

EFFICIENT COUNTING WITH OPTIMAL RESILIENCE*

CHRISTOPH LENZEN[†], JOEL RYBICKI[‡], AND JUKKA SUOMELA[§]

Abstract. Consider a complete communication network of n nodes, where the nodes receive a common clock pulse. We study the synchronous c -counting problem: given any starting state and up to f faulty nodes with arbitrary behaviour, the task is to eventually have all correct nodes labeling the pulses with increasing values modulo c in agreement. Thus, we are considering algorithms that are self-stabilising despite Byzantine failures. In this work, we give new algorithms for the synchronous counting problem that (1) are deterministic, (2) have optimal resilience, (3) have a linear stabilisation time in f (asymptotically optimal), (4) use a small number of states, and consequently, (5) communicate a small number of bits per round. Prior algorithms either resort to randomisation, use a large number of states and need high communication bandwidth, or have suboptimal resilience. In particular, we achieve an *exponential* improvement in both state complexity and message size for deterministic algorithms. Moreover, we present two complementary approaches for reducing the number of bits communicated during and after stabilisation.

Key words. self-stabilisation, Byzantine fault-tolerance

AMS subject classifications. 68M14, 68M15, 68Q25, 68W15

1. Introduction. In this work, we design space- and communication-efficient, self-stabilising, Byzantine fault-tolerant algorithms for the *synchronous counting problem*. We are given a complete communication network on n nodes, with arbitrary initial states. There are up to f faulty nodes. The task is to synchronise the nodes so that all non-faulty nodes will count rounds modulo c in agreement. For example, here is a possible execution for $n = 4$ nodes, $f = 1$ faulty node, and counting modulo $c = 3$; the execution stabilises after $t = 5$ rounds:

	Stabilisation	Counting
Node 1:	2 2 0 2 0	0 1 2 0 1 2 ...
Node 2:	0 2 0 1 0	0 1 2 0 1 2 ...
Node 3:	<i>faulty node, arbitrary behaviour</i> ...	
Node 4:	0 0 2 0 2	0 1 2 0 1 2 ...

Synchronous counting is a coordination primitive that can be used e.g. in large integrated circuits to synchronise subsystems to easily implement *mutual exclusion* and *time division multiple access* in a fault-tolerant manner. Note that in this context, it is natural to assume that a synchronous clock signal is available, but the clocking system usually does not provide explicit round numbers. Solving synchronous counting thus yields highly dependable round counters for subcircuits.

If we neglect communication, counting and consensus are essentially equivalent [13–15]. In particular, many lower bounds on (binary) consensus directly apply to the counting problem [16, 20, 27]. However, the known generic reduction of counting to consensus incurs a factor- f overhead in space and message size. In this work, we

*Submitted to the editors DATE. This paper is an extended and revised version of two preliminary conference reports [24, 26] that appeared in the Proceedings of the 34th Annual ACM Symposium on Principles of Distributed Computing (PODC 2015) and in the Proceedings of the 29th International Symposium on Distributed Computing (DISC 2015).

[†]Department of Algorithms and Complexity, Max Planck Institute for Informatics, Saarland Informatics Campus (clenzen@mpi-inf.mpg.de).

[‡]Current affiliation: Department of Biosciences, University of Helsinki (joel.rybicki@helsinki.fi).

[§]Helsinki Institute for Information Technology HIIT, Department of Computer Science, Aalto University (jukka.suomela@aalto.fi).

35 present techniques that reduce the number of bits nodes broadcast in each round to
 36 $O(\log^2 f + \log c)$.

37 **1.1. Contributions.** Our contributions constitute of two parts. First, we give
 38 novel space-efficient deterministic algorithms for synchronous counting with optimal
 39 resilience and fast stabilisation time. Second, we show how to extend these algorithms
 40 in a way that reduces the number of communicated bits *during* and *after* stabilisation.

41 *Space-efficient counting algorithms.* In this work, we take the following approach
 42 for devising *communication-efficient* counting algorithms: we first design *space-efficient*
 43 algorithms, that is, algorithms in which each node stores only a few bits between
 44 consecutive rounds. Space-efficient algorithms are particularly attractive from the
 45 perspective of fault-tolerant systems: if we can keep the number of state bits small,
 46 we can also reduce the overall complexity of the system, which in turn makes it easier
 47 to use highly reliable components for an implementation.

48 Once we have algorithms that only need a small number of bits to encode the
 49 local state of a node, we also get algorithms that use small messages: the nodes can
 50 simply broadcast their entire state to everyone. Our main result is summarised in the
 51 following theorem; here *f-resilient* means that we can tolerate up to f faulty nodes:

52 **THEOREM 1.1.** *For any integers $c, n > 1$ and $f < n/3$, there exists a deterministic*
 53 *f -resilient synchronous c -counter that runs on n nodes, stabilises in $O(f)$ rounds, and*
 54 *uses $O(\log^2 f + \log c)$ bits to encode the state of a node.*

55 Our main technical contribution is a recursive construction that shows how to
 56 “amplify” the resilience of a synchronous counting algorithm. Given a synchronous
 57 counter for some values of n and f , we will show how to design synchronous counters
 58 for larger values of n and f , with a very small increase in time and state complexity.
 59 This has two direct applications:

- 60 1. From a practical perspective, we can apply existing computer-designed algo-
 61 rithms (e.g. $n = 4$ and $f = 1$) as a building block in order to design efficient
 62 deterministic algorithms for a moderate number of nodes (e.g., $n = 36$ and
 63 $f = 7$).
- 64 2. From a theoretical perspective, we can design deterministic algorithms for
 65 synchronous counting for any n and any $f < n/3$, with a stabilisation time of
 66 $\Theta(f)$, and with only $O(\log^2 f)$ bits of state per node.

67 The state complexity and message size is an *exponential* improvement over prior work,
 68 and the stabilisation time is asymptotically optimal for deterministic algorithms [20].

69 *Reducing communication after stabilisation.* In our deterministic algorithms, each
 70 node only needs to store a few number of bits between consecutive rounds, and thus,
 71 a node can e.g. afford to broadcast its entire state to all other nodes in each round.
 72 Moreover, we present a technique to reduce the number of communicated bits further.

73 We give a deterministic construction in which *after* stabilisation each node broad-
 74 casts $O(1 + B \log B)$ bits every κ rounds, where $B = O(\log c / \log \kappa)$, for an essentially
 75 unconstrained choice of κ , at the expense of additively increasing the stabilisation
 76 time by $O(\kappa)$. In particular, for the special case of optimal resilience and polynomial
 77 counter size, we obtain the following result.

78 **COROLLARY 1.2.** *For any $n > 1$ and $c = n^{O(1)}$ that is an integer multiple of*
 79 *n , there exists a synchronous c -counter that runs on n nodes, has optimal resilience*
 80 *$f = \lfloor (n - 1)/3 \rfloor$, stabilises in $\Theta(n)$ rounds, requires $O(\log^2 n)$ bits to encode the state*
 81 *of a node, and for which after stabilisation correct nodes broadcast asymptotically*
 82 *optimal $O(1)$ bits per $\Theta(n)$ rounds.*

83 We remark that in the above result we simply reduce the frequency of com-
 84 munication and the size of messages instead of e.g. bounding the number of nodes
 85 communicating in any given round (known as broadcast efficiency) [28]. In our
 86 work, we exploit synchrony after stabilisation to schedule communication, and thus,
 87 our approach is to be contrasted with attempting to reduce the total number of
 88 communication partners or communicating nodes after stabilisation [9, 10, 28].

89 *Reducing the number of messages.* To substantiate the conjecture that finding
 90 algorithms with small state complexity may lead to highly communication-efficient
 91 solutions, we proceed to consider a slightly stronger synchronous *pulling model*. In
 92 this model, a node may send a request to another node and receive a response in
 93 a single round, based on the state of the responding node at the beginning of the
 94 round. The cost for the exchange is then attributed to the pulling node; in a circuit,
 95 this translates to each node being assigned an energy budget that it uses to “pay”
 96 for the communication it triggers. In this model, it is straightforward to combine
 97 our recursive construction used in [Theorem 1.1](#) with random sampling to obtain the
 98 following results:

- 99 1. We can achieve the same asymptotic running time and state complexity as
 100 the deterministic algorithm from [Theorem 1.1](#) with each node pulling only
 101 polylog n messages in each round. The price is that the resulting algorithm
 102 retains a probability of $n^{-\text{polylog } n}$ to fail in each round even after stabilisation
 103 and that the resilience is $f < n/(3 + \gamma)$ for any constant $\gamma > 0$.
- 104 2. If the failing nodes are chosen independently of the algorithm, we can fix the
 105 random choices. This results in a pseudorandom algorithm which stabilises
 106 with a probability of $1 - n^{-\text{polylog } n}$ and in this case keeps counting correctly.

107 **1.2. Our Approach.** Most prior deterministic algorithms for synchronous count-
 108 ing and closely-related problems utilise consensus protocols [14, 22]. Indeed, if we ignore
 109 space and communication, reductions exist both ways showing that the problems are
 110 more or less equivalent [12]; see [Section 2](#) for further discussion on prior work.

111 However, to construct fast space- and communication-efficient counters, we are
 112 facing a chicken-and-egg problem:

- 113 • **From counters to consensus:** If the correct nodes could agree on a counter,
 114 they could jointly run a *single* instance of synchronous consensus.
- 115 • **From consensus to counters:** If the nodes could run a consensus algorithm,
 116 they could agree on a counter.

117 A key step to circumvent this obstacle is the following observation:

- 118 • **From unreliable counters to consensus:** If the correct nodes can agree
 119 on a counter *at least for a while*, they can jointly run a single instance of
 120 consensus.
- 121 • **From consensus to reliable counters:** Consensus can be then used to
 122 facilitate agreement on the output counter, and it is possible to maintain
 123 agreement even if the underlying unreliable counters fail later on.

124 The task of constructing counters that are correct only once in a while is easier; in
 125 particular, it does not require that we solve consensus in the process. As our main
 126 technical result, we show how to “amplify” the resilience f , at a cost of losing some
 127 guarantees:

- 128 • **Input:** Two counters with a small f ; guaranteed to work permanently after
 129 stabilisation.
- 130 • **Output:** A counter with a large f ; guaranteed to work only once in a while.

131 This can be then used to jointly run a single instance of consensus and stabilise the

TABLE 1

Summary of counting algorithms for the case $c = 2$. For randomised algorithms, we list the expected stabilisation time. (*) The solution from [4] relies on a shared coin—details vary, but all known shared coins with large resilience require large states and messages.

resilience	stabilisation time	state bits	deterministic	reference
$f < n/3$	$O(1)$	$n^{O(1)}$	no	[4] (*)
$f < n/3$	$O(f)$	$O(f \log f)$	yes	[14]
$f < n/3$	$2^{2(n-f)}$	2	no	[17, 18]
$f < n/3$	$\min\{2^{2f+2} + 1, 2^{O(f^2/n)}\}$	1	no	[13]
$f = 1, n \geq 4$	7	2	yes	[13]
$f = n^{1-o(1)}$	$O(f)$	$O(\log^2 f / \log \log f)$	yes	[26]
$f < n/3$	$O(f)$	$O(\log^2 f)$	yes	this work

132 output. We show how to obtain such a counter based on simple local consistency
133 checks, timeouts, and threshold voting.

134 In the end, a recursive application of this scheme allows us to build space-efficient
135 counting algorithms for any n with optimal resilience. At each level of recursion, we
136 only need to run a single instance of consensus. As there will be $O(\log f)$ levels of
137 recursion, in total each node participates in only $O(\log f)$ consensus instances.

138 **1.3. Structure.** Section 2 reviews prior work on impossibility results and count-
139 ing algorithms. Section 3 provides a formal description of the basic model of computa-
140 tion and the synchronous counting problem. Section 4 gives the main technical result
141 on resilience boosting, and Section 5 applies it to construct fast and communication-
142 efficient algorithms. Section 6 shows how to reduce the number of bits communicated
143 during and after stabilisation. Section 7 discusses the pulling model and randomised
144 sampling.

145 **2. Related Work.** In this section, we first overview impossibility results related
146 to counting, and then discuss both deterministic and randomised algorithms for the
147 counting problem.

148 *Impossibility results.* As mentioned, counting is closely related to consensus as
149 reductions exist both ways [12]: consensus can be solved in time $O(T)$ tolerating f
150 faults if and only if counting can be solved in time $O(T)$ tolerating f faults.

151 With this equivalence in mind, several impossibility results for consensus directly
152 hold for counting as well. First, consensus cannot be solved in the presence of $n/3$ or
153 more Byzantine failures [27]. Second, any deterministic f -resilient consensus algorithm
154 needs to run for at least $f + 1$ communication rounds [20]. Third, it is known that the
155 connectivity of the communication network must be at least $2f + 1$ [11]. Finally, any
156 consensus algorithm needs to communicate at least $\Omega(nf)$ bits in total [16].

157 In terms of communication complexity, no better bound than $\Omega(nf)$ on the
158 total number of communicated bits is known. While non-trivial for consensus, this
159 bound turns out to be trivial for deterministic counting algorithms: a self-stabilising
160 algorithm needs to verify its output, and to do that, each of the n nodes needs to
161 receive information from at least $f + 1 = \Omega(f)$ other nodes to be certain that some
162 other non-faulty node has the same output value. Similarly, no non-constant lower
163 bounds on the number of *state bits* nodes are known; however, a non-trivial constant
164 lower bound for the case $f = 1$ is known [13].

165 *Prior algorithms.* There are several algorithms to the synchronous counting
 166 problem, with different trade-offs in terms of resilience, stabilisation time, space
 167 complexity, communication complexity, and the use of random bits. For a brief
 168 summary, see [Table 1](#).

169 Designing space-efficient *randomised* algorithms for synchronous counting is fairly
 170 straightforward [[13, 17, 18](#)]: for example, the nodes can simply choose random states
 171 until a clear majority of nodes has the same state, after which they start to follow
 172 the majority. Likewise, given a shared coin, one can quickly reach agreement by
 173 defaulting to the coin whenever no clear majority is observed [[4](#)]. However, existing
 174 highly-resilient shared coins are very inefficient in terms of communication or need
 175 additional assumptions, such as private communication links between correct nodes.
 176 Less resilient shared coins are easier to obtain: resilience $\Theta(\sqrt{n})$ is achieved by each
 177 node announcing the outcome of an independent coin flip and locally outputting the
 178 (observed) majority value. In addition, $\Omega(n/\log^2 n)$ -resilient Boolean functions give
 179 fast communication-efficient coins [[1](#)]. Designing quickly stabilising algorithms that
 180 are both communication-efficient and space-efficient has turned out to be a challenging
 181 task [[13–15](#)], and it remains open to what extent randomisation can help in designing
 182 such algorithms.

