in 5-second intervals, while at the same time we
 1063 crash and restart servers in intervals randomly distributed according to Poisson distribution with
 1064 the mean of 2 seconds. In the baseline test, we observe divergence in 2 out of 1,000 executions. In
 1065 addition to the divergence bug, one of the executions of the baseline test shows what seems to be a
 1066 new issue: at the end, one of the clients is unable to connect to any of the servers. We believe this
 1067 cannot be correct behavior. We refer to this issue as *client dropped*. The left of Table 6 summarizes
 1068 the results. The last row in the table shows executions that were unsuccessful: occasionally a client
 1069 fails to read a value from a server. These executions are more likely to be a result of our tool not
 1070 being perfectly robust than of an actual issue with Zookeeper.

1071 In our next experiment, we experimented with our tool in the context of a random enumeration
 1072 of harnesses. To restrict the space of harnesses, we fixed the number of servers to 3, and the number
 1073 of keys to 2. In one experiment, we additionally fixed the number of requests $req = 2$, the total
 1074 number of steps $p = 3$, and the fault budget $d = 4$. In another experiment, we fixed the additional

1075
 1076
 1077 ¹⁰<https://issues.apache.org/jira/browse/ZOOKEEPER-2832>
 1078

Table 6. Testing Zookeeper. On the left, we list the number of Zookeeper executions with harness H_{div} exhibiting bugs listed in the first column, for varying d (we ran 1,000 executions for each value of d and for the baseline test). On the right, the number of Zookeeper executions exhibiting bugs listed in the first column for randomly sampled harnesses. For each of the two choices of parameters we randomly sampled 12 harnesses and ran 1,000 executions per harness.

d	0	1	2	3	4	5	6	7	8	9	baseline
divergence	0	0	0	2	2	0	5	4	3	0	2
client dropped	0	0	0	0	0	0	0	0	0	0	1
unsuccessful	0	0	0	1	0	1	0	0	0	1	8

	$req = 2, p = 3$ $d = 4$	$req = 4, p = 5$ $d = 6$
divergence	15	13
failure of SC	0	1
client dropped	0	7
unsuccessful	4	8

Table 7. Number of Zookeeper executions with harness H_{sc} exhibiting bugs listed in the first column, for varying d . We ran 1,000 executions for each value of d and for the baseline test.

d	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	baseline
divergence	0	0	0	2	1	2	2	5	6	0	0	0	0	0	0	0	7
client dropped	0	0	0	0	0	0	2	0	0	0	0	0	0	0	0	0	2
unsuccessful	0	0	0	0	0	0	0	0	0	0	1	2	0	0	0	0	19

parameters as $req = 4, p = 5, d = 6$. For each choice of parameters, we sampled 12 harnesses and ran 1,000 executions per harness.

The highlight of our findings is that, in addition to observing more divergence and dropped clients, we observe a new issue: in 1 out of 12,000 executions with $req = 4, p = 5, d = 6$, the servers converge to the same state, but this state is not allowed under sequential consistency. We refer to the issue as *failure of sequential consistency*. We have created a test that deterministically reproduces the violating execution and reported the issue as [ZK-3875](#).¹¹ The issue occurs in Zookeeper 3.5.8, but not in the more recent branch 3.6 of stable releases. At the time of writing it was still unclear which change in the 3.6 branch seems to resolve the issue. The results are summarized on the right of Table 6.

In our final experiment, we isolated the harness that yielded the execution exhibiting the failure of sequential consistency:

$$H_{sc} = [s_1 : k_1 = 0 ? k_1 \leftarrow 101; \text{empty}; s_0 : k_1 = 101 ? k_0 \leftarrow 200; \\ s_1 : k_1 = 0 ? k_1 \leftarrow 301; s_0 : k_1 = 0 ? k_0 \leftarrow 400]$$

In the incorrect execution, the final state on all servers is $\{k_0 = 200, k_1 = 301\}$. In the experiment, we wanted to see if we can detect failure of sequential consistency again, either by our sampling algorithm or by the baseline test. Therefore, we fixed the harness to H_{sc} and varied the fault budget d from 0 to 15. We observe divergence in 0 to 6 executions for our sampling algorithm, and in 7 executions for the baseline test. We observe clients dropped in 2 executions, both in our sampling algorithm and the baseline test. However, were not able to catch the failure of sequential consistency again, which shows that it is a rare bug. Table 7 summarizes the results.

