

1 Lower bounds for authenticated randomized 2 Byzantine consensus under (partial) synchrony 3 the limits of standalone digital signatures

4 — Abstract —

5 This paper deals with lower bounds for *randomized* consensus against any *standard* adaptive
6 Byzantine Environment, that is, which can corrupt players at any time, up to a total of t , but *not*
7 delete messages sent. We assume that players can *sign* their messages, so that the Environment
8 cannot forge a message on behalf of a honest player. This assumption is known as the “authenticated
9 setting” and studied e.g. in the classical works of Pease et al, DLS, the lower bounds of Dolev-
10 Strong/Reischuk, Hadzilacos-Halpern etc. We rename it instead “standalone digital signatures”, not
11 to confuse it with the weaker model of plain authenticated channels. We do not make any further
12 assumption. By contrast, recall that Chen-Micali (Algorand) and Ouroboros Praos (EC’18) assume
13 the ability to publish public keys on a board (registered PKI), plus a global coin (a.k.a. CRS: the
14 hash of the previous block, or assuming a random value contained in the genesis block); alternatively:
15 Abraham et al (PODC’19) assume a trusted PKI setup. We show that:

- 16 ■ Under all partially synchronous models, especially those where messages are *never* lost, for
17 $n \geq 3t + 1$ players. Then (Theorem 4) we have that, with nonnegligible probability, forever
18 honest players need sending or forwarding a *quadratic* $\Omega(nt)$ number of *signatures* after the
19 moment when the network becomes synchronous (a.k.a. “global stabilization time” GST). This
20 also carries over the variant model where the maximum network delay is a priori unknown.
21 Particularly, this is not a consequence of the synchronous one of Abraham et al (PODC’19),
22 whose “rushing” Environment *can* in addition delete messages that were sent. This establishes a
23 *separation* with the synchronous setting, where King-Saia (PODC’10) exhibited a protocol in
24 $O(n\sqrt{n})$ bit complexity of communications (we assume like them that messages are confidential),
25 which furthermore does not even use digital signatures.
- 26 ■ Under the same model (Theorem 6): consider protocols in which so-far honest players send
27 a total *number* of messages after GST which is *linear* in the number n of players, then the
28 number of rounds after GST before every player outputs is at least in $\log(t)$. This establishes a
29 separation with the aforementioned message-linear and constant-round algorithms, which make
30 further assumptions to achieve random committee/leader sortition.
- 31 ■ Under *synchrony*, when honest players only *multicast* messages, then (Theorem 9) we have that
32 with nonnegligible probability, so-far honest players need sending a total *quadratic* number of
33 *messages*. This strictly reinforces Abraham et al’s (PODC’19) multicast lower bound, whose
34 argument does not handle the case where messages carry digital signatures. This makes a
35 separation with the aforementioned protocols, which also proceed by multicasts.

36 **2012 ACM Subject Classification** Computing methodologies → Distributed algorithms

37 **Keywords and phrases** Author: Please fill in \keywords macro

38 **Digital Object Identifier** 10.4230/LIPIcs...

39 **1** Model, motivations and results

40 **Confidential message-passing with public metadata** We consider players $\mathcal{P} := P_1, \dots, P_n$,
41 which are probabilistic machines whose identities are public and fixed. They are linked with
42 pairwise authenticated confidential communication channels. In particular, players receiving
43 a message automatically learn the identity of the sender. The Environment \mathcal{E} is a polynomial
44 probabilistic machine, to which the network leaks only the identity of the sender and of the
45 recipient of ever message sent. \mathcal{E} initializes every honest player with a binary input value.



:2 Lower bounds, limits of standalone digital signatures

46 **Standard adaptive Byzantine corruptions** The Environment can corrupt any honest player
47 at any time, up to a total of t corruptions. Corrupted players are *malicious*, in that they
48 share all their state with the Environment and are totally controlled by it. In Definitions 1
49 and 2 below we distinguish actions performed by *so far honest* players, that is, sent while they
50 are still honest. That is, we allow them to become corrupted later, as in [1, Theorem 3]. By
51 contrast, when restricting to *forever honest* players, we count only the messages/signatures
52 sent or forwarded by honest players who never become corrupted later in the execution.

53 Contrary to the lower bound [1, Theorem 1] (PODC'19), we do *not* assume that the
54 Environment is rushing: it cannot delete any message sent.

55 **Randomized** has a twofold meaning. *First*, each player has access to a *local* random source,
56 as [3]. But players do *not* have access for free to a global coin (that is: a public unpredictable
57 value, as known as “global coin” or CRS) as in Rabin [20] or [10]. *Second*, the aforementioned
58 specifications of consensus are relaxed, in that we can tolerate some negligible probability of
59 safety or liveness failure (as in [7, 1]), which we will specify.