183 In the case of *deterministic* algorithms, algorithm synthesis has been used for
 184 computer-aided design of optimal algorithms with resilience $f = 1$, but the approach
 185 does not scale due to the extremely fast-growing space of possible algorithms [[13](#)]. In
 186 general, many fast-stabilising algorithms build on a connection between Byzantine
 187 consensus and synchronous counting, but require a large number of states per node [[14](#)]
 188 due to, e.g., running a large number of consensus instances in parallel. Recently, in
 189 one of the preliminary conference reports [[26](#)] this paper is based on, we outlined a
 190 recursive approach where each node needs to participate in only $O(\log f / \log \log f)$
 191 parallel instances of consensus. However, this approach resulted in suboptimal resilience
 192 of $f = n^{1-o(1)}$.

193 Finally, we note that while counting algorithms are usually designed for the case
 194 of a fully-connected communication topology, the algorithms can be extended to use
 195 in a variety of other graph classes with high enough connectivity [[13](#)].

196 *Related problems.* Boczkowski et al. [[7](#)] study the synchronous c -counting problem
 197 (under the name self-stabilising clock synchronisation) with $O(\sqrt{n})$ Byzantine faults
 198 in a stochastic communication setting that resembles the pulling model we consider
 199 in [Section 7](#). However, their communication model is much more restricted: in every
 200 round, each node interacts with at most constantly many nodes which are chosen
 201 uniformly at random. Moreover, nodes only exchange messages of size $O(\log c)$ bits.

202 Without Byzantine (or other types of permanent) faults, self-stabilising counters
 203 and digital clocks have been studied as the *self-stabilising unison problem* [[2, 8, 21](#)].
 204 However, unlike in the fully-connected setting considered in this work, the underlying
 205 communication topology in the unison problem is typically assumed to be an arbitrary
 206 graph. In our model, in absence of permanent faults the problem becomes trivial, as
 207 nodes may simply reproduce the clock of a predetermined leader. The unison problem
 208 has also been studied in asynchronous models [[8, 19](#)]; this variant is also known as
 209 self-stabilising synchronisers [[3](#)].

210 **3. Preliminaries.** In this section, we define the model of computation and the
 211 counting problem.

212 **3.1. Model of Computation.** We consider a fully-connected synchronous message-
 213 passing network. That is, our distributed system consists of a network of n nodes,

214 where each node is a state machine and has communication links to all other nodes in
 215 the network. All nodes have a unique identifier from the set $[n] = \{0, 1, \dots, n - 1\}$.
 216 The computation proceeds in synchronous communication rounds. In each round, all
 217 nodes perform the following in a lock-step fashion:

- 218 1. *broadcast* a single message to all nodes,
- 219 2. *receive* messages from all nodes, and
- 220 3. *update* the local state.

221 We assume that the initial state of each node is arbitrary and there are up to
 222 f Byzantine nodes. A Byzantine node may have arbitrary behaviour, that is, it
 223 can deviate from the protocol in any manner. In particular, the Byzantine nodes
 224 can collude together in an adversarial manner and a single Byzantine node can send
 225 *different* messages to different correct nodes.

226 *Algorithms.* Formally, we define an algorithm as a tuple $\mathbf{A} = \langle X, g, p \rangle$, where
 227 X is the set of all states any node can have, $g: [n] \times X^n \rightarrow X$ is the *state transition*
 228 *function*, and $p: [n] \times X \rightarrow [c]$ is the *output function*. At each round when node v
 229 receives a vector $\mathbf{x} = \langle x_0, \dots, x_{n-1} \rangle \in X^n$ of messages, node v updates its state to
 230 $g(v, \mathbf{x})$ and outputs $p(v, x_v)$. As we consider c -counting algorithms, the set of output
 231 values is the set $[c] = \{0, 1, \dots, c - 1\}$ of counter values.

232 The tuples passed to the state transition function g are ordered according to the
 233 node identifiers. Put otherwise, the nodes can identify the sender of a message—this
 234 is frequently referred to as source authentication. Moreover, in the basic model, we
 235 assume that all nodes simply broadcast their state to all other nodes. Thus, the set of
 236 messages is the same as the set of possible states.

237 *Executions.* For any set of $\mathcal{F} \subseteq [n]$ of faulty nodes, we define a projection $\pi_{\mathcal{F}}$ that
 238 maps any state vector $\mathbf{x} \in X^n$ to a *configuration* $\pi_{\mathcal{F}}(\mathbf{x}) = \mathbf{e}$, where $e_v = *$ if $v \in \mathcal{F}$
 239 and $e_v = x_v$ otherwise. That is, the values given by Byzantine nodes are ignored
 240 and a configuration consists of only the states of correct nodes. A configuration \mathbf{d}
 241 is *reachable* from configuration \mathbf{e} if for every correct node $v \notin \mathcal{F}$ there exists some
 242 $\mathbf{x} \in X^n$ satisfying $\pi_{\mathcal{F}}(\mathbf{x}) = \mathbf{e}$ and $g(v, \mathbf{x}) = d_v$. An *execution* of an algorithm \mathbf{A}
 243 is an infinite sequence of configurations $\xi = \langle \mathbf{e}_0, \mathbf{e}_1, \dots \rangle$ where configuration \mathbf{e}_{r+1} is
 244 reachable from configuration \mathbf{e}_r .

245 **3.2. Synchronous Counters and Complexity Measures.** We say that an
 246 execution $\xi = \langle \mathbf{e}_0, \mathbf{e}_1, \dots \rangle$ of a counting algorithm \mathbf{A} *stabilises* in time T if there is
 247 some $k \in [c]$ such that for every correct node $v \in [n] \setminus \mathcal{F}$ it holds that

$$248 \quad p(v, e_{T+r,v}) = r - k \bmod c \quad \text{for all } r \geq 0,$$

249 where $e_{T+r,v} \in X$ is the state of node v in round $T + r$.

250 An algorithm \mathbf{A} is said to be a *synchronous c -counter with resilience f* that
 251 stabilises in time T , if for every $\mathcal{F} \subseteq [n]$, $|\mathcal{F}| \leq f$, all executions of algorithm \mathbf{A}
 252 stabilise within T rounds. In this case, we say that the *stabilisation time* $T(\mathbf{A})$ of
 253 \mathbf{A} is the minimal such T that all executions of \mathbf{A} stabilise in T rounds. The *state*
 254 *complexity* of \mathbf{A} is $S(\mathbf{A}) = \lceil \log |X| \rceil$, that is, the number of bits required to encode
 255 the state of a node between subsequent rounds. For brevity, we will often refer to
 256 $\mathcal{A}(n, f, c)$ as the family of synchronous c -counters over n nodes with resilience f . For
 257 example, $\mathbf{A} \in \mathcal{A}(4, 1, 2)$ denotes a synchronous 2-counter (i.e. a binary counter) over 4
 258 nodes tolerating one failure.

259 **4. Boosting Resilience.** In this section, we show how to use existing “small”
 260 synchronous counters to construct new “large” synchronous counters with a higher

261 resilience f and a larger number of nodes n ; we call this *resilience boosting*. We will
 262 then apply the idea recursively, with trivial counters as a base case.

263 **4.1. Road Map.** The high-level idea of resilience boosting is as follows. We
 264 start with counters that have a low resilience f' and use a small number of nodes n' .
 265 We use such counters to construct a new “weak” counter that has a higher resilience
 266 $f > f'$ and a large number of nodes $n > n'$ but only needs to behave correctly *once*
 267 *in a while* for sufficiently long. Once such a weak counter exists, it can be used to
 268 provide consistent round numbers for long enough to execute a *single* instance of a
 269 high-resilience consensus protocol. This can be used to reach agreement on the output
 270 counter.

271 *Constructing the Weak Counter.* For clarity, we will use here the term *strong*
 272 *counter* to refer to a self-stabilising fault-tolerant counter in the usual sense, and the
 273 term *weak counter* to refer to a counter that behaves correctly once in a while. We
 274 assume that f' -resilient strong counters for all $f' < f$ already exist, and we show how
 275 to construct an f -resilient weak counter that behaves correctly for at least τ rounds.
 276 Put slightly more formally, a weak τ -counter satisfies the following property: there
 277 exists a round r such that for all correct nodes $v, w \in V \setminus F$ satisfy

- 278 • $d(v, r) = d(w, r)$ and
- 279 • $d(v, r') = d(v, r' - 1) + 1 \pmod{\tau}$ for all $r' \in \{r + 1, \dots, r + \tau - 1\}$,

280 where $d(v, r)$ denotes the value of the weak counter at node v in round r . That is,
 281 eventually there will be τ consecutive rounds during which the (weak) counter values
 282 agree and are incremented by one modulo τ every round. However, after these τ
 283 rounds, the counters can behave arbitrarily.

284 Let $f_0 + f_1 + 1 = f$ and $n_0 + n_1 = n$. We take an f_0 -resilient strong 2τ -counter
 285 \mathbf{A}_0 with n_0 nodes and an f_1 -resilient strong 6τ -counter \mathbf{A}_1 with n_1 nodes, and use
 286 them to construct an f -resilient weak counter with n nodes.

287 We partition n nodes in disjoint “blocks”: block 0 runs \mathbf{A}_0 with n_0 nodes and
 288 block 1 runs \mathbf{A}_1 with n_1 nodes. At least one of the algorithms will eventually stabilise
 289 and count correctly. The key challenge is making sure that eventually all correct nodes
 290 (in both blocks!) will follow the same correct counter, at least for τ rounds.

291 To this end, each block maintains a *leader pointer*. The leader pointers are changed
 292 regularly: block 0 changes its leader pointer every τ rounds, and block 1 changes its
 293 leader pointer every 3τ rounds. If the leader pointers behave correctly, there will be
 294 regularly periods of τ rounds such that both of the leader pointers point to the same
 295 correct block.

296 If we had reliable counters, block i could simply use the current value of counter
 297 \mathbf{A}_i to determine the current value of its leader pointer. However, one of the counters
 298 might misbehave. As a remedy, each node v of block i checks if the output variable of
 299 counter \mathbf{A}_i increases by 1 in each round. If not, it will consider \mathbf{A}_i faulty for $\Theta(\tau)$
 300 rounds. The final output of a node is determined as follows:

- 301 • If node v in block i thinks that \mathbf{A}_i is faulty, it outputs the current value of
 302 counter \mathbf{A}_{1-i} .
- 303 • Otherwise, it uses the current value of \mathbf{A}_i to construct the leader pointer
 304 $\ell \in \{0, 1\}$, and it outputs the current value of counter \mathbf{A}_ℓ .

305 Note that the counter \mathbf{A}_i might seem to be behaving in a faulty manner if there has
 306 not been enough time for \mathbf{A}_i to stabilise. However, each node v of block i will consider
 307 a block to be faulty at most $\Theta(\tau)$ rounds before checking again whether the output
 308 of \mathbf{A}_i behaves consistently. Thus, if \mathbf{A}_i eventually stabilises, then eventually node v
 309 stops considering \mathbf{A}_i as faulty for good (at least until the next transient failure).

310 The above consistency check almost cuts it—except that two nodes $w \neq v$ of block
 311 i may have different opinions on the current value of \mathbf{A}_i . We clear this final hurdle
 312 by enlisting the help of *all* nodes for a majority vote on what the current value of \mathbf{A}_i
 313 actually is. Essentially, we use threshold voting; this way all nodes that think that \mathbf{A}_i
 314 behaves correctly will agree on a globally unique counter value α_i for \mathbf{A}_i .

315 If, for example, block 0 contains at most f_0 faulty nodes, all of this eventually
 316 entails the following:

- 317 1. Counter \mathbf{A}_0 stabilises, counts correctly, and all correct nodes agree on its
 318 counter value α_0 .
- 319 2. All correct nodes of block 0 think that block 0 is counting correctly. They
 320 use α_0 to derive the value of the leader pointer. Once in 2τ rounds, when the
 321 2τ -counter α_0 wraps around to 0, the pointer switches to 0, and the nodes
 322 will output the counter value α_0 for τ rounds.
- 323 3. Some correct nodes of block 1 may think that block 1 is counting correctly
 324 for $\Theta(\tau)$ rounds. While this is the case, all of them agree on a value α_1 that
 325 increases by 1 in each round. This value is used to derive the leader pointer
 326 of block 1. Once in 6τ rounds, when the 6τ -counter α_1 wraps around to 0,
 327 the pointer will switch to 0, and the nodes will output the value of α_0 for 3τ
 328 rounds (as the leader pointer does not change for 3τ rounds).
- 329 4. Some correct nodes of block 1 may detect that block 1 is faulty. Such nodes
 330 will output the value of α_0 for $\Theta(\tau)$ rounds.
- 331 5. In summary, eventually there will be τ consecutive rounds during which all
 332 correct nodes output the same counter value α_0 .

333 The other case (block 1 has at most f_1 faulty nodes) is analogous.

334 *Using the Weak Counter..* Now we have constructed a counter that will eventually
 335 produce a consistent output for at least τ rounds. We leverage this property to execute
 336 the *phase king* consensus protocol [6] to stabilise the output counters. The protocol
 337 will have the following crucial property: if all nodes agree on the output, then even if
 338 the round counter becomes inconsistent, the agreement on the output persists. Thus,
 339 it suffices for us that τ is large enough to enable the nodes to consistently execute the
 340 phase king algorithm once to reach agreement; $\tau = O(f)$ will do.