7.4 Summary of Evaluation

Our experimental evaluation shows that our algorithm can detect new bugs in large scale systems as well as reproduce known bugs. In our tests, small values of d and values of k allowing a client request to be processed between recovery points could successfully detect bugs. This confirms our hypothesis that uniform executions with a small number of isolations are sufficient to find many bugs. We discovered new bugs in the recent versions of Zookeeper and Ratis. We inspected the

¹¹<https://issues.apache.org/jira/browse/ZOOKEEPER-3875>

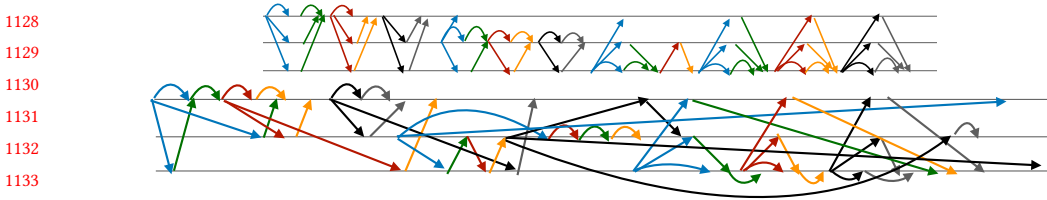


Fig. 6. A synchronous buggy trace sampled by our algorithm and a buggy trace sampled by PCT, both for Cassandra's Paxos bug.

buggy executions and have already reported some of them in the projects' issue tracker sites; some bugs in Ratis have already been fixed in the master branch of the project.

A limitation of some testing tools for distributed systems is the instrumentation burden. Our experimentation with different levels of precision on the identification of rounds and phases shows that sampling from uniform executions provides an effective approach for testing fault tolerance in general, even with coarse-grained instrumentation. All three versions of the implementation of our algorithm outperform a baseline random testing algorithm and can expose bugs in large scale systems.

Debuggability. We conclude by demonstrating that buggy executions detected by our algorithm can be easier to understand than the traces obtained by exploring all (asynchronous) executions. While the interpretability of the generated traces depends on the precision of the analysis of rounds in implementation, in general our algorithm produces execution traces that omit messages in a more structured way than reordering or dropping messages arbitrarily. Fig. 6 shows two buggy executions from Cassandra found by our algorithm and PCT [Kulahcioglu Ozkan et al. 2018], respectively. Our algorithm returns a synchronous execution trace which lists messages in the expected protocol order, making it more explicit which processes are isolated in each round. On the other hand, the programmer needs to follow the complicated message interleavings across phases to discover the delayed/dropped messages in the asynchronous trace.

8 RELATED WORK AND CONCLUSION

We have proposed a new testing methodology based on *communication-closure* as the starting point. Communication closure offers an elegant abstraction at the level of algorithm design, and our testing methodology uses the abstraction as a way to focus on a much smaller sample space of executions. For many common classes of distributed algorithms, the reduction remains complete. We have shown that exploring *uniform* executions with a small number of faults is sufficient to find bugs in production distributed systems like Cassandra, Zookeeper, or Ratis.

Using algorithmic insights into testing distributed systems to reduce the space of executions is a major point of departure from existing work in randomized or systematic testing of implementations of distributed systems. At the same time, our insight is orthogonal to the many reduction techniques already exploited in existing tools, such as depth bounding [Kulahcioglu Ozkan et al. 2018], partial order reduction [Kulahcioglu Ozkan et al. 2019; Yuan et al. 2018], or semantics-aware analyses [Leesatapornwongsa et al. 2014b; Lukman et al. 2019b].