60 **Partial synchrony [14, §2.2 (3)] and termination condition** the Environment controls the
61 delivery delay of every message sent on the network. It can deliver instantaneously messages
62 between corrupted players, and arbitrarily delay any message sent. The Environment may
63 however secretly commit, at some point called Global stabilization time (GST), to deliver all
64 messages sent, *even those previously delayed*, within a fixed delay Δ . The value of Δ is fixed
65 by the Environment at the beginning of the execution and made public to the players. We
66 assume that players have a global clock which enables them to communicate in the form of
67 successive rounds of duration Δ . Termination of consensus is required only *after* GST, if it
68 ever happens. Notice that the Environment may well never set GST and still schedule the
69 delivery of messages within Δ , as if synchrony held. Then, players playing consensus would
70 anyway also ultimately output, since these conditions are indistinguishable from if GST
71 did happen. Finally our communication lower bounds are interesting only in such partially
72 synchronous models where messages *cannot* be lost. When messages can be lost before GST,
73 then [14, §4.2 Remark 2] notice for any 1-resilient consensus algorithm, there always exists *a*
74 *run in which a player never stops sending messages*, even after GST.

75 **Partial synchrony [14, (2)] and termination condition** In this model, messages are always
76 delivered within a delay Δ that is set by the Environment at the beginning of the execution,
77 but which is unknown to players. Let us introduce a possibly new termination condition. We
78 say that, in an execution, honest players “*guess the actual value*” of Δ if, from some point in
79 this execution, players play the protocol round by round with delay Δ between each round.
80 Termination of consensus is then required only to happen *after* honest players have guessed
81 the actual value of Δ , if this ever happens. This is restrictive, so makes the lower bounds
82 stronger. This could be formalized with unreliable detectors controlled by the Environment.

83 **a Consensus protocol** as formalized in Ben Or [3] or DLS [14], is such that we have:

84 **(Consistency)** no two honest players output different values;

85 **(Termination)** after a polynomial number of rounds after GST (respectively: a polynomial
86 time after all players have guessed Δ), every player ultimately outputs a value (one could
87 specify up to some probability, as we do in Theorem 9, and could be easily done in
88 Theorems 4 and 6);

89 **(Strong/Weak Unanimity)** in the Strong variant, if all forever honest players have the same
 90 input value then this value should be the one they output. In the Weak variant, this is
 91 required only if, in addition, all players are forever honest.

92 \mathcal{F}_{ds} **the standalone digital signatures assumption (a.k.a. the authenticated setting)** In
 93 this paper we consider the classical model as described concretely e.g. in Dolev-Reischuk [12]:
 94 “we allow the processors to share a signature scheme that enables each one to sign its messages
 95 so that every receiver will recognize them as being signed by it and no one can change the
 96 contents of a message or the signature undetectably. Such a scheme is the one suggested
 97 in [...] and its use for Byzantine algorithms is described in [...]. We allow faulty processors
 98 to collude for cheating. Therefore every message that contains only signatures of faulty
 99 processors can be produced by them.” This is explicit in the classical works [13], [19, §4], [20],
 100 [14, §2.2], [15, §7], and also all works in which only the possibility to authenticate the original
 101 issuer of the message is actually used: PBFT [8], Cachin et al [6, §2.3.1] (nonwithstanding
 102 aggregation of multiple signatures) or the recent Kogias et al [18]: “we assume that messages
 103 are authenticated in a sense that if an honest party P_i receives a message m indicating that
 104 m was sent by an honest party P_j , then m was indeed generated by P_j at some prior time.”
 105 This assumption is traditionally called the “authenticated” setting, but we prefer here the
 106 terminology *standalone digital signatures*, in order to avoid confusion with the weaker plain
 107 setting where only the communication channels are authenticated (see the discussion about
 108 [1, Theorem 3] in §1.1). Importantly, players *do not* have access to a trusted bulletin board
 109 that publishes their public keys matched with their identities (see §1.1.0.1 for this stronger
 110 model). Let us abstract out the model as follows.

111 *Messages* allowed by the network are constituted of concatenations of *two* well-separate
 112 datatypes: strings of bits; and **signatures** (so that a message can well consist in a single
 113 signature). **signatures** are produced and used as follows:

- 114 ■ On request $\text{sign}(m)$ from any player P , where m is a message: \mathcal{F}_{ds} outputs a *new signature*
 115 to P , that we denote $s_{P,i}(m)$ (i the number of previous identical requests) and adds
 116 $(m, s_{P,i}(m), \text{“signed by P”})$ to its internal registry.
- 117 ■ On request $\text{check}(m, s, P)$ from any player Q with m a message and s a **signature**: if the
 118 registry contains $(m, s, \text{“signed by P”})$, then \mathcal{F}_{ds} outputs **true**; and **false** otherwise.