341 The stabilisation time on each level is the maximum of the stabilisation times
 342 of counters \mathbf{A}_i plus $O(\tau) = O(f)$; by choosing $f_1 \approx f_2 \approx f/2$, we can thus ensure
 343 an overall stabilisation time of $O(f)$, irrespectively of the number of recursion levels.
 344 Formally, we prove the following theorem:

345 **THEOREM 4.1.** *Let $c, n > 1$ and $f < n/3$. Define $n_0 = \lfloor n/2 \rfloor$, $n_1 = \lceil n/2 \rceil$,
 346 $f_0 = \lfloor (f-1)/2 \rfloor$, $f_1 = \lceil (f-1)/2 \rceil$, and $\tau = 3(f+2)$. If for $i \in \{0, 1\}$ there exist
 347 synchronous counters $\mathbf{A}_i \in \mathcal{A}(n_i, f_i, c_i)$ such that $c_i = 3^i \cdot 2\tau$, then there exists a
 348 synchronous c -counter $\mathbf{B} \in \mathcal{A}(n, f, c)$ that*

- 349 • stabilises in $T(\mathbf{B}) = \max\{T(\mathbf{A}_0), T(\mathbf{A}_1)\} + O(f)$ rounds, and
- 350 • has state complexity of $S(\mathbf{B}) = \max\{S(\mathbf{A}_0), S(\mathbf{A}_1)\} + O(\log f + \log c)$ bits.

351 We fix the notation of this theorem for the remainder of this section. More-
 352 over, for notational convenience we abbreviate $T = \max\{T(\mathbf{A}_0), T(\mathbf{A}_1)\}$ and $S =$
 353 $\max\{S(\mathbf{A}_0), S(\mathbf{A}_1)\}$.

354 **4.2. Agreeing on a Common Counter (Once in a While).** In this part, we
 355 construct a counter that will eventually count consistently at all nodes for τ rounds.
 356 The τ -counter then will be used as a common clock for executing the phase king
 357 algorithm.

358 We partition the set of nodes $V = V_0 \cup V_1$ such that $V_0 \cap V_1 = \emptyset$, $|V_0| = n_0$ and
 359 $|V_1| = n_1$. We refer to the set V_i as *block i* . For each $i \in \{0, 1\}$, the nodes in set V_i
 360 execute the algorithm \mathbf{A}_i . In case block i has more than f_i faults, we call the block i
 361 faulty. Otherwise, we say that block i is correct. By construction, at least one of the
 362 blocks is correct. Hence, there is a correct block i for which \mathbf{A}_i stabilises within T
 363 rounds, that is, nodes in block i output a consistent c_i -counter in rounds $r \geq T$.

364 LEMMA 4.2. *For some $i \in \{0, 1\}$, block i is correct.*

365 *Proof.* By choice of f_i , we have $f = f_0 + f_1 + 1$. Hence, at least one of the sets V_i
 366 will contain at most f_i faults. \square

367 Next, we apply the typical threshold voting mechanism employed by most Byzan-
 368 tine tolerant algorithms in order to filter out differing views of counter values that are
 369 believed to be consistent. This is achieved by broadcasting candidate counter values
 370 and applying a threshold of $n - f$ as a consistency check, which guarantees that at
 371 most one candidate value from the set $[c]$ can remain. In case the threshold check fails,
 372 a fallback value $\perp \notin [c]$ is used to indicate an inconsistency. This voting scheme is
 373 applied for both blocks concurrently, and all nodes participate in the process, so we
 374 can be certain that fewer than one third of the voters are faulty.

375 In addition to passing this voting step, we require that the counters also have
 376 behaved consistently over a sufficient number of rounds; this is verified by the obvious
 377 mechanism of testing whether the counter increases by 1 each round and counting the
 378 number of rounds since the last inconsistency was detected.

379 In the following, nodes frequently examine a set of values, one broadcast by each
 380 node, and determine majority values. Note that *Byzantine nodes may send different*
 381 *values to different nodes*, that is, it may happen that correct nodes output different
 382 values from such a vote. We refer to a *strong majority* as at least $n - f$ nodes supporting
 383 the same value, which is then called the *majority value*. If a node does not see a strong
 384 majority, it outputs the symbol \perp instead. Clearly, this procedure is well-defined for
 385 $f < n/2$.

386 We will refer to this procedure as a *majority vote*, and slightly abuse notation by
 387 saying “majority vote” when, precisely, we should talk of “the output of the majority
 388 vote at node v ”. Since we require that $f < n/3$, the following standard argument
 389 shows that for each vote, there is a unique value such that each node either outputs
 390 this value or \perp .

391 LEMMA 4.3. *If $v, w \in V \setminus \mathcal{F}$ both observe a strong majority, they output the same*
 392 *majority value.*

393 *Proof.* Fix any set A of $n - f$ correct nodes. For v and w to observe strong
 394 majorities for different values, for each value A must contain $n - 2f$ nodes supporting
 395 it. However, as correct nodes broadcast the same value to each node, this leads to the
 396 contradiction that $|A| \geq 2(n - 2f) = n - f + (n - 3f) > n - f = |A|$. \square

397 We now put this principle to use. In the following, we will use the notation
 398 $x(v, r)$ to refer to the value of local variable x of node v in round r . As we consider
 399 self-stabilising algorithms, the nodes themselves are not aware of what is the value of
 400 r . We introduce the following local variables for each node $v \in V$, block $i \in \{0, 1\}$,
 401 and round $r > 0$ (see Tables 2 and 3):

- 402 • $m_i(v, r)$ stores the most frequent counter value in block i in round r , which
 403 is determined from the broadcasted output variables of \mathbf{A}_i with ties broken
 404 arbitrarily,
- 405 • $M_i(v, r)$ stores the majority vote on $m_i(v, r - 1)$,

TABLE 2
The local state variables used in the boosting construction.

Variable	Range	Description
$m_i(v, r)$	$[c_i]$	the most frequent value observed for the \mathbf{A}_i counter of block i
$M_i(v, r)$	$[c_i] \cup \{\perp\}$	the result of majority vote on $m_i(\cdot, r - 1)$ values
$w_i(v, r)$	$[c_1 + 1]$	“cooldown counter” that is reset if block i behaved inconsistently
$d_i(v, r)$	$[c_i] \cup \{\perp\}$	observation on what seems to be the counter output of block i
$\ell_i(v, r)$	$\{0, 1, \perp\}$	the value of the “leader pointer” for block i
$\ell(v, r)$	$\{0, 1, \perp\}$	leader pointer used by node v
$d(v, r)$	$[\tau]$	once-in-a-while round counter for clocking phase king
$a(v, r)$	$[c] \cup \{\infty\}$	the output of the new c -counter we are constructing
$b(v, r)$	$\{0, 1\}$	helper variable for the phase king algorithm

TABLE 3
Behaviour of local state variables; pointers switch once in $3^i\tau$ rounds.

Variable	Block i is correct	Block i is faulty
$m_i(v, r)$	consistent counter	arbitrary values
$M_i(v, r)$	consistent counter	\perp or some consistent value
$d_i(v, r)$	consistent counter	\perp or some consistent counter
$\ell_i(v, r)$	consistent pointer	\perp or some consistent pointer

406 • $w_i(v, r)$ is a cooldown counter which is reset to $2c_1$ whenever the node perceives
 407 the counter of block i behaving inconsistently, that is, $M_i(v, r) \neq M_i(v, r -$
 408 $1) + 1 \bmod c_i$. Note that this test will automatically fail if either value is \perp .
 409 Otherwise, if the counter behaves consistently, $w_i(v, r) = \max\{w_i(v, r - 1) -$
 410 $1, 0\}$.

411 Clearly, these variables can be updated based on the local values from the previous
 412 round and the states broadcasted at the beginning of the current round. This requires
 413 nodes to store $O(\log c_i) = O(\log f)$ bits.

414 Furthermore, we define the following derived variables for each $v \in V$, block
 415 $i \in \{0, 1\}$, and round r (see Tables 2 and 3):

- 416 • $d_i(v, r) = M_i(v, r)$ if $w_i(v, r) = 0$, otherwise $d_i(v, r) = \perp$,
- 417 • $\ell_i(v, r) = \lfloor d_i(v, r) / (3^i\tau) \rfloor$ if $d_i(v, r) \neq \perp$, otherwise $\ell_i(v, r) = \perp$,
- 418 • for $v \in V_i$, $\ell(v, r) = \ell_i(v, r)$ if $\ell_i(v, r) \neq \perp$, otherwise $\ell(v, r) = \ell_{1-i}(v, r)$, and
- 419 • $d(v, r) = d_{\ell(v, r)}(v, r) \bmod \tau$ if $\ell(v, r) \neq \perp$, otherwise $d(v, r) = 0$.

420 These can be computed locally, without storing or communicating additional values.
 421 The variable $\ell(v, r)$ indicates the block that node v currently considers leader. Note
 422 that some nodes may use $\ell_0(\cdot, r)$ as the leader pointer while some other nodes may
 423 use $\ell_1(\cdot, r)$ as the leader pointer, but this is fine:

- 424 • all nodes v that use $\ell(v, r) = \ell_0(v, r)$ observe the same value $\ell_0(\cdot, r) \neq \perp$,
- 425 • all nodes w that use $\ell(w, r) = \ell_1(w, r)$ observe the same value $\ell_1(\cdot, r) \neq \perp$,
- 426 • eventually $\ell_0(\cdot, r)$ and $\ell_1(\cdot, r)$ will point to the same correct block for τ rounds.

427 We now verify that $\ell(v, r)$ indeed has the desired properties. To this end, we
 428 analyse $d_i(v, r)$. We start with a lemma showing that eventually a correct block’s
 429 counter will be consistently observed by all correct nodes.

430 LEMMA 4.4. *Suppose block $i \in \{0, 1\}$ is correct. Then for all $v, w \in V \setminus \mathcal{F}$, and*

431 rounds $r \geq R = T + O(f)$ it holds that $d_i(v, r) = d_i(w, r)$ and $d_i(v, r) = d_i(v, r - 1) +$
 432 $1 \bmod c_i$.

433 *Proof.* Since block i is correct, algorithm \mathbf{A}_i stabilises within $T(\mathbf{A}_i)$ rounds. As
 434 $f_i < n_i/3$, we will observe correctly $m_i(v, r+1) = m_i(v, r) + 1 \bmod c_i$ for all $r \geq T(\mathbf{A}_i)$.
 435 Consequently, $M_i(v, r+1) = M_i(v, r) + 1 \bmod c_i$ for all $r \geq T(\mathbf{A}_i) + 1$. Therefore,
 436 $w_i(v, r)$ cannot be reset in rounds $r \geq T(\mathbf{A}_i) + 2$, yielding that $w_i(v, r) = 0$ for all
 437 $r \geq T(\mathbf{A}_i) + 2 + 2c_1 = T + O(f)$. The claim follows from the definition of variable
 438 $d_i(v, r)$. \square

439 The following lemma states that if a correct node v does not detect an error in
 440 a block's counter, then any other correct node w that considers the block's counter
 441 correct in any of the last $2c_1$ rounds has a counter value that agrees with v .

442 **LEMMA 4.5.** *Suppose for $i \in \{0, 1\}$, $v \in V \setminus \mathcal{F}$, and $r \geq 2c_1 = O(f)$ it holds that*
 443 $d_i(v, r) \neq \perp$. *Then for each $w \in V \setminus \mathcal{F}$ and each $r' \in \{r - 2c_1 + 1, \dots, r\}$ either*

- 444 • $d_i(w, r') = d_i(v, r) - (r - r') \bmod c_i$, or
- 445 • $d_i(w, r') = \perp$.

446 *Proof.* Suppose $d_i(w, r') \neq \perp$. Thus, $d_i(w, r') = M_i(w, r') \neq \perp$. By **Lemma 4.3**,
 447 either $M_i(v, r') = \perp$ or $M_i(v, r') = M_i(w, r')$. However, $M_i(v, r') = \perp$ would imply
 448 that $w_i(v, r') = 2c_1$ and thus

$$449 \quad w_i(v, r) \geq w_i(v, r') + r' - r = 2c_1 + r' - r > 0,$$

450 contradicting the assumption that $d_i(v, r) \neq \perp$. Thus, $M_i(v, r') = M_i(w, r') =$
 451 $d_i(w, r')$. More generally, we get from $r - r' < 2c_1$ and $w_i(v, r) = 0$ that $w_i(v, r'') \neq 2c_1$
 452 for all $r'' \in \{r', \dots, r\}$. Therefore, we have that $M_i(v, r'' + 1) = M_i(v, r'') + 1 \bmod c$
 453 for all $r'' \in \{r', \dots, r - 1\}$, implying

$$454 \quad d_i(v, r) = M_i(v, r) = M_i(v, r') + r - r' = d_i(w, r') + r - r',$$

455 proving the claim of the lemma. \square

456 The above properties allow us to prove a key lemma: within $T + O(f)$ rounds,
 457 there will be τ consecutive rounds during which the variable $\ell(v, r)$ points to the same
 458 correct block for all correct nodes.

459 **LEMMA 4.6.** *Let R be as in **Lemma 4.4**. There is a round $r \leq R + O(f) = T + O(f)$
 460 and a correct block i so that for all $v \in V \setminus \mathcal{F}$ and $r' \in \{r, \dots, r + \tau - 1\}$ it holds that
 461 $\ell(v, r') = i$.*

462 *Proof.* By **Lemma 4.2**, there exists a correct block i . Thus by **Lemma 4.4**, variable
 463 $d_i(v, r)$ counts correctly during rounds $r \geq R$. If there is no round $r \in \{R, \dots, R + c_i - 1\}$
 464 such that some $v \in V \setminus \mathcal{F}$ has $\ell_{1-i}(v, r) \neq \perp$, then $\ell(v, r) = \ell_i(v, r)$ for all such v and
 465 r and the claim of the lemma holds true by the definition of $\ell_i(v, r)$ and the fact that
 466 $d_i(v, r)$ counts correctly and consistently.

467 Hence, assume that $r_0 \in \{R, \dots, R + c_i - 1\}$ is minimal with the property that
 468 there is some $v \in V \setminus \mathcal{F}$ so that $\ell_{1-i}(v, r_0) \neq \perp$. Therefore, $d_{1-i}(v, r_0) \neq \perp$ and, by
 469 **Lemma 4.5**, this implies for all $w \in V \setminus \mathcal{F}$ and all $r \in \{r_0, \dots, r_0 + 2c_1 - 1\}$ that either
 470 $d_{1-i}(w, r) = \perp$ or $d_{1-i}(w, r) = d_{1-i}(v, r_0) + r - r_0$. In other words, there is a “virtual
 471 counter” that equals $d_{1-i}(v, r_0)$ in round r_0 so that during rounds $\{r_0, \dots, r_0 + 2c_1 - 1\}$
 472 all $d_{1-i}(\cdot, \cdot)$ variables that are not \perp agree with this counter.

473 Consequently, it remains to show that both ℓ_i and the variable ℓ_{1-i} derived
 474 from this virtual counter are equal to i for τ consecutive rounds during the interval

475 $I = \{r_0, \dots, r_0 + 2c_1 - 1\}$, as then $\ell(v, r') = i$ for $v \in V \setminus \mathcal{F}$ and all such rounds r' .