Several execution prioritization techniques are designed for efficient analysis of concurrent software [Thomson et al. 2014]. Context bounding [Qadeer and Rehof 2005] or preemption-bounding [Musuvathi and Qadeer 2007] are designed for shared memory programs, defining a prioritization scheme based on multithreading concepts. While delay bounding [Emmi et al. 2011] or probabilistic prioritization in PCT [Burckhardt et al. 2010] are applicable to message

1177 passing systems, they consider the state space of message reorderings, hence parameterize the
1178 set of asynchronous executions. In this work, we provide an approach for exploring the set of
1179 synchronous executions of a distributed system. Note that we are not aware of any notion similar
1180 to communication closure that applies to shared-memory programs.

1181 While we address fault tolerance bugs due to message losses in this work, a related source of
1182 bugs is erroneous crash recovery of servers [Gao et al. 2018; Gunawi et al. 2015; Lu et al. 2019].
1183 Erroneous recovery causes the servers not to restart properly and leads to bugs in the system.
1184 Since message losses in network and server crashes are orthogonal sources of faults, producing
1185 executions with both kinds of faults may be promising for more extensive testing.

1186 Our work is inspired by the quest for an easier to understand subset of representative asyn-
1187 chronous executions, and simpler proofs of algorithms, which led to the communication closure
1188 property. Communication-closed layered systems [Charron-Bost and Schiper 2009; Chou and
1189 Gafni 1988; Gafni 1998; Moses and Rajsbaum 2002; Santoro and Widmayer 1989] capture both
1190 lossy synchronous and lossy asynchronous behaviors and solve consensus under the partial syn-
1191 chrony network assumption [Dwork et al. 1988]. They rely on easier to interpret synchronous
1192 lock-step executions and simpler proof arguments. For example an equivalence relation between
1193 asynchronous and communication closed executions is established for systems that solve consensus
1194 in [Chaouch-Saad et al. 2009; Elrad and Francez 1982; Moses and Rajsbaum 2002].

1195 Motivated by the impossibility of solving consensus over asynchronous faulty networks [?]
1196 synchronous abstractions offer an alternative view of distributed systems. They have been studied
1197 to simplify programming distributed, concurrent, and parallel systems, e.g., virtual synchrony [?],
1198 bulk programming [?], for designing theoretical solutions for consensus [Dwork et al. 1988], and
1199 to simplify reasoning about a system's traces [Elrad and Francez 1982]. Implementations of con-
1200 sensus protocols have been proposed for these synchronous programming paradigms, e.g., virtual
1201 synchrony [?] or PSync [Dragoi et al. 2016] (a programming paradigm based on communication-
1202 closure). However, in production asynchronous state machine replication systems are still to be
1203 understood if they have an implementation in synchronous programming models. In contrast,
1204 using communication-closure in testing increases the confidence we have in production systems
1205 without having to reimplement them. In [Damian et al. 2019] communication-closure is defined
1206 based on conditions on the sequential code independently of the specification of the systems and it
1207 is applied to semi-automatically prove correct several consensus protocols. The complexity and
1208 scale of the verified code is far from production system. No previous work studies the relation
1209 between communication closure and testing distributed systems.

1210 Finally, recent developments in verifying replicated state machine and consensus protocols [Chaud-
1211 huri et al. 2010; Hawblitzel et al. 2015; Padon et al. 2017; von Gleissenthall et al. 2019; Wilcox et al.
1212 2015] allow fully verified implementations to be developed. However, these verified implemen-
1213 tations lack the performance of production systems, are small scale implementations that have
1214 prototype clients and minimal deployment. Formalization is important, however bugs may still
1215 arise [Fonseca et al. 2017; Sutra 2019].

1216
1217

1218 ACKNOWLEDGMENTS

1219 Kulahcioglu Ozkan and Majumdar were supported in part by the Deutsche Forschungsgemeinschaft
1220 project 389792660 TRR 248 and by the European Research Council under the Grant Agreement
1221 610150 (ERC Synergy Grant ImPACT). This work was done mainly when Cezara Drăgoi was
1222 affiliated with INRIA supported by the French National Research Agency ANR project SAFTA
1223 (12744-ANR-17-CE25-0008-01).