119 **signatures** are objects which can be *copied*, but which are otherwise *immutable*. **signatures** all
 120 look like *identical* from the points of view of players: the *only* information that players can
 121 extract from a given **signature** is the output of the **check** request to \mathcal{F}_{ds} . New signatures can
 122 *only* be created by \mathcal{F}_{ds} , from the **sign** request. [To fix ideas in a more concrete way, one could
 123 consider that on every **sign** request, then the **signature** output is a uniform random string in
 124 $\{0, 1\}^\kappa$, with κ a security parameter. But since these strings are mutable, this would then
 125 require tedious discussions on what happens if a signature is sent into (possibly incomplete)
 126 pieces. The “number of signatures sent” would then be a fractional number, equal to the
 127 total quantity of information on these strings which is sent].

128 *In particular, the Environment can simulate a fairly sampled execution of a protocol in*
 129 *the \mathcal{F}_{ds} model, up to a polynomial time/number of rounds, without actually corrupting the*
 130 *players involved.*

131 We say that a s is a *valid signature* of P for some message m if and only if s is the output
 132 of a request $\text{sign}(m)$ from P ; and it is *invalid* otherwise. So by definition, $\text{check}(m, s, P)$
 133 returns **true** if and only if s is a valid signature of P on m . **signatures issued** by P are by
 134 definition those which are valid for P on some message m . the definition of \mathcal{F}_{ds} implies that
 135 the only way the Environment can obtain a true answer on a **check** query, is to have an

136 actual corrupted player make the query with an actual valid signature s from P on m . Said
 137 otherwise, the Environment cannot guess if some message m was signed by P by brute-forcing
 138 many $\text{check}(m, *, P)$ (which would not be tractable anyway since it is polynomial). In this
 139 model, what corresponds to a “signed message” in the concrete sense may be obtained as a
 140 message equal to a concatenation of the form $m' = m \| s$, where s is a valid signature on m .
 141 So by definition $\text{check}(m', s, P)$ returns false, since s is thus *not* a valid signature on m' .

142 Number of signatures and number of messages complexities (but not bit-complexity)

143 Let us define a measurement that formalizes, in our \mathcal{F}_{dis} model, the concrete [12]: “Since
 144 there exists an algorithm for reaching Byzantine Agreement without using signatures, the
 145 lower bound is meaningless unless we somehow count the messages that do not contain
 146 signatures. We make the technical assumption that every message in an authenticated
 147 algorithm carries at least the signature of its sender. Alternatively, the lower bound given
 148 below holds if one counts the number of signatures together with the number of messages
 149 without signatures.” For every player Q , in an execution e , let us denote $\text{Sent}_e(Q)$ the sum
 150 over the messages m' sent by Q : after GST (resp.: after players have guessed Δ), if he is
 151 (forever / so far honest), of:

- 152 ■ the number of signatures contained in m' , $\left[\text{Notice that if } m' = s_R(s_{Z_1}(m_1) \| \dots \| s_{Z_\ell}(m_\ell)), \right.$
 153 then by definition this number is 1 and not $\ell + 1$ because the signatures $s_{Z_j}(m_j)$ are *not*
 154 in the concatenation of objects forming m' (for instance if $j = 1$, notice that, concretely,
 155 someone receiving solely m would *not* be able to extract $s_{Z_1}(m_1)$ in any manner, nor
 156 even guess that this is the object signed by s_R , since she could not perform any successful
 157 check request without the $s_{Z_1}(m_1)$ at hand) $\left. \right]$
- 158 ■ plus 1 if m' is not of the form $m' = m \| s$, where s is a valid signature of Q on m (that is:
 159 if m' is “unsigned” in the concrete sense).

160 The rule of counting, as precised in the first item, makes our measurement comparable with
 161 the bit complexity (see §1.1): one signature in a message counts no more than one unit.
 162 The cost for this optimistic measurement is that this will require an additional “technical
 163 assumption” for the result of Theorem 4.

164 ► **Definition 1** (the “number of signatures sent” by (forever / so far) honest players after GST).
 165 is $\sum_{Q \in \mathcal{P}} \text{Sent}_e(Q)$.

166 The next measurement is (strictly) more optimistic than the *total number of messages*
 167 *sent* in the concrete sense of [12], since we count only once the possibly multiple messages
 168 sent from Q to R in a given round.

169 ► **Definition 2** (the “number of messages sent” by (forever / so far) honest players in an
 170 execution after GST). is the sum over all rounds r of the execution taking place after GST,
 171 over all (forever / so far) honest player Q , of all players R such that, in round r , Q sends
 172 at least one message to R on the network.