476 Clearly, the c_1 -counter consecutively counts from 0 to $c_1 - 1$ at least once during
 477 the interval $I = \{r_0, \dots, r_0 + 2c_1 - 1\}$. Recalling that $c_1 = 6\tau$, we see that $\ell_1(v, r) = i$
 478 for all $v \in V \setminus \mathcal{F}$ with $\ell_1(v, r) \neq \perp$ for some interval $I_1 \subset I$ of 3τ consecutive
 479 rounds. As $c_0 = 2\tau$, we have that $\ell_0(v, r) = i$ for all $v \in V \setminus \mathcal{F}$ with $\ell_0(v, r) \neq \perp$
 480 for τ consecutive rounds during this subinterval I_1 . Thus, we have an interval
 481 $I_0 = \{r, \dots, r + \tau - 1\} \subseteq I_1$ such that for all $r' \in I_0$ we have $\ell_0(v, r'), \ell_1(v, r') \in \{i, \perp\}$
 482 and $\ell_0(v, r') \neq \perp$ or $\ell_1(v, r') \neq \perp$ yielding $\ell(v, r') = i$ for each correct node. Because
 483 $r < r_0 + 2c_1 - 1 < R + 3c_1 = T + O(f)$, this completes the proof. \square

484 Using the above lemma, we get a counter where all nodes eventually count correctly
 485 and consistently modulo τ for at least τ rounds.

486 **COROLLARY 4.7.** *There is a round $r = T + O(f)$ so that for all $v, w \in V \setminus \mathcal{F}$ it*
 487 *holds that*

- 488 1. $d(v, r) = d(w, r)$ and
- 489 2. for all $r' \in \{r + 1, \dots, r + \tau - 1\}$ we have $d(v, r') = d(v, r' - 1) + 1 \pmod{\tau}$.

490 *Proof.* By [Lemma 4.6](#), there is a round $r = T + O(f)$ and a correct block i such
 491 that for all $v \in V \setminus \mathcal{F}$ we have $\ell(v, r') = i$ for all $r' \in \{r, \dots, r + \tau - 1\}$. Moreover, r is
 492 sufficiently large to apply [Lemma 4.4](#) to $d_i(v, r') = d(v, r')$ for $r' \in \{r + 1, \dots, r + \tau - 1\}$,
 493 yielding the claim. \square

494 **4.3. Reaching Consensus.** [Corollary 4.7](#) guarantees that all correct nodes
 495 eventually agree on a common counter for τ rounds, i.e., we have a weak counter. We
 496 will now use the weak counter to construct a strong counter.

497 Our construction uses a non-self-stabilising consensus algorithm. The basic idea
 498 is that the weak counter serves as the “round counter” for the consensus algorithm.
 499 Hence we will reach agreement as soon as the weak counter is counting correctly. The
 500 key challenge is to make sure that agreement *persists* even if the counter starts to
 501 misbehave. It turns out that a straightforward adaptation of the classic phase king
 502 protocol [6] does the job. The algorithm has the following properties:

- 503 • the algorithm tolerates $f < n/3$ Byzantine failures,
- 504 • the running time of the algorithm is $O(f)$ rounds and it uses $O(\log c)$ bits of
 505 state,
- 506 • if node k is correct, then agreement is reached if all correct nodes execute
 507 rounds $3k, 3k + 1$, and $3k + 2$ consecutively in this order,
- 508 • once agreement is reached, it will persist even if nodes execute *different* rounds
 509 in arbitrary order.

510 We now describe the modified phase king algorithm that will yield a c -counting
 511 algorithm. Denote by $a(v, r) \in [c] \cup \{\infty\}$ the output value of the algorithm at round
 512 r . Here ∞ is used as a “reset state” similarly to \perp in the previous section. There is
 513 also an auxiliary binary value $b(v, r) \in \{0, 1\}$. Define the following short-hand for the
 514 increment operation modulo c :

$$515 \quad x \oplus 1 = \begin{cases} x + 1 \pmod{c} & \text{if } x \neq \infty, \\ \infty & \text{if } x = \infty. \end{cases}$$

516 For $k \in [f + 2]$, we define the instruction sets listed in [Table 4](#). Recall that in
 517 the model of computation that we use in this work, in each round all nodes first
 518 broadcast their current state (in particular, the current value of a), then they receive
 519 the messages, and finally they update their local state. The instruction sets pertain to

TABLE 4
The instruction sets for node $v \in V$ in the phase king protocol.

Set	Instructions for round $r > 0$
I_{3k} :	0a. If fewer than $n - f$ nodes sent $a(v, r - 1)$, set $a(v, r) = \infty$. 0b. Otherwise, $a(v, r) = a(v, r - 1) \oplus 1$.
I_{3k+1} :	1a. Let $z_j = \{u \in V : a(u, r - 1) = j\} $ be the number of j values received. 1b. If $z_{a(v, r-1)} \geq n - f$, set $b(v, r) = 1$. Otherwise, set $b(v, r) = 0$. 1c. Let $z = \min\{j : z_j > f\}$. 1d. Set $a(v, r) = z \oplus 1$.
I_{3k+2} :	2a. If $a(v, r - 1) = \infty$ or $b(v, r - 1) = 0$, set $a(v, r) = \min\{c - 1, a(k, r - 1)\} \oplus 1$. 2b. Otherwise, $a(v, r) = a(v, r - 1) \oplus 1$. 2c. Set $b(v, r) = 1$.

520 the final part—how to update the local state variables a and b based on the messages
 521 received from the other nodes.

522 First, we show that if the instruction sets are executed in the right order by all
 523 correct nodes for a correct leader node $k \in [f + 2]$, then agreement on a counter value
 524 is established.

525 LEMMA 4.8. *Suppose that for some correct node $k \in [f + 2]$ and a round $r > 2$, all*
 526 *non-faulty nodes execute instruction sets I_{3k} , I_{3k+1} , and I_{3k+2} in rounds $r - 2$, $r - 1$,*
 527 *and r , respectively. Then $a(v, r) = a(u, r) \neq \infty$ for any two correct nodes $u, v \in V$.*
 528 *Moreover, $b(v, r + 1) = 1$ at each correct node $v \in V$.*

529 *Proof.* This is essentially the correctness proof for the phase king algorithm.
 530 Without loss of generality, we can assume that the number of faulty nodes is exactly
 531 f . Since we have $f < n/3$, it is not possible that two correct nodes $u, v \in V \setminus \mathcal{F}$
 532 both satisfy $a(v, r - 2) \neq a(u, r - 2)$ and $a(v, r - 2), a(u, r - 2) \in [c]$: otherwise, on
 533 round $r - 2$, nodes u and v would have observed different majority values contradicting
 534 Lemma 4.3. Therefore, there exists some $x \in [c]$ such that $a(v, r - 2) \in \{x, \infty\}$ for all
 535 $v \in V \setminus \mathcal{F}$. Checking I_{3k+1} we get that $a(v, r - 1) \in \{x + 1 \bmod c, \infty\}$, as no node can
 536 see values other than x or ∞ more than f times when executing instruction 1c.

537 To prove the claim, it remains to consider two cases when executing instructions
 538 in I_{3k+2} . In the first case, all non-faulty nodes execute instruction 2a on round r .
 539 Then $a(u, r) = a(v, r) = \min\{c - 1, a(k, r - 1)\} \oplus 1 \in [c]$ for any $u, v \in V \setminus \mathcal{F}$.

540 In the second case, there is some node v not executing instruction 2a. Hence,
 541 $a(v, r - 1) \neq \infty$ and $b(v, r - 1) = 1$, implying that v computed $z_{a(v, r-2)} \geq n - f$ on
 542 round $r - 1$. Consequently, at least $n - 2f > f$ correct nodes u satisfy $a(u, r - 2) =$
 543 $a(v, r - 2) \neq \infty$. We can now infer that $a(u, r - 1) = a(v, r - 1) = a(v, r - 2) + 1 \bmod c$ for
 544 all correct nodes u : instruction 1c must evaluate to $a(v, r - 1) \in [c]$ at all correct nodes,
 545 because we know that no correct node u satisfies that both $a(u, r - 2) \neq a(v, r - 2)$
 546 and $a(u, r - 2) \neq \infty$. This implies that $a(u, r) = a(v, r) \neq \infty$ for all correct nodes u ,
 547 regardless of whether they execute instruction 2a. Trivially, $b(v, r) = 1$ at each correct
 548 node v due to instruction 2c. \square

549 Next, we argue that once agreement is established, it persists—it does not matter
 550 any more which instruction sets are executed.

551 LEMMA 4.9. *Assume that $a(v, r) = x \in [c]$ and $b(v, r) = 1$ for all correct nodes v*
 552 *in some round r . Then $a(v, r + 1) = x + 1 \bmod c$ and $b(v, r + 1) = 1$ for all correct*

553 nodes v .

554 *Proof.* Each node observes at least $n - f$ nodes with counter value $x \in [c]$, and
 555 hence at most f nodes with some value $y \neq x$. Let v be a correct node and consider
 556 all possible instruction sets it may execute.

557 First, consider the case where instruction set I_{3k} is executed. In this case, v
 558 increments x , resulting in $a(v, r + 1) = x + 1 \bmod c$ and $b(v, r + 1) = 1$. Second,
 559 executing I_{3k+1} , node v evaluates $z_x \geq n - f$ and $z_y \leq f$ for all $y \neq x$. Hence it sets
 560 $b(v, r + 1) = 1$ and $a(v, r + 1) = x + 1 \bmod c$. Finally, when executing I_{3k+2} , node v
 561 skips instruction 2a and sets $a(v, r + 1) = x + 1 \bmod c$ and $b(v, r + 1) = 1$. \square

562 **4.4. Proof of Theorem 4.1.** We now have all the building blocks to devise
 563 an f -resilient c -counter running on n nodes. The idea is as follows: first, we use
 564 the construction given in Section 4.2 to get a weak τ -counter that eventually counts
 565 correctly for $\tau = 3(f + 2)$ rounds. Concurrently, all nodes execute the modified phase
 566 king algorithm given in Section 4.3 which by Lemma 4.8 and Lemma 4.9 guarantees
 567 that all nodes will establish and maintain agreement on the output variable for the
 568 c -counter.

569 **THEOREM 4.1.** *Let $c, n > 1$ and $f < n/3$. Define $n_0 = \lfloor n/2 \rfloor$, $n_1 = \lceil n/2 \rceil$,
 570 $f_0 = \lfloor (f - 1)/2 \rfloor$, $f_1 = \lceil (f - 1)/2 \rceil$, and $\tau = 3(f + 2)$. If for $i \in \{0, 1\}$ there exist
 571 synchronous counters $\mathbf{A}_i \in \mathcal{A}(n_i, f_i, c_i)$ such that $c_i = 3^i \cdot 2\tau$, then there exists a
 572 synchronous c -counter $\mathbf{B} \in \mathcal{A}(n, f, c)$ that*

- 573 • stabilises in $T(\mathbf{B}) = \max\{T(\mathbf{A}_0), T(\mathbf{A}_1)\} + O(f)$ rounds, and
- 574 • has state complexity of $S(\mathbf{B}) = \max\{S(\mathbf{A}_0), S(\mathbf{A}_1)\} + O(\log f + \log c)$ bits.

575 *Proof.* First, we apply the construction underlying Corollary 4.7. Then we have
 576 every node $v \in V$ in each round r execute the instructions for round $d(v, r)$ of
 577 the phase king algorithm from Section 4.3. It remains to show that this yields a
 578 correct algorithm \mathbf{B} with stabilisation time $T(\mathbf{B}) = T + O(f)$ and state complexity
 579 $S(\mathbf{B}) = S + O(\log f + \log c)$, where $T = \max\{T(\mathbf{A}_i)\}$ and $S = \max\{S(\mathbf{A}_i)\}$.

580 By Corollary 4.7, there exists a round $r = T + O(f)$ so that the variables $d(v, r)$
 581 behave as a consistent τ -counter during rounds $\{r, \dots, r + \tau - 1\}$ for all $v \in V \setminus \mathcal{F}$.
 582 As there are at most f faulty nodes, there exist at least two correct nodes $v \in [f + 2]$.
 583 Since $\tau = 3(f + 2)$, then for at least one correct node $k \in [f + 2] \setminus \mathcal{F}$, there is a
 584 round $r \leq r_k \leq r + \tau - 3$ such that $d(w, r_k + h) = 3k + h$ for all $w \in V \setminus \mathcal{F}$ and
 585 $h \in \{0, 1, 2\}$. Therefore, by Lemma 4.8 and Lemma 4.9, the output variables satisfy
 586 $a(v, r') = a(w, r') \in [c]$ for all correct nodes and rounds $r' \geq r_k + 3$. Thus, the
 587 algorithm stabilises in $r_v + 3 \leq r + \tau = r + O(f) = T + O(f)$ rounds.

588 The bound for the state complexity follows from the facts that, at each node, we
 589 need at most S bits to store the state of \mathbf{A}_i and $O(\log \tau + \log c) = O(\log f + \log c)$
 590 bits to store the variables listed in Table 2. \square

591 **5. Deterministic Counting.** In this section, we use the construction given in
 592 the previous section to obtain algorithms that only need a small number of state bits.
 593 Essentially, all that remains is to recursively apply Theorem 4.1. Each step of the
 594 recursion roughly doubles the resilience in an optimal manner: if we start with an
 595 optimally resilient algorithm, we get a new algorithm with higher, but still optimal,
 596 resilience. Therefore, to get any desired resilience of $f > 0$, it suffices to repeat the
 597 recursion for $\Theta(\log f)$ many steps. Figure 1 illustrates how we can recursively apply
 598 Theorem 4.1.

599 We now analyse the correctness, time and state complexity of the resulting
 600 algorithms.




FIG. 1. An example on how to recursively construct a 5-resilient algorithm running on 16 nodes. The small circles represent the nodes. Each group of four nodes runs a 1-resilient counter $\mathbf{A}(4, 1)$. On top of this, each larger group of 8 nodes runs a 2-resilient counter $\mathbf{A}(8, 2)$ attained from the first step of recursion. At the top-most layer, all of the 16 nodes run a 5-resilient counter $\mathbf{A}(16, 5)$. Faulty nodes are black and faulty blocks are gray.