1224
1225

REFERENCES

- 1226
1227 Apache. 2013. CASSANDRA-6023: CAS should distinguish promised and accepted ballots. Retrieved January 26, 2020 from
1228 <http://issues.apache.org/jira/browse/CASSANDRA-6023>
- 1229 Apache. 2020. *Apache Ratis*. Retrieved May 14, 2020 from <http://ratis.incubator.apache.org/>
- 1230 Sebastian Burckhardt, Pravesh Kothari, Madanlal Musuvathi, and Santosh Nagarakatte. 2010. A randomized scheduler with
1231 probabilistic guarantees of finding bugs. In *Proceedings of the 15th International Conference on Architectural Support for*
1232 *Programming Languages and Operating Systems, ASPLOS 2010, Pittsburgh, Pennsylvania, USA, March 13-17, 2010*, James C.
1233 Hoe and Vikram S. Adve (Eds.). ACM, 167–178. <https://doi.org/10.1145/1736020.1736040>
- 1234 Tushar Deepak Chandra, Robert Griesemer, and Joshua Redstone. 2007. Paxos made live: an engineering perspective. In
1235 *Proceedings of the Twenty-Sixth Annual ACM Symposium on Principles of Distributed Computing, PODC 2007, Portland,*
1236 *Oregon, USA, August 12-15, 2007*, Indranil Gupta and Roger Wattenhofer (Eds.). ACM, 398–407. <https://doi.org/10.1145/1281100.1281103>
- 1237 Mouna Chaouch-Saad, Bernadette Charron-Bost, and Stephan Merz. 2009. A Reduction Theorem for the Verification of
1238 Round-Based Distributed Algorithms. In *Reachability Problems, 3rd International Workshop, RP 2009, Palaiseau, France,*
1239 *September 23-25, 2009. Proceedings (Lecture Notes in Computer Science, Vol. 5797)*, Olivier Bournez and Igor Potapov (Eds.).
1240 Springer, 93–106. https://doi.org/10.1007/978-3-642-04420-5_10
- 1241 Charalambos A Charalambides. 2018. *Enumerative combinatorics*. Chapman and Hall/CRC.
- 1242 Bernadette Charron-Bost and André Schiper. 2009. The Heard-Of model: computing in distributed systems with benign
1243 faults. *Distributed Comput.* 22, 1 (2009), 49–71. <https://doi.org/10.1007/s00446-009-0084-6>
- 1244 Kaustuv Chaudhuri, Damien Doligez, Leslie Lamport, and Stephan Merz. 2010. Verifying Safety Properties with the TLA+
1245 Proof System. In *Automated Reasoning, 5th International Joint Conference, IJCAR 2010, Edinburgh, UK, July 16-19, 2010.*
1246 *Proceedings (Lecture Notes in Computer Science, Vol. 6173)*, Jürgen Giesl and Reiner Hähnle (Eds.). Springer, 142–148.
1247 https://doi.org/10.1007/978-3-642-14203-1_12
- 1248 Ching-Tsun Chou and Eli Gafni. 1988. Understanding and Verifying Distributed Algorithms Using Stratified Decomposition.
1249 In *Proceedings of the Seventh Annual ACM Symposium on Principles of Distributed Computing, Toronto, Ontario, Canada,*
1250 *August 15-17, 1988*, Danny Dolev (Ed.). ACM, 44–65. <https://doi.org/10.1145/62546.62556>
- 1251 Andrei Damian, Cezara Dragoi, Alexandru Militaru, and Josef Widder. 2019. Communication-Closed Asynchronous
1252 Protocols. In *Computer Aided Verification - 31st International Conference, CAV 2019, New York City, NY, USA, July 15-18,*