173 1.1 Overview of the results

174 **First** we are unaware of any lower bound in the literature specific to the communication
 175 complexity under partial synchrony. Recall that this model was introduced by Dwork et
 176 al [14] to circumvent the impossibility of fully asynchronous consensus. For instance, this
 177 setting is handled by popular state machine replication protocols such as Castro-Liskov’s [8]
 178 or Hotstuff [23] (PODC’19), which are both deterministic and with a static environment.
 179 Precisely, recall that termination in these protocols relies on timing assumptions: if a honest

180 leader does not benefit from a “synchronous enough” network, then he may be replaced
 181 before he could enforce a decision. Recall ([14, Theorem 4.4]) that authenticated Byzantine
 182 consensus under partial synchrony is possible only if $t < \frac{n}{3}$.

183 In Theorem 4 we show that in any randomized protocol for binary consensus resilient
 184 under standard adaptive Byzantine Environments, then there exists an Environment under
 185 which forever honest players always send a quadratic ($\Omega(nt)$) number of signatures after GST
 186 in any execution, in the sense of Definition 1. This is not a consequence of [1, Theorem 1]
 187 because in our model, the only possibility for the Environment to remove forever one message
 188 is to delay it forever and set $\text{GST}=\infty$, but then players are not required anymore to output,
 189 contrary to the synchronous [1, Theorem 1]. The proof combines [12, Theorem 1] (recalled
 190 in §A.2.1), the simulation argument of e.g. [1, Theorem 3], and in addition the ability of the
 191 Environment to delay until GST any message sent before. Let us derive the consequences.
 192 In a concrete protocol in the sense of [12], when Q forwards to R the signatures of x distinct
 193 players P_i , then this represents at least x bits sent from Q to R . This means that the bit
 194 complexity is at least as large as the message complexity, and thus is also *quadratic*. By
 195 comparison, the synchronous consensus protocol of King-Saia [17] has a bit complexity of
 196 $O(n\sqrt{n})$ which is *strictly lower*, while under a *strictly weaker* model (apart synchrony):
 197 same standard adaptive Environments and confidential messages, no setup (the global coins
 198 are emulated), while assuming *no digital signatures at all*. Thus, Theorem 4 establishes a
 199 *strict separation* between the bit complexity of partial synchrony and synchrony. The above
 200 argument shows that closing this gap would require to enrich the model with the possibility
 201 to *aggregate* multiple signatures into one message of constant ($O(1)$) bit complexity. The
 202 state of the art implementation is [21], but makes the strong assumption of a trusted setup
 203 (see below). Alternatively, the asynchronous [18] achieves this assuming only standalone
 204 digital signatures as here, but with a $O(n^4)$ bit complexity for setup. Finally, the closest
 205 solution to the King-Saia setting may be [5, §5.2], which aggregates batches of signatures of
 206 logarithmic sizes, assuming the intermediate “registered PKI” model (see 1.1.0.1).

207 **Second** achieving a communication size which is *linear* in the number of players against an
 208 adaptive environment, plus a *constant* expected number of rounds, was recently popularized by
 209 Chen-Micali [10] (Algorand), Ouroboros-Praos [11], Abraham et al [1, Theorem 2] (PODC’19)
 210 and Chan et al [16] (EC’19). These papers adapt consensus protocols to a setting where only
 211 a small random set of players talks in each round. Precisely, these random “committees” are
 212 sampled with the help verifiable random functions (VRF). These functions enable a player
 213 to prove that his secret signature key satisfies some rare condition (analogous to a hash with
 214 many zeros in Bitcoin), without revealing his secret key. When reading the last two papers
 215 quickly, it might seem that a Public Key Infrastructure (PKI) is sufficient to achieve such
 216 protocols. Indeed, a PKI is the sole assumption in their [1, Theorem 2], while [16, Theorem
 217 1.1] specifies only “standard bilinear group assumption”. But actually if corrupted players
 218 could *choose* their secret keys before the protocol starts, then they could brute force it, and
 219 obtain membership in many committees. This would e.g. ruin the consensus of [1, §6.2] (to
 220 stick with the partially synchronous setting). This is why these papers actually assume that
 221 a *trusted* dealer gives the keys to the players (mentioned as the “Trusted PKI setup” page
 222 19 of [16]). Another way around consists in assuming that an unpredictable random string
 223 is learned *after* the keys are set up. In [10] this is assumed to be satisfied by the hash of
 224 the previous block. In Praos [11] this is formalized as the ideal functionality \mathcal{F}_{INIT} page
 225 10, that creates a public random string $\eta \leftarrow \{0, 1\}^\lambda$ (in a “genesis block”). This has to be
 226 done everytime new players receive a big amount of stake. A CRS actually also turns out to

227 be assumed in [16, p19]. To be sure, although Hotstuff [23] has linear bit complexity, recall
 228 that termination there is also conditioned to a honest leader, which could be adaptively
 229 corrupted in our setting. Let alone that this protocol satisfies only Weak unanimity.