601 THEOREM 1.1. For any integers $c, n > 1$ and $f < n/3$, there exists a deterministic
 602 f -resilient synchronous c -counter that runs on n nodes, stabilises in $O(f)$ rounds, and
 603 uses $O(\log^2 f + \log c)$ bits to encode the state of a node.

604 *Proof.* We show the claim by induction on f . The induction hypothesis is that
 605 for all $f > f' \geq 0$, $c > 1$, and $n' > 3f'$, we can construct $\mathbf{B} \in \mathcal{A}(f', n', c)$ with

$$606 \quad T(\mathbf{B}) \leq 1 + \alpha f' \sum_{k=0}^{\lceil \log f' \rceil} (1/2)^k \quad \text{and} \quad S(\mathbf{B}) \leq \beta(\log^2 f' + \log c),$$

607 where α and β are sufficiently large constants and for $f' = 0$ the sum is empty, that is,
 608 $T(\mathbf{B}) \leq 1$. As $\sum_{k=0}^{\infty} (1/2)^k = 2$, the time bound will be $O(f')$.

609 Note that for $f \geq 0$ it is sufficient to show the claim for $n(f) = 3f + 1$, as we can
 610 easily generalise to any $n > n(f)$ by running \mathbf{B} on the first $n(f)$ nodes and letting
 611 the remaining nodes follow the majority counter value among the first $n(f)$ nodes
 612 executing the algorithm; this increases the stabilisation time by one round and induces
 613 no memory overhead.

614 For the base case, observe that a 0-resilient c -counter of $n(0) = 1$ node is trivially
 615 given by the node having a local counter. It stabilises in 0 rounds and requires $\lceil \log c \rceil$
 616 state bits. As pointed out above, this implies a 0-resilient c -counter for any n with
 617 stabilisation time 1 and $\lceil \log c \rceil$ bits of state.

618 For the inductive step to f , we apply [Theorem 4.1](#) with the parameters $n_0 = \lfloor n/2 \rfloor$,
 619 $n_1 = \lceil n/2 \rceil$, $f_0 = \lfloor (f-1)/2 \rfloor$, $f_1 = \lceil (f-1)/2 \rceil$, $\tau = 3(f+2)$ and $c_i = 3^i \cdot 2\tau$. Since
 620 $f_i \leq f/2$ and $n_i > 3f_i$, for $i \in \{0, 1\}$, the induction hypothesis gives us algorithms
 621 $\mathbf{A}_i(n_i, f_i, c_i)$. Now by applying [Theorem 4.1](#) we get an algorithm \mathbf{B} with

$$622 \quad T(\mathbf{B}) = \max\{T(\mathbf{A}_0), T(\mathbf{A}_1)\} + O(f)$$

$$623 \quad \leq 1 + \frac{\alpha f}{2} \sum_{k=0}^{\lceil \log f/2 \rceil} \left(\frac{1}{2}\right)^k + O(f)$$

$$624 \quad = 1 + \alpha f \sum_{k=1}^{\lceil \log f \rceil} \left(\frac{1}{2}\right)^k + O(f)$$

$$625 \quad \leq 1 + \alpha f \sum_{k=0}^{\lceil \log f \rceil} \left(\frac{1}{2}\right)^k,$$
 626

627 where in the second to last step we use that α is a sufficiently large constant. Since
 628 the sum is at most 2, we get that $T(\mathbf{B}) = O(f)$. Moreover, the state complexity is
 629 bounded by

$$\begin{aligned}
 630 \quad S(\mathbf{B}) &= \max\{S(\mathbf{A}_0), S(\mathbf{A}_1)\} + O(\log f + \log c) \\
 631 \quad &\leq \beta \left(\log^2 \frac{f}{2} + \log \frac{f}{2} \right) + O(\log f + \log c) \\
 632 \quad &\leq \beta (\log^2 f + \log c), \\
 633
 \end{aligned}$$

634 where we exploit that β is a sufficiently large constant. Hence, $S(\mathbf{B}) = O(\log^2 f + \log c)$,
 635 the induction step succeeds, and the proof is complete. \square

636 **6. Reducing the Number of Bits Communicated.** In this section, we dis-
 637 cuss how to reduce the number of bits broadcast by a node *after* stabilisation. We
 638 consider the following extension of the model of computation: instead of a node always
 639 broadcasting its current state, we allow it to broadcast an arbitrary message (including
 640 an empty message) each round. Formally, this entails that we extend the definition of
 641 an algorithm by (1) introducing a new function $\mu: [n] \times X \rightarrow \mathcal{M}$ that maps the current
 642 state x to a message $\mu(x)$ which is broadcast and (2) modify the state transition
 643 function to map the old internal state and the vector of received messages to a new
 644 state, that is, the new state transition function has the form $g': [n] \times X \times \mathcal{M}^n \rightarrow X$.

645 First, we show how to construct counters that only send $O(1 + B \log B)$ bits every
 646 κ rounds, where $B = O(\log c / \log \kappa)$, while increasing the stabilisation time only by
 647 an additive $O(\kappa)$ term, where $\kappa = \Omega(f)$ is a parameter. In particular, we show that
 648 for polynomial-sized counters with optimal resilience, the algorithm only needs to
 649 communicate an asymptotically optimal number of bits after stabilisation:

650 **COROLLARY 6.1.** *For any $n > 1$ and $c = n^{O(1)}$ that is an integer multiple of*
 651 *n , there exists a synchronous c -counter that runs on n nodes, has optimal resilience*
 652 *$f = \lfloor (n-1)/3 \rfloor$, stabilises in $\Theta(n)$ rounds, requires $O(\log^2 n)$ bits to encode the state*
 653 *of a node, and for which after stabilisation correct nodes broadcast asymptotically*
 654 *optimal $O(1)$ bits per $\Theta(n)$ rounds.*

655 We start by outlining the high-level idea of the approach, then give a detailed
 656 description of the construction we use, and finally prove the main results of this section.

657 **6.1. High-Level Idea.** The techniques we use are very similar to the ones we
 658 used for deriving [Theorem 1.1](#). Essentially, we devise a “silencing wrapper” for
 659 algorithms given by [Theorem 1.1](#). Let \mathbf{A} be such a counting algorithm. The high-level
 660 idea and the key ingredients are the following:

- 661 • The goal is that nodes eventually become *happy*: they assume stabilisation
 662 has occurred and check for counter consistency only every κ rounds (as self-
 663 stabilising algorithms always need to verify their output).
- 664 • Happy nodes do not execute the underlying algorithm \mathbf{A} .
- 665 • Using a cooldown counter with similar effects as shown in [Lemma 4.5](#), we
 666 enforce that all happy nodes output consistent counters.
- 667 • We override the phase king instruction of \mathbf{A} if at least $n - 2f \geq f + 1$ nodes
 668 claim to be happy and propose a counter value x . In that case nodes adjust
 669 their counter output to match x . If there is no strong majority of happy nodes
 670 supporting a counter value, either all nodes become unhappy or all correct
 671 nodes reach agreement and start counting correctly.

- 672 • If all correct nodes are unhappy, they execute \mathbf{A} “as is” reaching agreement
673 eventually.
- 674 • The counters are used to make all nodes concurrently switch their state to
675 being happy, in a way that does not interfere with the above stabilisation
676 process.

677 We will show that happy nodes can communicate their counter values very effi-
678 ciently in a manner that self-stabilises within κ rounds. As their counter increases by
679 1 modulo c every round (or they become unhappy), they can use multiple rounds to
680 encode a counter value; the recipient simply counts locally in the meantime.

681 **6.2. The Silencing Wrapper.** Let $\mathbf{A} \in \mathcal{A}(n, f, c)$ be an algorithm given by
682 [Theorem 1.1](#) and let $c = j\kappa$ for any $j > 0$ and $\kappa > T(\mathbf{A})$. We use the short-hand
683 $T = T(\mathbf{A})$ throughout this section. Let $a(v, r)$ be the output of the synchronous
684 counting algorithm for node v in round r . Recall that by a *strong majority* we mean
685 that at least $n - f$ received messages support a value. We now modify \mathbf{A} so that it
686 meets the additional requirement of little communication after stabilisation.

687 We introduce two new variables: a cooldown counter $t(v, r) \in [T + 1]$ and a
688 “happiness” indicator $h(v, r) \in \{0, 1\}$. These are updated according to the following
689 rules in every round $r > 0$:

- 690 1. Set $t(v, r) = T$ if there was no strong majority of nodes w with $a(w, r - 1) =$
691 $a(v, r - 1)$ or $a(v, r) \neq a(v, r - 1) + 1 \pmod{c}$. Otherwise, decrement the counter,
692 that is, $t(v, r) = \max\{0, t(v, r - 1) - 1\}$.
- 693 2. Set $h(v, r) = 0$ if $h(v, r - 1) = 1$, but there was no strong majority of nodes w
694 with $h(w, r - 1) = 1$ and $a(w, r - 1) = a(v, r - 1)$, or if $t(v, r) > 0$. Set $h(v, r) = 1$
695 if $t(v, r - 1) = 0$ and $a(v, r - 1) = 0 \pmod{\kappa}$. Otherwise, $h(v, r) = h(v, r - 1)$.
- 696 3. If $h(v, r) = 0$, execute a single step of \mathbf{A} *except* for the phase king instructions
697 given in [Table 4](#). The counter value $a(v, r + 1)$ is updated according to the
698 next rule.
- 699 4. If received $n - 2f$ times a value $a(w, r) = x$ from nodes with $h(w, r) = 1$, set
700 $a(v, r + 1) = x + 1 \pmod{c}$; if there are two such values x , it does not matter
701 which is chosen. Otherwise, execute *only* the phase king instructions of \mathbf{A}
702 given in [Table 4](#) as indicated by the once-in-a-while round counter $d(v, r)$ as
703 usual; in particular, this determines $a(v, r + 1)$.

704 In the following, we say that a node $v \in V \setminus \mathcal{F}$ with value $h(v, r) = 1$ is *happy* in
705 round r and *unhappy* if $h(v, r) = 0$. Moreover, the counters *converge* in round r if for
706 all $v, w \in V \setminus \mathcal{F}$, it holds that $a(v, r) = a(w, r)$. The idea is to show that not only do the
707 counters converge (and then count correctly), but also all correct nodes become happy.
708 As a happy node that remains happy simply increases its counter value by 1 modulo c ,
709 there is no need to explicitly communicate this except for verification purposes. It is
710 straightforward to exploit this to ensure that the algorithm communicates very little
711 (explicitly) once all nodes are happy; we will discuss this after showing stabilisation of
712 the routine.

713 **6.3. Proof of Stabilisation.** Let us first establish that if the counters converge,
714 they will keep counting correctly and correct nodes will become happy within $O(\kappa + T)$
715 additional rounds for any parameter $\kappa > T$.

716 **LEMMA 6.2.** *If the counters converge in round r , then $a(u, r') = a(v, r') = a(u, r) +$
717 $(r - r') \pmod{c}$ for all $u, v \in V \setminus \mathcal{F}$ and $r' \geq r$.*

718 *Proof.* Since the counters have converged, there is a strong majority of nodes
719 supporting the same value. Hence, variable $a(u, r')$ is updated according to Rule 4.

720 As all counter values from correct nodes are identical, it does not matter whether
 721 these nodes are happy or not; either way, the counters are increased by 1 modulo c
 722 (cf. Lemma 4.9). \square

723 LEMMA 6.3. *If the counters converge in round r , then for all rounds $r' \geq r + T + \kappa$
 724 and all nodes $v \in V \setminus \mathcal{F}$ we have $h(v, r') = 1$.*

725 *Proof.* By Lemma 6.2, the agreement on output values will persist once reached.
 726 Hence, at all nodes $v \in V \setminus \mathcal{F}$ we have $t(v, r') = 0$ in all rounds $r' \geq r + T$ by Rule 1.
 727 Therefore, there is a round $r' \leq r + T + \kappa$ so that $t(v, r') = 0$ and $a(v, r') = 0 \pmod{\kappa}$
 728 at all such v . Consequently, all correct nodes jointly set $h(v, r' + 1) = 1$. By induction
 729 on the round number, we see that no such node sets $h(v, r'') = 0$ for $r'' > r' + 1$, as
 730 there is always a strong majority of $n - f$ happy and correct nodes supporting the
 731 (joint) counter value. \square

732 We now proceed to show that the counters converge within $O(\kappa + T)$ rounds.
 733 The first step is to observe that if no correct node is happy, then algorithm **A** is run
 734 without modification, and hence, the counters converge in T rounds.

735 LEMMA 6.4. *Let $r \geq T$. If for all $v \in V \setminus \mathcal{F}$ and $r' \in \{r - T + 1, \dots, r\}$, we have
 736 $h(v, r') = 0$, then the counters converge in round $r + 1$.*

737 *Proof.* Since $h(v, r') = 0$, each node v applies Rule 3 in any such round r' . As
 738 there are no happy nodes in round r' , a node can never receive the same counter value
 739 from more than f nodes that (claim to be) happy. Hence, Rule 4 boils down to just
 740 updating $a(v, r')$ according to the rules of **A**. As $T = T(\mathbf{A})$, algorithm **A** stabilises
 741 and thus $a(v, r) = a(w, r)$ for all $v, w \in V \setminus \mathcal{F}$. \square

742 To deal with the case that some nodes may be happy (which entails that not all
 743 nodes may execute **A** correctly, destroying its guarantees), we argue that ongoing
 744 happiness also implies that the counters converge. To this end, we first show that
 745 the cooldown counters $t(v, r)$ ensure that correct nodes whose counters are 0 count
 746 correctly and agree on their counter values. This is shown analogously to Lemma 4.5.

747 LEMMA 6.5. *Let $r > T$ and $v, w \in V \setminus \mathcal{F}$. If $t(v, r) = t(w, r') = 0$ for $r' \in$
 748 $\{r - T + 1, \dots, r\}$, then $a(v, r) = a(w, r') + r - r' \pmod{c}$.*

749 *Proof.* Since $t(v, r) = 0$, by Rule 1 it holds that $t(v, r') \leq r - r' < T$. Hence,
 750 both v and w saw a strong majority of nodes u with $a(u, r' - 1) = a(v, r' - 1)$ and
 751 $a(u, r' - 1) = a(w, r' - 1)$, respectively. By Lemma 4.3, it follows that $a(v, r' -$
 752 $1) = a(w, r' - 1)$. Likewise, $t(v, r'') \neq T$ for rounds $r' < r'' \leq r$, implying that
 753 $a(v, r) = a(v, r') + r - r' \pmod{c}$, and $a(w, r') = a(w, r' - 1) + 1 \pmod{c} = a(v, r')$. \square

754 Except for the initial rounds, the above lemma implies that happy nodes always
 755 have the same counter value: by Rule 2, a node v with $h(v, r) = 1$ must have $t(v, r) = 0$.
 756 A node remaining happy thus entails that every node receives the same counter value
 757 from at least $n - 2f \geq f + 1$ happy nodes, and no other counter value with the same
 758 property may be perceived. In other words, a node staying happy implies that the
 759 counters converge.