1253 *2019, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 11562)*, Isil Dillig and Serdar Tasiran (Eds.). Springer,
1254 344–363. https://doi.org/10.1007/978-3-030-25543-5_20
- 1255 Ankush Desai, Shaz Qadeer, and Sanjit A. Seshia. 2015. Systematic testing of asynchronous reactive systems. In *Proceedings*
1256 *of the 2015 10th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2015, Bergamo, Italy, August 30 - September*
1257 *4, 2015*, Elisabetta Di Nitto, Mark Harman, and Patrick Heymans (Eds.). ACM, 73–83. <https://doi.org/10.1145/2786805.2786861>
- 1258 Cezara Dragoi, Thomas A. Henzinger, and Damien Zufferey. 2016. PSync: a partially synchronous language for fault-tolerant
1259 distributed algorithms. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming*
1260 *Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, Rastislav Bodik and Rupak Majumdar (Eds.). ACM,
1261 400–415. <https://doi.org/10.1145/2837614.2837650>
- 1262 Cynthia Dwork, Nancy A. Lynch, and Larry J. Stockmeyer. 1988. Consensus in the presence of partial synchrony. *J. ACM*
1263 35, 2 (1988), 288–323. <https://doi.org/10.1145/42282.42283>
- 1264 Tzilla Elrad and Nissim Francez. 1982. Decomposition of Distributed Programs into Communication-Closed Layers. *Sci.*
1265 *Comput. Program.* 2, 3 (1982), 155–173. [https://doi.org/10.1016/0167-6423\(83\)90013-8](https://doi.org/10.1016/0167-6423(83)90013-8)
- 1266 Michael Emmi, Shaz Qadeer, and Zvonimir Rakamaric. 2011. Delay-bounded scheduling. In *Proceedings of the 38th ACM*
1267 *SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011,*
1268 Thomas Ball and Mooly Sagiv (Eds.). ACM, 411–422. <https://doi.org/10.1145/1926385.1926432>
- 1269 Alan Fekete and Nancy A. Lynch. 1990. The Need for Headers: An Impossibility Result for Communication over Unreliable
1270 Channels. In *CONCUR '90, Theories of Concurrency: Unification and Extension, Amsterdam, The Netherlands, August 27-30,*
1271 *1990, Proceedings (Lecture Notes in Computer Science, Vol. 458)*, Jos C. M. Baeten and Jan Willem Klop (Eds.). Springer,
1272 199–215. <https://doi.org/10.1007/BFb0039061>
- 1273 Pedro Fonseca, Kaiyuan Zhang, Xi Wang, and Arvind Krishnamurthy. 2017. An Empirical Study on the Correctness of
1274 Formally Verified Distributed Systems. In *Proceedings of the Twelfth European Conference on Computer Systems, EuroSys*
2017, Belgrade, Serbia, April 23-26, 2017. ACM, 328–343. <https://doi.org/10.1145/3064176.3064183>
- Eli Gafni. 1998. Round-by-Round Fault Detectors: Unifying Synchrony and Asynchrony (Extended Abstract). In *Proceedings*
of the Seventeenth Annual ACM Symposium on Principles of Distributed Computing, PODC '98, Puerto Vallarta, Mexico,
June 28 - July 2, 1998, Brian A. Coan and Yehuda Afek (Eds.). ACM, 143–152. <https://doi.org/10.1145/277697.277724>
- Yu Gao, Wensheng Dou, Feng Qin, Chushu Gao, Dong Wang, Jun Wei, Ruirui Huang, Li Zhou, and Yongming Wu. 2018. An
empirical study on crash recovery bugs in large-scale distributed systems. In *Proceedings of the 2018 ACM Joint Meeting*