230 By contrast we show in Theorem 6 that, assuming standalone digital signatures only,
 231 then, for every protocol in which the number of *messages* sent by so-far honest (in the sense
 232 of Definition 2) players after GST is *linear*, then there exists an Environment under which the
 233 number of rounds before termination is at least *logarithmic*. The proof follows the pattern of
 234 the previous, but this time, the baseline Environments also adaptively corrupt players who
 235 receive many signatures. Let us point the related randomized lower bound of Bar Joseph
 236 and Ben Or [2] in $\Omega(\sqrt{n})$ rounds, but assuming that the Environment *knows* the full internal
 237 state of players, including the local randomness they sample before they send their messages.
 238 Notice that Theorem 6 considers the number of messages complexity, which is the most
 239 optimistic one. For instance, sending an aggregated signature anyway counts as one message.
 240 Considering existence of the previous committee-based algorithms, a direct consequence of
 241 Theorem 6 is:

242 ► **Corollary 3.** *The functionality \mathcal{F}_{mine} , as defined in [1, §C.2] and which returns eligibility*
 243 *of a vote, cannot be implemented in constant round and linear communication complexity*
 244 *under standalone digital signatures, even assuming aggregated signatures.*

245 **Third** recall the *synchronous* lower bound of [1, Theorem 3], that states that any broadcast
 246 protocol in which players *multicast* messages only, and resilient against a standard adaptive
 247 Environment, has *quadratic* message complexity. The problem is that this is introduced
 248 by “we additionally investigate whether the remaining setup PKI assumption is necessary”,
 249 which seems to refer to the *trusted setup model* of their [1, Theorem 2], whereas their lower
 250 bound *does not* even assumes digital signatures (“plain authenticated channels” only). We do
 251 not see why ruling out this weak model would make necessary the very strong assumption of
 252 a trusted setup. So in this paper we investigate the question for the intermediate standalone
 253 digital signatures. In §A.2.2 we revisit the proof of [1, Theorem 3], which we show would fail
 254 if assuming digital signatures on the messages sent or forwarded by honest players. We then
 255 show in Theorem 9 that this quadratic lower bound does also hold *even if* assuming digital
 256 signatures, for the problem of consensus, .

257 1.1.0.1 Extension to stronger models

258 The one called *registered PKI* in Boyle-Cohen-Goel [5, p9], for which they refer to e.g.
 259 Boldyreva [4] (PKC’03), is strictly stronger than our “standalone digital signatures”. There,
 260 players can publish a public key on a public bulletin board before the execution starts
 261 (before they receive their input), which makes it known to all players at the beginning of the
 262 execution. From this, and additional cryptographic assumptions, players can exhibit random
 263 strings provably derived from the corresponding secret key, without disclosing it. This is
 264 leveraged in [11] and in [10] (which considers furthermore hashes of previous blocks as an
 265 unpredictable CRS) and in [1]&[16, p19] (assuming furthermore a trusted key setup). Our
 266 lower bounds *do* also hold in this model, but provided an *unlimited* Environment. Indeed
 267 in the proofs of Theorems 4, 6 and 9, it could then sample a distribution of secret keys
 268 conditioned to the public keys on the board.

269 One could also consider an intermediate model, without any public key board but with a
 270 trusted authority that delivers certificates to players, that match their public keys with their
 271 identities. A player P sending for the first time a signed message to player Q , with his public

272 key, may then exhibit this certificate to Q . Our lower bounds would also hold in this model
 273 if the authority accepted to certify *multiple* keys from the same player. Indeed in our proofs,
 274 a freshly corrupted player could then ask for a *new* certificate related to the public/private
 275 key pair of the execution simulated by the Environment, use this new public key towards the
 276 isolated player p , while continuing to use the old one with the remaining honest players.

277 1.2 Roadmap

278 In §2.1 and §2.2 we make more explicit the model of behavior of players and provide additional
 279 notations for the proofs. We state precisely and prove Theorems 4, 6 and 9 in §2.3, §2.4 and
 280 §2.5. In the Appendix we revisit known synchronous lower bounds. Particularly, in §A.2.3
 281 we take the opportunity to expose a detailed proof of the [12, Theorem 2], and provide what
 282 we think are fixes for two existing arguments in the litterature.