760 LEMMA 6.6. *If $h(v, r - 1) = h(v, r) = 1$ for some $v \in V \setminus \mathcal{F}$ and $r > 3$, then the
 761 counters converge in round $r + 1$.*

762 *Proof.* By Rule 2, any node w with $h(v, r) = 1$ satisfies $t(w, r) = 0$. We apply
 763 Lemma 6.5 to see that, for any $w \in V \setminus \mathcal{F}$ that is happy in round $r - 1$, we have
 764 that $a(v, r - 1) = a(w, r - 1)$. As $h(v, r) = h(v, r - 1) = 1$, node v observed a strong
 765 majority of happy nodes w with $a(v, r - 1) = a(w, r - 1)$ in round $r - 1$, implying

766 that all nodes received this counter value from at least $n - 2f \geq f + 1$ happy nodes.
 767 Together with Rule 4, these observations imply that $a(u, r) = a(v, r - 1) + 1 \pmod c$
 768 for all $u \in V \setminus \mathcal{F}$. \square

769 Using these lemmas and the fact that nodes may become happy only after counting
 770 consistently for sufficiently long *and* when their counters are 0 modulo $\kappa > T$, we can
 771 show that the counters converge in all cases.

772 **LEMMA 6.7.** *Within $O(\kappa)$ rounds, the counters converge.*

773 *Proof.* Either all $v \in V \setminus \mathcal{F}$ with $h(v, 3) = 1$ set $h(v, 4) = 0$ or **Lemma 6.6** shows the
 774 claim. If there are no nodes v with $h(v, r) = 1$ for $r \in \{4, \dots, T + 3\}$, then **Lemma 6.4**
 775 shows the claim. Hence, assume that there is some node v with $h(v, r) = 1 \neq h(v, r - 1)$
 776 for some minimal $r \in \{4, \dots, T + 3\}$. Again, either $h(v, r + 1) = 0$ for all such nodes
 777 or we can apply **Lemma 6.6**; thus assume the former in the following.

778 Suppose for contradiction that there is a node w with $h(w, r') = 1$ for a minimal
 779 $r' \in \{r + 1, \dots, r + T\}$. As r' is minimal and all nodes with $h(v, r) = 1$ have
 780 $h(v, r + 1) = 0$, it must hold that $h(w, r' - 1) = 0$. Hence, $t(w, r' - 1) = 0 = t(v, r - 1)$.
 781 By **Lemma 6.5**, this implies that $a(w, r' - 1) = a(v, r - 1) + r - r' \pmod c$. However,
 782 $\kappa > T$, $0 < r - r' \leq T$, and $a(v, r - 1) = 0 \pmod \kappa$, implying that $a(w, r' - 1) \neq 0 \pmod \kappa$,
 783 which (by Rule 2) is a contradiction to $h(w, r') = 1 \neq h(w, r' - 1)$.

784 We conclude that $h(v, r') = 0$ for all v and $r' \in \{r + 1, \dots, r + T\}$. The claim
 785 follows by applying **Lemma 6.4**. \square

786 We now can conclude that within $O(\kappa)$ rounds, the algorithm stabilises in the
 787 sense that all nodes become happy and count correctly and consistently.

788 **COROLLARY 6.8.** *There exists a round $R = O(\kappa)$ such that for all $v \in V \setminus \mathcal{F}$ and
 789 $r \geq R$, it holds that $h(v, r) = 1$, and $a(v, r) = a(v, r - 1) + 1 \pmod c$, and $a(v, r) = a(w, r)$
 790 for all $w \in V \setminus \mathcal{F}$.*

791 *Proof.* By **Lemma 6.7** we get that there exists a round $r' = O(\kappa)$ in which the
 792 counters converge. Since $r' + T + \kappa = O(\kappa)$, happiness follows from **Lemma 6.3** and
 793 agreement follows from **Lemma 6.2**. \square

794 **6.4. Reducing the Communication Complexity after Stabilisation.** As
 795 noted earlier, the counter variables for happy nodes count modulo c . Hence, it is trivial
 796 to deduce the counter value of a happy node from its counter value in an earlier round.
 797 Moreover, happy nodes do not execute algorithm **A**. Therefore, we can change the
 798 encoding of the happy nodes' counter values to reduce the communication complexity
 799 after stabilisation.

800 **COROLLARY 6.9.** *Suppose happy nodes communicate their counter values by any
 801 method that stabilises in κ rounds, then the algorithm presented in this section retains
 802 its properties, except that its stabilisation time increases by an additive κ rounds.*

803 The above immediately implies that happy nodes v could simply transmit the
 804 $a(v, r)$ only in rounds r when $a(v, r) \pmod \kappa = 0$ and perform no other communication.
 805 The fact that v does not transmit readily implies that it is happy, permitting to derive
 806 its counter value by counting from the most recent value v transmitted. Moreover, by
 807 **Lemma 6.5** the output counters of happy nodes agree after $O(1)$ rounds. Thus, a single
 808 local counter suffices for verification yielding a cost of using only $\lceil \log c \rceil$ additional
 809 bits of memory per node.

810 Clearly, this trivial encoding mechanism stabilises in κ rounds. However, we can
 811 do much better. For simplicity, we do not try to give a tight bound here.

812 LEMMA 6.10. *Happy nodes can communicate their counter values by sending only*
 813 *$O(1 + B \log B)$ bits per κ rounds, where $B = O(\log c / \log \kappa)$, in a way that stabilises*
 814 *in κ rounds.*

815 *Proof.* First, we fix two unique bit strings HAPPY and UNHAPPY both having
 816 a length of $O(1)$ bits. We mark all messages from unhappy nodes with the header
 817 UNHAPPY. Happy nodes $v \in V \setminus \mathcal{F}$ send the bit string HAPPY in rounds r when
 818 $a(v, r) \bmod \kappa = 0$. In this and the subsequent $\kappa - 1$ rounds, they furthermore send
 819 up to b bits in order to encode the value of $a(v, r) \in [c]$, where they avoid the two
 820 excluded unique bit strings HAPPY and UNHAPPY. Since we are only interested in the
 821 asymptotic behaviour, we may neglect these possible collisions and determine how
 822 large b must be so that in κ rounds we can encode c different values.

823 Since there are κ rounds in which to broadcast a message, we can think each round
 824 as being a bin containing the bits broadcast by a node. Suppose we have $B = b / \log b$
 825 uniquely labelled balls that we can place in κ different bins. This way we can encode
 826 B -length strings over an alphabet of size κ by interpreting each ball in a bin $i \in [\kappa]$ as
 827 giving the indices for the symbol i . This allows us to encode a total of κ^B distinct
 828 values.

829 Since encoding the unique label of a single ball takes $O(\log B)$ bits and we can
 830 use constant-sized delimiters when encoding the set of balls in a single bin, we need
 831 $O(B \log B)$ bits to encode all the values. Thus, each node communicates a total of
 832 $O(B \log B) = O(b)$ bits during the course of κ rounds. In order to encode c different
 833 values, it suffices to satisfy $c \leq \kappa^B$. This can be done by choosing $B \geq \log c / \log \kappa$.
 834 Taking into account the bits for delimiters and the HAPPY string, the claim follows. \square

835 Overall, we obtain the following theorem.

836 THEOREM 6.11. *For any integers $n > 1$, $f < n/3$, $\kappa = \Omega(f)$, and $c = \kappa j$ for*
 837 *$j > 0$, there exists an f -resilient synchronous c -counter that runs on n nodes, stabilises*
 838 *in $O(\kappa)$ rounds, and requires $O(\log^2 f + \log c)$ bits to encode the state of a node.*
 839 *Moreover, once stabilised, nodes send only $O(1 + B \log B)$ bits per κ rounds, where*
 840 *$B = O(\log c / \log \kappa)$.*

841 *Proof.* Let $\mathbf{A} \in \mathcal{A}(n, f, c)$ be an algorithm given by [Theorem 1.1](#). As $T(\mathbf{A}) = \Theta(f)$,
 842 for any $\kappa > T(\mathbf{A})$, the claim now directly follows from [Corollaries 6.8](#) and [6.9](#) and
 843 [Lemma 6.10](#), where we note that only a constant number of variables of size at most
 844 $\max\{T(\mathbf{A}), c\}$ need to be encoded in the state of a node. \square

845 We remark that since $\kappa > T(\mathbf{A}) = \Theta(f)$, in case of optimal resilience and $c = n^{O(1)}$, it
 846 holds that $B = O(1)$, and thus also, $O(1 + B \log B) = O(1)$.

847 COROLLARY 6.12. *For any $n > 1$ and $c = n^{O(1)}$ that is an integer multiple of*
 848 *n , there exists a synchronous c -counter that runs on n nodes, has optimal resilience*
 849 *$f = \lfloor (n - 1)/3 \rfloor$, stabilises in $\Theta(n)$ rounds, requires $O(\log^2 n)$ bits to encode the state*
 850 *of a node, and for which after stabilisation correct nodes broadcast asymptotically*
 851 *optimal $O(1)$ bits per $\Theta(n)$ rounds.*

852 *Proof.* All properties except for the optimality of the last point follow from the
 853 choice of parameters by picking $\kappa = \Theta(n)$ in [Theorem 6.11](#). The claimed optimality
 854 follows from the fact that in order to prove to a node that its counter value is
 855 inconsistent with that of others, it must receive messages from at least $f + 1 = \Theta(n)$
 856 nodes; to guarantee stabilisation in $O(n)$ rounds, this must happen every $\Omega(n)$ rounds
 857 for each correct node. \square

858 **7. Sending Fewer Messages.** So far we have considered the size of messages
 859 nodes need to broadcast every round. In the case of the algorithm given in [Theorem 1.1](#),
 860 every node will send $S = O(\log^2 f + \log c)$ bits in each round. As there are $\Theta(n^2)$
 861 communication links, the total number of communicated bits in each round is $\Theta(S \cdot n^2)$.
 862 In this section, we consider a randomised variant of the algorithm that achieves better
 863 message and bit complexities in a slightly different communication model.

864 **7.1. Pulling Model.** Throughout this section we consider the following variant
 865 of our communication model, where in every synchronous round t each correct node v :
 866 1. contacts a subset $C(v, t) \subseteq V$ of other nodes to *pull* information from,
 867 2. pulls a response message $r_u \in \mathcal{M}$ from every contacted node $u \in C(v, t)$,
 868 3. updates its local state according to its current state and the responses it
 869 received.

870 Thus, every round t node v obtains a message vector $\mathbf{m} = \langle m_0, \dots, m_{n-1} \rangle$, where
 871 $m_u = r_u$ if $u \in C(v, t)$ and $m_u = \perp$, otherwise. Besides this modification, the model of
 872 computation is as before: node v updates its state using the state transition function
 873 $g: [n] \times X \times \mathcal{M}^n \rightarrow X$ and a correct node u in state x_u responds with the message
 874 $\mu(x_u)$, where $\mu: X \rightarrow \mathcal{M}$ maps the internal state of a node to a message. However
 875 in the pulling model, the algorithm also needs to specify the set $C(v, t)$ of nodes it
 876 contacts every round. We assume that every correct node chooses this set randomly
 877 independent of its internal state.

878 As before, faulty nodes may respond with arbitrary messages that can be different
 879 for different pulling nodes. We define the (per-node) message and bit complexities of
 880 the algorithm as the maximum number of messages and bits, respectively, pulled by a
 881 non-faulty node in any round.

882 This model is motivated by the challenges of designing energy-limited fault-tolerant
 883 circuits. We suggest the approach in which each node that makes a request for data
 884 also has to provide the energy resources for processing and answering the request.
 885 This way by limiting the energy supply of each individual node, we can also effectively
 886 limit the total amount of energy wasted due to the actions of the Byzantine nodes.
 887 However, to make this approach feasible, we have to design an algorithm in which
 888 each non-faulty node needs to make only a few requests for data. In this section we
 889 design a randomised algorithm that satisfies this property.

890 **7.2. High-Level Idea of the Probabilistic Construction.** To keep the num-
 891 ber of pulls, and thus number of messages sent, small, we modify the construction of
 892 [Theorem 4.1](#) to use random sampling where useful. Essentially, the idea is to show that
 893 *with high probability* a small set of sampled messages accurately represents the current
 894 state of the system and the randomised algorithm will behave as the deterministic
 895 one. There are two steps where the nodes rely on information broadcast by the all the
 896 nodes: the majority voting scheme over the blocks and the variant of the phase king
 897 algorithm. In the following, both are shown to work under the sampling scheme with
 898 high probability by using concentration bound arguments.

899 More specifically, here *with high probability* means that for any constant $k \geq 1$ the
 900 probability of failure is bounded above by η^{-k} when sampling $K = \Theta(\log \eta)$ messages
 901 (where the constants in the asymptotic notation may depend on k); here η denotes the
 902 total number of nodes in the system after the recursive application of the resilience
 903 boosting procedure described in [Section 5](#). The idea is to use a union bound over all
 904 levels of recursion, nodes, and considered rounds, to show that the sampling succeeds
 905 with high probability in all cases. For the randomised variant of [Theorem 1.1](#), we will
 906 require the following additional constraint: when constructing a counter on n nodes,

907 the total number of failures is bounded by $f < \frac{n}{3+\gamma}$, where $\gamma > 0$ is constant.

908 This allows us to construct *probabilistic synchronous c-counters* in the sense that
 909 we say that the counter stabilises in time T , if for each round $t \geq T$ all non-faulty
 910 nodes count correctly with probability $1 - \eta^{-k}$.