- 1275 on *European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/SIGSOFT*
1276 *FSE 2018, Lake Buena Vista, FL, USA, November 04-09, 2018*. 539–550. <https://doi.org/10.1145/3236024.3236030>
- 1277 Haryadi S. Gunawi, Thanh Do, Agung Laksono, Mingzhe Hao, Tanakorn Leesatapornwongsa, Jeffrey F. Lukman, and Riza O.
1278 Suminto. 2015. What Bugs Live in the Cloud?: A Study of Issues in Scalable Distributed Systems. *login Usenix Mag.* 40, 4
(2015). <https://www.usenix.org/publications/login/aug15/gunawi>
- 1279 Chris Hawblitzel, Jon Howell, Manos Kapritsos, Jacob R. Lorch, Bryan Parno, Michael L. Roberts, Srinath T. V. Setty, and Brian
1280 Zill. 2015. IronFleet: proving practical distributed systems correct. In *Proceedings of the 25th Symposium on Operating*
1281 *Systems Principles, SOSP 2015, Monterey, CA, USA, October 4-7, 2015*, Ethan L. Miller and Steven Hand (Eds.). ACM, 1–17.
1282 <https://doi.org/10.1145/2815400.2815428>
- 1283 Patrick Hunt, Mahadev Konar, Flavio Paiva Junqueira, and Benjamin Reed. 2010. ZooKeeper: Wait-free Coordination for
1284 Internet-scale Systems. In *2010 USENIX Annual Technical Conference, Boston, MA, USA, June 23-25, 2010*.
- 1285 Yury Izrailevsky and Ariel Tseitlin. 2011. The Netflix Simian army. *The Netflix Tech Blog* (2011).
- 1286 Flavio Paiva Junqueira, Benjamin C. Reed, and Marco Serafini. 2011. Zab: High-performance broadcast for primary-backup
1287 systems. In *Proceedings of the 2011 IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2011,*
1288 *Hong Kong, China, June 27-30 2011*. IEEE Compute Society, 245–256. <https://doi.org/10.1109/DSN.2011.5958223>
- 1289 Victor Kac and Pokman Cheung. 2001. *Quantum calculus*. Springer Science & Business Media.
- 1290 Charles Edwin Killian, James W. Anderson, Ranjit Jhala, and Amin Vahdat. 2007. Life, Death, and the Critical Transition:
1291 Finding Liveness Bugs in Systems Code (Awarded Best Paper). In *4th Symposium on Networked Systems Design and*
1292 *Implementation (NSDI 2007), April 11-13, 2007, Cambridge, Massachusetts, USA, Proceedings*, Hari Balakrishnan and Peter
1293 Druschel (Eds.). USENIX. <http://www.usenix.org/events/nsdi07/tech/killian.html>
- 1294 Kyle Kingsbury. 2013–2018. *Jepsen*. Retrieved January 26, 2020 from <http://jepsen.io/>
- 1295 Burcu Kulahcioglu Ozkan, Rupak Majumdar, Filip Niksic, Mitra Tabaei Befrouei, and Georg Weissenbacher. 2018. Randomized
1296 testing of distributed systems with probabilistic guarantees. *PACMPL* 2, OOPSLA (2018), 160:1–160:28.
- 1297 Burcu Kulahcioglu Ozkan, Rupak Majumdar, and Simin Oraee. 2019. Trace aware random testing for distributed systems.
1298 *PACMPL* 3, OOPSLA (2019), 180:1–180:29.
- 1299 Avinash Lakshman and Prashant Malik. 2010a. Cassandra: a decentralized structured storage system. *Operating Systems*
1300 *Review* 44, 2 (2010), 35–40. <https://doi.org/10.1145/1773912.1773922>
- 1301 Avinash Lakshman and Prashant Malik. 2010b. Cassandra: a decentralized structured storage system. *ACM SIGOPS Operating*