283 2 Precise statements and proofs

284 2.1 Disclaimer on a simplification made in the statements of Thms 4 285 & 6 and in the proof of Thm 9

286 For simplicity in the statements of Theorems 4 and 6, we consider protocols that always
 287 satisfy Termination and (Strong or Weak) unanimity, that is, with 100% probability. So
 288 that we can concentrate on exhibiting Consistency failures in the proof. Of course when
 289 considering protocols for which Termination or Unanimity can fail up to some probability,
 290 then our arguments as such would not work. But then, these failures sum up in the total
 291 probability of failure anyway. Adapting this intuition rigorously would follow exactly the
 292 same pattern as in the proofs of [1, Theorems 1 & 3]. As for Theorem 9, since the proof
 293 follows exactly the same pattern as the [1, Thm 3], then we allow ourselves to state it with
 294 the same probability of overall failure ($\leq 1/6$), although, for simplicity in the proof, we do as
 295 if Termination and Strong unanimity always held, leaving the reader to check the detailed
 296 proof of [1, Thm 3] in case he needed to be convinced about the $1/6$.

297 2.2 General ingredients and precisions

298 We denote $|S|$ the cardinality of a set S .

299 **the State** of a honest player Q in a given execution e up to time τ , and denoted $e_\tau(Q)$, is
 300 his history up to τ , and consists of:

- 301 ■ His initial binary input value given by the Environment
- 302 ■ The messages he received so far, along with the time at which he received them, and the
 303 identity of their sender. We insist that players react identically to all **signatures** received
 304 in messages (or copied or obtained from \mathcal{F}_{ds}). The only extractable information from
 305 them are the outputs of **check** .
- 306 ■ The random coins he tossed so far (which are purely local computations), along with the
 307 time at which they were tossed
- 308 ■ Although these are local computations in the concrete sense, we specify that the state
 309 also contains the results of his **check** requests.

310 We insist that players do *not* receive any other incoming information than their input,
 311 messages from other players, and results of **check** . E.g. they do *not* receive any notification

:8 Lower bounds, limits of standalone digital signatures

312 of public keys from some board as in §1.1.0.1 (let alone of a private key from a trusted
313 dealer).

314 For instance, a player P receiving a message m_1 from P_1 containing a signature of P_2
315 on some message m_2 enclosed, will add everything to his current state: the time of delivery
316 and issuer P_1 , and the *whole* content of m_1 , including the signature of P_2 . . . irrespective to
317 whether or not this content itself would be specifically of the form $m_1 = m \parallel s_{P_1}(m)$ (in the
318 concrete terminology: where m would be the “body” of m_1 and $s_{P_1}(m)$ a valid signature of
319 P_1 “on m_1 ”, see Definition 1).

320 **Actions** taken by a honest player P (in the [14, §2.2 (3)]: at the beginning of a round) are
321 completely *determined* by his current state, and consist of:

- 322 ■ local computations, including: tossing coins, performing sign requests, copying signatures,
323 creating messages as concatenations of bitstrings and signatures etc.
- 324 ■ sending messages
- 325 ■ output a value
- 326 ■ check requests (although these are local computations in the concrete sense)

327 **Random variables** Let us transpose the vocabulary of the deterministic [12]. For a given
328 protocol, we will often identify a given environment \mathcal{E} with the probabilistic set of executions
329 under \mathcal{E} . Then, for instance, the “number of signatures or messages sent” (Definitions 1
330 and 2) under \mathcal{E} denotes actually the corresponding random variable with respect to this
331 probabilistic set of executions. We call execution e *restricted* to a honest player, or group of
332 honest players Q , and denote $e(Q) := e_\infty(Q)$, the history of internal states of players in Q
333 during the whole execution (or up to some time τ when precised). For a *fixed* environment
334 \mathcal{E} and for a player, or group of players, Q , we note $\mathcal{E}(Q)$ the probabilistic set of executions
335 restricted to Q under environment \mathcal{E} . In our arguments, we will meet the situation where
336 distributions $\mathcal{E}(Q)$ and $\mathcal{E}'(Q)$ coincide on a probabilistic set of measure *proba*. In this event,
337 then the probabilistic distribution of outputs of Q (if any) is the same. On the other hand,
338 $(1 - \text{proba})$ accounts for the events where the two sets of executions may *not* be equally
339 distributed. Typically, because too much players interact with Q , and so \mathcal{E}' is not able to
340 “fake” the distribution \mathcal{E} towards Q .