911 **7.3. Sampling Communication Channels.** As discussed, there are two steps
 912 in the construction of [Theorem 4.1](#) where we rely on broadcasting: (1) the majority
 913 voting scheme for electing a leader block and counter, and (2) the execution of the
 914 phase king protocol. For the sake of clarity, we only focus on modifying the basic
 915 algorithm, where the nodes broadcast their entire state each round. We start with a
 916 sampling lemma we use for both steps. First, recall the following concentration bound
 917 for the sum of independent random binary variables:

918 **LEMMA 7.1** (Chernoff's bound). *Let $X = \sum X_i$ be a sum of independent random*
 919 *variables $X_i \in \{0, 1\}$. Then for $0 < \delta < 1$,*

$$920 \quad \Pr[X \leq (1 - \delta) \mathbf{E}[X]] \leq \exp\left(-\frac{\delta^2}{2} \mathbf{E}[X]\right).$$

921 **LEMMA 7.2.** *Let $U \subseteq V$ be a non-empty set of nodes such that the fraction of faulty*
 922 *nodes in U is strictly less than $1/(3 + \gamma)$. Suppose we sample K nodes v_0, \dots, v_{K-1}*
 923 *uniformly at random from the set U . For a given local variable $x(\cdot, r)$ encoded in the*
 924 *nodes' local state on round $r \geq 0$ and a value y , define the random variable*

$$925 \quad X_i = \begin{cases} 1 & \text{if } x(v_i, r) = y \text{ and } v_i \notin \mathcal{F}, \\ 0 & \text{otherwise} \end{cases}$$

926 for each $i \in [K]$ and let $X = \sum_{i=0}^{K-1} X_i$ be the number of y values sampled from correct
 927 nodes. There exists $K_0(\eta, k, \gamma) = \Theta(\log \eta)$ such that $K \geq K_0$ implies the following
 928 with high probability:

- 929 (a) If $x(u, r) = y$ for all $u \in U \setminus \mathcal{F}$, then $X \geq 2K/3$.
 930 (b) If a majority of nodes $u \in U \setminus \mathcal{F}$ have $x(u, r) = y$, then $X \geq K/3$.
 931 (c) If $X \geq 2K/3$, then $|\{x(u, r) = y : u \in U \setminus \mathcal{F}\}| \geq |U \setminus \mathcal{F}|/2$.

932 *Proof.* Define $\delta = 1 - \frac{2}{3} \cdot \frac{3+\gamma}{2+\gamma}$ and let $\rho < 1/(3 + \gamma)$ be the fraction of faulty nodes
 933 in U .

- 934 (a) If all correct nodes $u \in U \setminus \mathcal{F}$ agree on value $x(u, r) = y$, then

$$935 \quad \mathbf{E}[X] = (1 - \rho) K > \frac{2 + \gamma}{3 + \gamma} K.$$

936 As δ satisfies $(1 - \delta) \mathbf{E}[X] > 2K/3$, it follows from Chernoff's bound that

$$937 \quad \Pr\left[X < \frac{2K}{3}\right] \leq \Pr[X < (1 - \delta) \mathbf{E}[X]] \\ 938 \quad \leq \exp\left(-\frac{\delta^2}{2} \mathbf{E}[X]\right) \\ 939 \quad \leq \exp\left(-\delta^2 \frac{2 + \gamma}{2(3 + \gamma)} K\right). \\ 940$$

941 If $K_0(\eta, k, \gamma) = \Theta(\log \eta)$ is sufficiently large, $K \geq K_0(\eta, k, \gamma)$ implies that this proba-
 942 bility is bounded by η^{-k} .

943 (b) If a majority of non-faulty nodes u have value $x(u, r) = y$, then

$$944 \quad \mathbf{E}[X] \geq \frac{1}{2}(1 - \rho)K > \frac{1}{2} \cdot \frac{2 + \gamma}{3 + \gamma}K.$$

945 As above, by picking the right constants and using concentration bounds, we get that

$$\begin{aligned} 946 \quad \Pr \left[X \leq \frac{K}{3} \right] &\leq \Pr[X < (1 - \delta) \mathbf{E}[X]] \\ 947 \quad &\leq \exp \left(-\frac{\delta^2}{2} \mathbf{E}[X] \right) \\ 948 \quad &\leq \exp \left(-\delta^2 \frac{2 + \gamma}{4(3 + \gamma)} K_0 \right) \leq \eta^{-k}. \\ 949 \end{aligned}$$

950 (c) Suppose the majority of correct nodes have values different from y . Define

$$951 \quad \bar{X}_i = \begin{cases} 1 & \text{if } x(v_i, r) \neq y \text{ and } v_i \notin \mathcal{F}, \\ 0 & \text{otherwise.} \end{cases}$$

952 and $\bar{X} = \sum_{i=0}^{K-1} \bar{X}_i$ as the random variable counting the number of samples with
953 values different from y and arguing as for (b), we see that

$$954 \quad \Pr \left[X \geq \frac{2K}{3} \right] = \Pr \left[\bar{X} < \frac{K}{3} \right] \leq \eta^{-k},$$

956 where again we assume that $K_0(\eta, k, \gamma) = \Theta(\log \eta)$ is sufficiently large. Thus, $X \geq$
957 $2K/3$ implies with high probability that the majority of correct nodes have value y . \square

958 *Randomised Majority Voting.* Recall that in the majority voting scheme, there
959 are four local variables, two for each $i \in \{0, 1\}$, whose values depend directly on the
960 messages broadcast by all nodes:

- 961 • $m_i(v, r)$ stores the most frequent counter value in block i in round r , which
962 is determined from the broadcasted output variables of \mathbf{A}_i with ties broken
963 arbitrarily, and
- 964 • $M_i(v, r)$ stores the majority vote on $m_i(v, r - 1)$.

965 Throughout the remainder of this section, we let $K = \Theta(\log \eta)$ such that $K \geq K_0$
966 as given by [Lemma 7.2](#). Let $m_i^*(v, r)$ be the sampled version of $m_i(v, r)$; here the value
967 is determined by taking a random sample of size K from the set V_i . Analogously, the
968 variable $M_i^*(v, r)$ is determined by taking a random sample of size K from the set V
969 and taking the value that appears at least $2K/3$ times in the sample.

970 *Remark 7.3.* It holds that $f_i/n_i < 1/(3 + \gamma)$ for $i \in \{0, 1\}$.

971 **LEMMA 7.4.** *Suppose block $i \in \{0, 1\}$ is correct. Then for all $v \in V \setminus \mathcal{F}$ and*
972 *$r \geq T(\mathbf{A}_i)$, we have*

$$\begin{aligned} 973 \quad m_i^*(v, r) &= m_i(v, r) \\ 974 \quad M_i^*(v, r + 1) &= M_i(v, r + 1) \end{aligned}$$

976 *with high probability.*

977 *Proof.* To show the claim, we will apply [Lemma 7.2](#) with $U = V$ and $U = V_i$.
978 Before this, note that the fraction of faulty nodes in both V and V_i is less than $1/(3 + \gamma)$:

979 by assumption we have $f/n < 1/(3 + \gamma)$ and by Remark 7.3 yields $f_i/n_i < 1/(3 + \gamma)$.
 980 Thus, in both cases, we satisfy the first condition of Lemma 7.2.

981 For the claim regarding variable m_i , we apply Lemma 7.2 with $U = V_i$, that
 982 is, sample the subset $V_i \subseteq V$ consisting of nodes in block i . Since $|V_i| = n_i$ and i
 983 is a correct block, the set V_i contains at most f_i faulty nodes and all correct nodes
 984 output the same value $y \in [c_i]$, as \mathbf{A}_i has stabilised by round $r \geq T(\mathbf{A}_i)$. Moreover,
 985 $f_i/n_i < \frac{1}{3+\gamma}$ by Remark 7.3, so statement (a) of Lemma 7.2 yields that with high
 986 probability at least a fraction of $2/3$ of the sampled nodes output y .

987 To show the claim for variable M_i^* , note that by the previous case, $m_i^*(v, r) =$
 988 $m_i(v, r)$ holds for all correct nodes v with high probability. Applying Statement (a) of
 989 Lemma 7.2 to the set V and variable $m_i^*(v, r)$, we get that at least a fraction of $2/3$ of
 990 the samples have the same value. \square

991 From Lemma 7.4 it follows that we get probabilistic—in the sense that the claims
 992 hold with high probability—variants of Lemma 4.4, Lemma 4.5, and Lemma 4.6.
 993 These, in turn, yield the following probabilistic variant of Corollary 4.7.

994 COROLLARY 7.5. *There is a round $r = T + O(f)$ so that for all $v, w \in V \setminus \mathcal{F}$ with*
 995 *high probability it holds that*

- 996 1. $d(v, r) = d(w, r)$ and
- 997 2. for all $r' \in \{r + 1, \dots, r + \tau - 1\}$ we have $d(v, r') = d(v, r' - 1) + 1 \pmod{\tau}$.

998 *Randomised Phase King.* To obtain a randomised variant of the phase king
 999 algorithm, we modify the threshold votes used in the algorithm as follows. Instead of
 1000 checking whether at least $n - f$ of all messages have the same value, we check whether
 1001 at least a fraction of $2/3$ of the sampled messages have the same value. Similarly, when
 1002 checking for at least $f + 1$ values, we check whether a fraction of $1/3$ of the sampled
 1003 messages have this value.

1004 As a corollary, we get that when using the sampling scheme in the pulling model,
 1005 the execution of the phase king essentially behaves as in the deterministic broadcast
 1006 model.

1007 COROLLARY 7.6. *When executing the randomised variant of the phase king protocol*
 1008 *from Section 4 for $\eta^{O(1)}$ rounds, the statements of Lemma 4.8 and Lemma 4.9 hold*
 1009 *with high probability.*

1010 *Proof.* The modified phase king algorithm given in Section 4.3 uses two thresholds,
 1011 $n - f$ and $f + 1$. As discussed, these are replaced with threshold values of $2K/3$ and
 1012 $K/3$ when taking $K \geq K_0(\eta, k, \gamma)$ samples. Using the statements of Lemma 7.2, we
 1013 can argue analogously to the proofs of Lemma 4.8 and Lemma 4.9.

1014 First, to see that Lemma 4.8 holds with high probability, note that from statements
 1015 (b) and (c) of Lemma 7.2, it follows that if a node samples $2K/3$ times value y , then
 1016 w.h.p. other nodes sample at least $K/3$ times the same value (that is, we get the
 1017 probabilistic version of Lemma 4.3). Now we can follow the same reasoning as in
 1018 Lemma 4.8.

1019 Similarly, it is straightforward to check that Lemma 4.9 holds with high probability:
 1020 if all correct nodes agree on $a(\cdot)$, then all correct nodes sample at least $2K/3$ times
 1021 the same value w.h.p. by statement (a) of Lemma 7.2. Thus, analogously as in the
 1022 proof of Lemma 4.9, we get that the agreement persists when executing I_{3k} , I_{3k+1} , or
 1023 I_{3k+2} with high probability.

1024 Finally, we can apply the union bound over all $\eta^{O(1)}$ rounds and samples taken by
 1025 correct nodes ($n - f \leq \eta$ per round), that is, in total over $\eta^{O(1)}$ events. By choosing
 1026 large enough $k = O(1)$, we get that the claim holds with probability $1 - \eta^{-k}$. \square

1027 **7.4. Randomised Resilience Boosting.** It remains to formulate the proba-
 1028 bilistic variant of [Theorem 4.1](#). To this end, define $\mathcal{P}(n, f, c, \eta, k)$ as the family of
 1029 probabilistic synchronous c -counters on n nodes of resilience f . Here, probabilistic
 1030 means that an algorithm $\mathbf{P} \in \mathcal{P}(n, f, c, \eta, k)$ with stabilisation time $T(\mathbf{P})$ merely
 1031 guarantees that it counts correctly with probability $1 - \eta^{-k}$ in any given round
 1032 $t \geq T(\mathbf{P})$.

1033 Let $P(\mathbf{P})$ denote the number of messages pulled *per node* by a probabilistic
 1034 counter $\mathbf{P} \in \mathcal{P}(n, f, c, \eta, k)$. For any deterministic algorithm $\mathbf{A} \in \mathcal{A}(n, f, c)$, we define
 1035 $P(\mathbf{A}) = n$.

1036 **THEOREM 7.7.** *Let $c, n > 1$ and $f < n/(3 + \gamma)$, where $\gamma > 0$ and $n \leq \eta$. Define*
 1037 $n_0 = \lfloor n/2 \rfloor$, $n_1 = \lceil n/2 \rceil$, $f_0 = \lfloor (f - 1)/2 \rfloor$, $f_1 = \lceil (f - 1)/2 \rceil$ and $\tau = 3(f + 2)$. *If for*
 1038 $i \in \{0, 1\}$ *there exist synchronous counters $\mathbf{A}_i \in \mathcal{A}(n_i, f_i, c_i)$ such that $c_i = 3^i \cdot 2\tau$, then*
 1039 *for any sufficiently large $k = O(1)$, there exists a probabilistic synchronous c -counter*
 1040 $\mathbf{B} \in \mathcal{P}(n, f, c, \eta, k)$ *that*

- 1041 • *stabilises in $T(\mathbf{B}) = \max\{T(\mathbf{A}_0), T(\mathbf{A}_1)\} + O(f)$ rounds,*
- 1042 • *has state complexity of $S(\mathbf{B}) = \max\{S(\mathbf{A}_0), S(\mathbf{A}_1)\} + O(\log f + \log c)$ bits,*
 1043 *and*
- 1044 • *each node pulls at most $P(\mathbf{B}) = \max\{P(\mathbf{A}_0), P(\mathbf{A}_1)\} + O(\log \eta)$ messages per*
 1045 *round.*

1046 *Proof.* The proof proceeds analogously to the proof of [Theorem 4.1](#). First, we
 1047 apply [Corollary 7.5](#) to get a round counter that works once in a while with high
 1048 probability. We can then use this to clock the randomised phase king and [Corollary 7.6](#)
 1049 implies that the new output counter will reach agreement in $O(f)$ rounds with high
 1050 probability. The time and state complexities are as in the proof of [Theorem 4.1](#).