1302 *Systems Review* 44, 2 (2010), 35–40.
- 1303 Leslie Lamport. 2005. *Generalized Consensus and Paxos*. Technical Report MSR-TR-2005-33. 60 pages. <https://www.microsoft.com/en-us/research/publication/generalized-consensus-and-paxos/>
- 1304 Tanakorn Leesatapornwongsa, Mingzhe Hao, Pallavi Joshi, Jeffrey F. Lukman, and Haryadi S. Gunawi. 2014a. SAMC:
1305 Semantic-Aware Model Checking for Fast Discovery of Deep Bugs in Cloud Systems. In *11th USENIX Symposium on*
1306 *Operating Systems Design and Implementation, OSDI '14, Broomfield, CO, USA, October 6-8, 2014*, Jason Flinn and Hank
1307 Levy (Eds.). USENIX Association, 399–414. <https://www.usenix.org/conference/osdi14/technical-sessions/presentation/leesatapornwongsa>
- 1308 Tanakorn Leesatapornwongsa, Mingzhe Hao, Pallavi Joshi, Jeffrey F. Lukman, and Haryadi S. Gunawi. 2014b. SAMC:
1309 Semantic-Aware Model Checking for Fast Discovery of Deep Bugs in Cloud Systems. In *11th USENIX Symposium on*
1310 *Operating Systems Design and Implementation, OSDI '14, Broomfield, CO, USA, October 6-8, 2014*. 399–414.
- 1311 Jie Lu, Chen Liu, Lian Li, Xiaobing Feng, Feng Tan, Jun Yang, and Liang You. 2019. CrashTuner: detecting crash-recovery
1312 bugs in cloud systems via meta-info analysis. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles,*
1313 *SOSP 2019, Huntsville, ON, Canada, October 27-30, 2019*. 114–130. <https://doi.org/10.1145/3341301.3359645>
- 1314 Jeffrey F. Lukman, Huan Ke, Cesar A. Stuardo, Riza O. Suminto, Daniar H. Kurniawan, Dikaimin Simon, Satria Priambada,
1315 Chen Tian, Feng Ye, Tanakorn Leesatapornwongsa, Aarti Gupta, Shan Lu, and Haryadi S. Gunawi. 2019a. FlyMC: Highly
1316 Scalable Testing of Complex Interleavings in Distributed Systems. In *Proceedings of the Fourteenth EuroSys Conference*
1317 *2019, Dresden, Germany, March 25-28, 2019*, George Candea, Robbert van Renesse, and Christof Fetzer (Eds.). ACM,
1318 20:1–20:16. <https://doi.org/10.1145/3302424.3303986>
- 1319 Jeffrey F. Lukman, Huan Ke, Cesar A. Stuardo, Riza O. Suminto, Daniar H. Kurniawan, Dikaimin Simon, Satria Priambada,
1320 Chen Tian, Feng Ye, Tanakorn Leesatapornwongsa, Aarti Gupta, Shan Lu, and Haryadi S. Gunawi. 2019b. FlyMC: Highly
1321 Scalable Testing of Complex Interleavings in Distributed Systems. In *Proceedings of the Fourteenth EuroSys Conference*
1322 *2019, Dresden, Germany, March 25-28, 2019*. 20:1–20:16.
- 1323 Nancy A. Lynch. 1996. *Distributed Algorithms*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- 1324 Iulian Moraru, David G. Andersen, and Michael Kaminsky. 2013. There is more consensus in Egalitarian parliaments.
1325 In *ACM SIGOPS 24th Symposium on Operating Systems Principles, SOSP '13, Farmington, PA, USA, November 3-6, 2013*,
1326 Michael Kaminsky and Mike Dahlin (Eds.). ACM, 358–372. <https://doi.org/10.1145/2517349.2517350>
- 1327 Yoram Moses and Sergio Rajsbaum. 2002. A Layered Analysis of Consensus. *SIAM J. Comput.* 31, 4 (2002), 989–1021.
1328 <https://doi.org/10.1137/S0097539799364006>