341 **Players interacting with Q** For a given player Q , in an execution e , let us consider the sets

- 342 ■ of players which receive a message from Q or a message containing a **signature** issued by
343 Q (in the same strict sense than in Definition 1, we recall this in the next item). We call
344 them the *players which are reached by Q* , and denote them $\text{RB}_e(Q)$;
- 345 ■ of players such that Q either receives a message from them, or receives some message
346 containing a **signature** issued by them. We call them the *players from which Q hears*
347 *of*, and denote them $\text{HO}_e(Q)$ (compare with [9]). [Notice that if Q receives the sole
348 message $m = s_{R_2}(s_{R_3}(m_3))$, then by definition R_3 is *not* in the set, because the **signature**
349 $s_{R_3}(m_3)$ is *not* in the concatenation of objects forming m (notice that, concretely, someone
350 receiving solely m would *not* be able to extract $s_{R_3}(m_3)$ in any manner, nor even guess
351 that this is the object signed by s_{R_2} , since she could not perform any successful **check**
352 request without s_{R_2} at hand)]
- 353 ■ the union of the two previous sets is denoted $A_e(Q)$ (as in [12, Theorem 1]). We call
354 them the *players interacting with Q* in execution e , and denote them $A_e(Q)$.

355 Under a probabilistic distribution of executions, in particular an Environment \mathcal{E} , then the
 356 previous become random variables: $\text{RB}_{\mathcal{E}}(Q)$, $\text{HO}_{\mathcal{E}}(Q)$ and $|A_{\mathcal{E}}(Q)|$. By definition we have

$$357 \quad |A_{\mathcal{E}}(Q)| \leq |\text{HO}_{\mathcal{E}}(Q)| + |\text{RB}_{\mathcal{E}}(Q)| \quad (1)$$

358 **Signatures received (and sent)** In addition to the $\text{Sent}_e(Q)$ of Definition 1, let us define
 359 for any player Q and execution e , the total number of signatures received by Q while he is
 360 honest, that we denote $\text{Received}_e(Q)$. That is, we count every signature contained in messages
 361 received by Q , plus 1 for every message received which is not of the form $m' = m \parallel s_Q(m)$
 362 (which means an “unsigned” message in the concrete sense). On the one hand we have:

$$363 \quad |\text{HO}_e(Q)| \leq \text{Received}_e(Q) \quad (2)$$

364 but on the other hand there is no comparison between $\text{Sent}_e(Q)$ and $|\text{RB}_e(Q)|$.

365 2.3 Theorem 4

366 Of course, under partial synchrony, the results are meaningful only for $t < \frac{n}{3}$, otherwise no
 367 consensus protocol exists at all. The following theorem holds under all partially synchronous
 368 models, especially those where *every* message sent is ultimately delivered. We consider the
 369 two “optimistic” Environments \mathcal{E}_0 and \mathcal{E}_1 which: leave *all* players honest, set GST from the
 370 beginning, set the same public Δ network delay from the beginning (that is: $\text{GST}=0$), and
 371 give input 0, resp. 1, to all the players. Under these environments, then the Weak unanimity
 372 condition imposes that all players ultimately output 0, resp. 1.

373 **A technical assumption** We assume for Theorem 4 that honest players Q do not send
 374 signatures $s_P(m_P)$ in their messages to player R , such that the message m_P contains itself
 375 a signature $s_T(m_T)$ that Q did not send so far to R . Informally, Q does not send some
 376 signature to R on which R is unable to perform a check that returns true (even with
 377 unlimited computational power), from what Q sent to far to him.

378 **► Theorem 4.** *Consider the partially synchronous model [14, §2.3 (3)]. Assume a standard*
 379 *adaptive Byzantine Environment and standalone digital signatures. Then, for every $\eta > 0$*
 380 *(small enough), consider a randomized protocol that solves consensus with Weak unanimity*
 381 *after GST and such that, under both the two optimistic environments \mathcal{E}_0 and \mathcal{E}_1 , we have*
 382 *that: Unanimity and Termination hold with 100% probability, and the expected number of*
 383 *signatures sent by forever honest players after GST is $\leq \frac{1}{4}\eta nt$. Then the probability of a*
 384 *Consistency violation is at least $1 - \eta$.*

385 *The same result holds under the model [14, (2)]: in the statement, replace “after GST”*
 386 *by “after honest players have guessed the actual network delivery delay”.*

387 We provide the detailed proof under the model [14, §2.3 (3)] (GST). At the end, we will
 388 briefly explain how to adapt the argument to the other model [14, (2)].

389 Notice that, under both optimistic environments \mathcal{E}_0 and \mathcal{E}_1 , since all players are honest
 390 and all messages are delivered, we have that for every execution e , the number of signatures
 391 sent by (honest) players (Definition 1) is then the same as the number of signatures received
 392 by (honest) players:

$$393 \quad \sum_{Q \in \mathcal{P}} \text{Sent}_e(Q) = \sum_{Q \in \mathcal{P}} \text{Received}_e(Q) \quad (3)$$

394 Also, since all players are honest, we have that:

$$395 \quad \sum_{Q \in \mathcal{P}} |\text{RB}_e(Q)| \leq \sum_{P \in \mathcal{P}} \text{Sent}_e(P) \quad (4)$$

396 Indeed, for every player R reached by a given Q , we have at least one of the two events:
 397 either R receives a signature from Q sent by some honest player P , or R receives a message
 398 from the honest $P = Q$. Both cases count as 1 in the definition (Def 1) of $\text{Sent}_e(P)$. When
 399 Q and/or R vary, then those events are *disjunct*. Thus, taking the sum over all R reached by
 400 Q , then over all Q , gives the inequality.