1051 To analyse the number of pulls, observe that in [Lemma 7.4](#) each node samples
 1052 twice $K = O(\log \eta)$ messages (from both V_0 and V_1) and [Corollary 7.6](#) samples $O(\log \eta)$
 1053 messages from all the nodes. Thus, in total, a node $v \in V_i$ samples $O(\log \eta)$ messages
 1054 in addition to the messages pulled when executing \mathbf{A}_i . \square

1055 Note that we can choose to replace $\mathbf{A} \in \mathcal{A}(n, f, c)$ by $\mathbf{Q} \in \mathcal{P}(n, f, c, \eta, k)$ when
 1056 applying this theorem, arguing that with high probability it *behaves* like a corresponding
 1057 algorithm $\mathbf{A} \in \mathcal{A}(n, f, c)$ for polynomially many rounds. Furthermore, note that it is
 1058 also possible to boost the probability of success, and thus the period of stability, by
 1059 simply increasing the sample size. For instance, sampling $\text{polylog } \eta$ messages yields
 1060 an error probability of $\eta^{-\text{polylog } \eta}$ in each round, whereas in the extreme case, by
 1061 “sampling” all nodes the algorithm reduces to the deterministic case.

1062 Using [Theorem 7.7](#) recursively as in [Section 5](#) for $O(\log f)$ steps, we get the
 1063 following result.

1064 **THEOREM 7.8.** *For any integers $c, n > 1$, $f < n/(3 + \gamma)$, there exists an f -resilient*
 1065 *probabilistic synchronous c -counter that runs on n nodes, requires $O(\log^2 f + \log c)$ bits*
 1066 *to encode the state of a node, has each node pull $O(\log f \log n)$ messages per round,*
 1067 *and stabilises in $O(f)$ rounds with probability $1 - n^{-k}$, where $k > 0$ is a freely chosen*
 1068 *constant.*

1069 **7.5. Oblivious Adversary.** Finally, we remark that under an *oblivious adver-*
 1070 *sary*, that is, an adversary that picks the set of faulty nodes independently of the
 1071 randomness used by the non-faulty nodes, we get *pseudorandom* synchronous counters
 1072 satisfying the following: (1) the execution stabilises with high probability and (2) if
 1073 the execution stabilises, then all non-faulty nodes will deterministically count correctly.
 1074 Put otherwise, we can fix the random bits used by the nodes to sample the communica-

1075 tion links *once*, and with high probability we sample sufficiently many communication
 1076 links to non-faulty nodes for the algorithm to (deterministically) stabilise. This gives
 1077 us the following result.

1078 **COROLLARY 7.9.** *For any integers $c, n > 1$, $f < n/(3 + \gamma)$, there exists a pseudo-*
 1079 *random synchronous c -counter with resilience f against an oblivious fault pattern that*
 1080 *runs on n nodes, requires $O(\log^2 f + \log c)$ bits to encode the state of a node, has each*
 1081 *node pull $O(\log f \log n)$ messages per round, and stabilises in $O(f)$ rounds.*

1082 **8. Conclusions.** In this work, we showed that there exist algorithms for syn-
 1083 chronous counting that (1) are deterministic, (2) tolerate the optimal number of
 1084 faults, (3) have asymptotically optimal stabilisation time, and (4) need to store *and*
 1085 communicate a very small number of bits between consecutive rounds—something no
 1086 prior algorithms have been able to do.

1087 In addition, we discussed two complementary approaches on how to further reduce
 1088 the total number of communicated bits in the network. The first one is a deterministic
 1089 construction that lets the nodes communicate only few bits after stabilisation, in order
 1090 to verify that stabilisation has occurred and that the counters agree. The construction
 1091 retains all properties (1)–(4), and in particular, when constructing polynomially-sized
 1092 counters with linear resilience, the algorithm communicates an asymptotically optimal
 1093 number of bits after stabilisation.

1094 The second technique for reducing the amount of communication is based on
 1095 random sampling of communication channels. Here, we employed randomisation so
 1096 that each node needs to communicate only with $\text{polylog } n$ instead of $n - 1$ other nodes
 1097 in the system, thus reducing the number of messages sent from $\Theta(n^2)$ to $\Theta(n \text{ polylog } n)$.
 1098 The trade-off here is that the resulting algorithm has *slightly* suboptimal resilience
 1099 of $f < n/(3 + \gamma)$, where $\gamma > 0$ is a constant, and is merely guaranteed to work for
 1100 polynomially many rounds with high probability before a new stabilisation phase is
 1101 required. The latter issue disappears when employing pseudorandomness. In this case,
 1102 one may simply fix a random topology and the algorithm will not fail again after
 1103 stabilisation; naturally, this necessitates that the Byzantine faulty nodes are chosen in
 1104 an oblivious manner, i.e., independently of the topology.

1105 We can also combine both techniques to attain probabilistic counters that dur-
 1106 ing stabilisation communicate $\Theta(n \text{ polylog } n)$ bits each round and after stabilisation
 1107 asymptotically optimal $O(1)$ bits every $\Theta(n)$ rounds.

1108 To conclude the paper, we now wish to highlight some interesting problems that
 1109 still remain open:

- 1110 Q1. Our solutions are not adaptive (as defined in [23]), as their stabilisation time
 1111 is not bounded by a function of the number of *actual* permanent faults. Can
 1112 this be achieved?
- 1113 Q2. Are there algorithms that satisfy (1)–(3), but need to store and communicate
 1114 substantially fewer than $\log^2 f$ bits? This question has been partially answered
 1115 in follow-up work [25], showing that $O(\log f)$ bits suffice. However, no non-
 1116 trivial lower bound is known, so it remains open whether $o(\log f)$ bits suffice.
- 1117 Q3. Can the ideas presented in this paper be applied to *randomised* consensus
 1118 routines in order to achieve sublinear stabilisation time with high resilience
 1119 and small communication overhead? Again, a partial answer is provided in [25]:
 1120 this is possible, but the given solutions may still fail *after* stabilisation (with
 1121 a very small probability per round). The question thus remains open w.r.t.
 1122 the original problem definition, which requires that after stabilisation the
 1123 algorithm keeps counting correctly.

1124 Finally, we point out that the recursive approach we employ in this paper can be
 1125 interpreted as an extension of its similar use in synchronous consensus routines [5, 6],
 1126 where the shared round counter is implicitly given by the synchronous start.

1127 Q4. Can a similar recursive approach also be used for deriving improved *pulse*
 1128 *synchronisation* [14, 18] algorithms?

1129 Interestingly, no reduction from consensus to pulse synchronisation is known, so there
 1130 is still hope for efficient deterministic pulse synchronisation algorithms that stabilise
 1131 in sublinear time.

1132 **Acknowledgements.** We thank all the anonymous reviewers for helpful com-
 1133 ments.

1134 REFERENCES

- 1135 [1] M. AJTAI AND N. LINIAL, *The influence of large coalitions*, *Combinatorica*, 13 (1993), pp. 129–145.
 1136 doi:10.1007/BF01303199.
- 1137 [2] A. ARORA, S. DOLEV, AND M. G. GOUDA, *Maintaining digital clocks in step*, *Parallel Processing*
 1138 *Letters*, 1 (1991), pp. 11–18.
- 1139 [3] B. AWERBUCH, S. KUTTEN, Y. MANSOUR, B. PATT-SHAMIR, AND G. VARGHESE, *A time-optimal*
 1140 *self-stabilizing synchronizer using a phase clock*, *IEEE Transactions on Dependable and*
 1141 *Secure Computing*, 4 (2007), pp. 180–190.
- 1142 [4] M. BEN-OR, D. DOLEV, AND E. N. HOCH, *Fast self-stabilizing Byzantine tolerant digital clock*
 1143 *synchronization*, in *Proc. 27th Annual ACM Symposium on Principles of Distributed*
 1144 *Computing (PODC 2008)*, ACM Press, 2008, pp. 385–394. doi:10.1145/1400751.1400802.
- 1145 [5] P. BERMAN, J. A. GARAY, AND K. J. PERRY, *Bit optimal distributed consensus*, in *Computer*
 1146 *Science: Research and Applications*, Springer, pp. 313–321. doi:10.1007/978-1-4615-3422-8-
 1147 27.
- 1148 [6] P. BERMAN, J. A. GARAY, AND K. J. PERRY, *Towards optimal distributed consensus*, in *Proc.*
 1149 *30th Annual Symposium on Foundations of Computer Science (FOCS 1989)*, IEEE, 1989,
 1150 pp. 410–415. doi:10.1109/SFCS.1989.63511.
- 1151 [7] L. BOCZKOWSKI, A. KORMAN, AND E. NATALE, *Minimizing message size in stochastic com-*
 1152 *munication patterns: Fast self-stabilizing protocols with 3 bits*, in *Proc. 28th Annual*
 1153 *ACM-SIAM Symposium on Discrete Algorithms (SODA 2017)*, SIAM, 2017, pp. 2540–2559.
 1154 doi:10.1137/1.9781611974782.168.
- 1155 [8] C. BOULINIER, F. PETIT, AND V. VILLAIN, *Synchronous vs. asynchronous unison*, *Algorithmica*,
 1156 51 (2008), pp. 61–80. doi:10.1007/s00453-007-9066-x.
- 1157 [9] C. DELPORTE-GALLET, S. DEVISMES, AND H. FAUCONNIER, *Robust stabilizing leader election*, in
 1158 *Proc. 9th Symposium on Stabilization, Safety, and Security of Distributed Systems*, 2007,
 1159 pp. 219–233.
- 1160 [10] S. DEVISMES, T. MASUZAWA, AND S. TIXEUIL, *Communication efficiency in self-stabilizing silent*
 1161 *protocols*, in *Proc. 29th Conference on Distributed Computing Systems (ICDCS)*, 2009,
 1162 pp. 474–481.
- 1163 [11] D. DOLEV, *The Byzantine generals strike again*, *Journal of Algorithms*, 3 (1982), pp. 14–30.
- 1164 [12] D. DOLEV, M. FÜGGER, C. LENZEN, U. SCHMID, AND A. STEININGER, *Fault-tolerant distributed*
 1165 *systems in hardware*, *Bulletin of the EATCS*, (2015). [http://bulletin.eatcs.org/index.php/](http://bulletin.eatcs.org/index.php/beatcs/issue/view/18)
 1166 [beatcs/issue/view/18](http://bulletin.eatcs.org/index.php/beatcs/issue/view/18).
- 1167 [13] D. DOLEV, K. HELJANKO, M. JÄRVISALO, J. H. KORHONEN, C. LENZEN, J. RYBICKI, J. SUOMELA,
 1168 AND S. WIERINGA, *Synchronous counting and computational algorithm design*, 2015. arXiv:
 1169 1304.5719v2.
- 1170 [14] D. DOLEV AND E. N. HOCH, *On self-stabilizing synchronous actions despite Byzantine at-*
 1171 *tacks*, in *Proc. 21st International Symposium on Distributed Computing (DISC 2007)*,
 1172 vol. 4731 of *Lecture Notes in Computer Science*, Springer, 2007, pp. 193–207. doi:
 1173 10.1007/978-3-540-75142-7_17.
- 1174 [15] D. DOLEV, J. H. KORHONEN, C. LENZEN, J. RYBICKI, AND J. SUOMELA, *Synchronous counting*
 1175 *and computational algorithm design*, in *Proc. 15th International Symposium on Stabilization,*
 1176 *Safety, and Security of Distributed Systems (SSS 2013)*, vol. 8255 of *Lecture Notes in*
 1177 *Computer Science*, Springer, 2013, pp. 237–250. doi:10.1007/978-3-319-03089-0_17. arXiv:
 1178 1304.5719v1.
- 1179 [16] D. DOLEV AND R. REISCHUK, *Bounds on information exchange for Byzantine agreement*, *Journal*

- 1180 of the ACM, 32 (1985), pp. 191–204. doi:10.1145/2455.214112.
- 1181 [17] S. DOLEV, *Self-Stabilization*, The MIT Press, Cambridge, MA, 2000.
- 1182 [18] S. DOLEV AND J. L. WELCH, *Self-stabilizing clock synchronization in the presence of Byzantine*
- 1183 *faults*, Journal of the ACM, 51 (2004), pp. 780–799. doi:10.1145/1017460.1017463.
- 1184 [19] S. DUBOIS, M. POTOP-BUTUCARU, M. NESTERENKO, AND S. TIXEUIL, *Self-stabilizing byzantine*
- 1185 *asynchronous unison*, Journal of Parallel and Distributed Computing, 72 (2012), pp. 917–
- 1186 923.
- 1187 [20] M. J. FISCHER AND N. A. LYNCH, *A lower bound for the time to assure interactive consistency*,
- 1188 Information Processing Letters, 14 (1982), pp. 183–186. doi:10.1016/0020-0190(82)90033-3.
- 1189 [21] M. G. GOUDA AND T. HERMAN, *Stabilizing unison*, Information Processing Letters, 35 (1990),
- 1190 pp. 171–175.
- 1191 [22] E. HOCH, D. DOLEV, AND A. DALIOT, *Self-stabilizing Byzantine digital clock synchronization*,
- 1192 in Proc. 8th International Symposium on Stabilization, Safety, and Security of Distributed
- 1193 Systems (SSS 2006), vol. 4280, 2006, pp. 350–362.
- 1194 [23] S. KUTTEN AND B. PATT-SHAMIR, *Adaptive stabilization of reactive protocols*, in Proc. 24th
- 1195 Conference on Foundations of Software Technology and Theoretical Computer Science
- 1196 (FSTTCS), 2005, pp. 396–407.
- 1197 [24] C. LENZEN AND J. RYBICKI, *Efficient counting with optimal resilience*, in Proc. 29th International
- 1198 Symposium on Distributed Computing (DISC 2015), Springer, 2015, pp. 16–30. doi:
- 1199 10.1007/978-3-662-48653-5_2.
- 1200 [25] C. LENZEN AND J. RYBICKI, *Near-optimal self-stabilising counting and firing squads*, in Proc.
- 1201 18th Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS), 2016,
- 1202 pp. 263–280.
- 1203 [26] C. LENZEN, J. RYBICKI, AND J. SUOMELA, *Towards optimal synchronous counting*, in Proc. 34th
- 1204 Annual ACM Symposium on Principles of Distributed Computing (PODC 2015), ACM
- 1205 Press, 2015, pp. 441–450. doi:10.1145/2767386.2767423.
- 1206 [27] M. C. PEASE, R. E. SHOSTAK, AND L. LAMPORT, *Reaching agreement in the presence of faults*,
- 1207 Journal of the ACM, 27 (1980), pp. 228–234. doi:10.1145/322186.322188.
- 1208 [28] T. TAKIMOTO, F. OOSHITA, H. KAKUGAWA, AND T. MASUZAWA, *Communication-efficient self-*
- 1209 *stabilization in wireless networks*, in Proc. 14th Conference on Stabilization, Safety, and
- 1210 Security of Distributed Systems (SSS), 2012, pp. 1–15.