- 1324 Madanlal Musuvathi and Shaz Qadeer. 2007. Iterative context bounding for systematic testing of multithreaded programs.
1325 In *Proceedings of the ACM SIGPLAN 2007 Conference on Programming Language Design and Implementation, San Diego,*
1326 *California, USA, June 10-13, 2007*, Jeanne Ferrante and Kathryn S. McKinley (Eds.). ACM, 446–455. <https://doi.org/10.1145/1250734.1250785>
- 1327 Brian M. Oki and Barbara Liskov. 1988. Viewstamped Replication: A General Primary Copy. In *Proceedings of the Seventh*
1328 *Annual ACM Symposium on Principles of Distributed Computing, Toronto, Ontario, Canada, August 15-17, 1988*, Danny
1329 Dolev (Ed.). ACM, 8–17. <https://doi.org/10.1145/62546.62549>
- 1330 Diego Ongaro and John K. Ousterhout. 2014. In Search of an Understandable Consensus Algorithm. In *2014 USENIX Annual*
1331 *Technical Conference, USENIX ATC '14, Philadelphia, PA, USA, June 19-20, 2014*, Garth Gibson and Nickolai Zeldovich
1332 (Eds.). USENIX Association, 305–319. <https://www.usenix.org/conference/atc14/technical-sessions/presentation/ongaro>
- 1333 Burcu Kulahcioglu Ozkan, Rupak Majumdar, Filip Nicksic, Mitra Tabaei Befrouei, and Georg Weissenbacher. 2018. Randomized
1334 testing of distributed systems with probabilistic guarantees. *Proc. ACM Program. Lang.* 2, OOPSLA (2018), 160:1–160:28.
<https://doi.org/10.1145/3276530>
- 1335 Oded Padon, Giuliano Losa, Mooly Sagiv, and Sharon Shoham. 2017. Paxos made EPR: decidable reasoning about distributed
1336 protocols. *Proc. ACM Program. Lang.* 1, OOPSLA (2017), 108:1–108:31. <https://doi.org/10.1145/3140568>
- 1337 Shaz Qadeer and Jakob Rehof. 2005. Context-Bounded Model Checking of Concurrent Software. In *Tools and Algorithms for*
1338 *the Construction and Analysis of Systems, 11th International Conference, TACAS 2005, Held as Part of the Joint European*
1339 *Conferences on Theory and Practice of Software, ETAPS 2005, Edinburgh, UK, April 4-8, 2005, Proceedings (Lecture Notes in*
1340 *Computer Science, Vol. 3440)*, Nicolas Halbwachs and Lenore D. Zuck (Eds.). Springer, 93–107. https://doi.org/10.1007/978-3-540-31980-1_7
- 1341 Nicola Santoro and Peter Widmayer. 1989. Time is Not a Healer. In *STACS 89, 6th Annual Symposium on Theoretical Aspects*
1342 *of Computer Science, Paderborn, FRG, February 16-18, 1989, Proceedings (Lecture Notes in Computer Science, Vol. 349)*,
1343 Burkhard Monien and Robert Cori (Eds.). Springer, 304–313. <https://doi.org/10.1007/BFb0028994>
- 1344 Pierre Sutra. 2019. On the correctness of Egalitarian Paxos. *CoRR* abs/1906.10917 (2019). [arXiv:1906.10917](http://arxiv.org/abs/1906.10917) <http://arxiv.org/abs/1906.10917>
- 1345 Paul Thomson, Alastair F. Donaldson, and Adam Betts. 2014. Concurrency testing using schedule bounding: an empirical
1346 study. In *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPoPP '14, Orlando, FL, USA,*
1347 *February 15-19, 2014*, José E. Moreira and James R. Larus (Eds.). ACM, 15–28. <https://doi.org/10.1145/2555243.2555260>
- 1348 Klaus von Gleissenthall, Rami Gökhan Kici, Alexander Bakst, Deian Stefan, and Ranjit Jhala. 2019. Pretend synchrony:
1349 synchronous verification of asynchronous distributed programs. *Proc. ACM Program. Lang.* 3, POPL (2019), 59:1–59:30.
<https://doi.org/10.1145/3290372>
- 1350 James R. Wilcox, Doug Woos, Pavel Panchekha, Zachary Tatlock, Xi Wang, Michael D. Ernst, and Thomas E. Anderson.
1351 2015. Verdi: a framework for implementing and formally verifying distributed systems. In *Proceedings of the 36th ACM*
1352 *SIGPLAN Conference on Programming Language Design and Implementation, Portland, OR, USA, June 15-17, 2015*, David
1353 Grove and Steve Blackburn (Eds.). ACM, 357–368. <https://doi.org/10.1145/2737924.2737958>
- 1354 Xinhao Yuan, Junfeng Yang, and Ronghui Gu. 2018. Partial Order Aware Concurrency Sampling. In *Computer Aided*
1355 *Verification - 30th International Conference, CAV 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford,*
1356 *UK, July 14-17, 2018, Proceedings, Part II* 317–335.

1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372