401 For each $b \in \{0, 1\}$, let us sum the *expectations* E_b under \mathcal{E}_b of the two equal quantities
 402 (signatures sent and signatures received) of (3). Then sum over $b \in \{0, 1\}$. By assumption,
 403 the total is thus lower than $2 \times 2 \times \eta n \frac{t}{4} = \eta nt$. Applying (2) and (2), we obtain:

$$404 \quad \eta nt \geq \sum_{Q \in \mathcal{P}} E_0(|\text{RB}_0(Q)|) + E_0(|\text{HO}_0(Q)|) + E_1(|\text{RB}_1(Q)|) + E_1(|\text{HO}_1(Q)|) \quad (5)$$

405 Dividing by n , we deduce that there exists a player p , such that in expectation:

$$406 \quad \eta t \geq E_0(|\text{RB}_0(p)|) + E_0(|\text{HO}_0(p)|) + E_1(|\text{RB}_1(p)|) + E_1(|\text{HO}_1(p)|) . \quad (6)$$

407 ► **Lemma 5.** *Sample two independent executions: one under \mathcal{E}_0 , the other one under \mathcal{E}_1 .
 408 Note $A(p)$ the set of players, cumulated over the two executions, such that either p is sent or
 409 forwarded a signature issued by some player in $A(p)$, or one of the players in $A(p)$ is sent or
 410 forwarded a signature issued by p . Then, with probability $1 - \eta$, we have that the cardinality
 411 $|A(p)| \leq t$.*

412 **Proof.** First, notice that $A(p)$ is by definition lower than the RHS of equation (6). Then,
 413 the rest is a straightforward application of the Markov bound: divide the LHS of Equation
 414 (6) by η . ◀

415 **Idea of the proof** We consider the following Environment \mathcal{E} . It takes advantage that, under
 416 the optimistic environment \mathcal{E}_0 , since we have GST=0 and unanimity of inputs, then all
 417 players must ultimately output 0 (respectively: 1 under \mathcal{E}_1). Its strategy will consist in having
 418 two groups of honest players: $\{p\}$ and $\mathcal{P} \setminus \{p \cup A(p)\}$ not interact with each other, and believe
 419 that they are facing respectively \mathcal{E}_0 and \mathcal{E}_1 , and thus to ultimately output. Recall that in the
 420 classical deterministic proof of [12, Theorem 1] (which we adapt for convenience to consensus
 421 in A.2.1), the Environment knows these sets prior to the execution, so can corrupt $A(p)$ from
 422 the beginning, and succeed in his strategy with certainty. In our probabilistic setting instead,
 423 the Environment \mathcal{E} *discovers* on the fly players which are sending messages to p , and those
 424 to which p sends messages. To achieve his strategy, \mathcal{E} will: adaptively delay forever those
 425 problematic messages (so will never actually set GST), corrupt on the fly sufficiently many
 426 players (the set $A(p)$), and make them continue interacting with both $\{p\}$ and $\mathcal{P} \setminus \{p \cup A(p)\}$
 427 as if they were still facing \mathcal{E}_0 and \mathcal{E}_1 (in particular: make them believe that GST holds).

428 **Proof** We consider the following Environment \mathcal{E} . It initially performs (a polynomial number
 429 of) simulations of executions (up to a polynomial time or number of rounds) under \mathcal{E}_0 and
 430 \mathcal{E}_1 , to identify with overwhelming probability a player p such that (6) holds up to the time
 431 τ when all players output in both executions (which is a fortiori guaranteed by (6)). It
 432 gives input 1 to all players, which are all honest at the beginning, except to p , to which he

433 assigns input 0. It sets a public Δ eventual delivery delay, the same as in the two optimistic
 434 environments \mathcal{E}_0 and \mathcal{E}_1 . Then, \mathcal{E} starts simulating the players in $\mathcal{P} \setminus \{p\}$ in this head, giving
 435 them inputs 0. Let us call this the *simulated* execution.

436 Then the real execution starts and runs at the same pace as the simulated execution,
 437 which is a little ahead in time, so that the Environment has the time to perform the following
 438 actions. Let $A(p)$ be the set of corrupted players so far (initially empty). Corrupted players
 439 continue their initial thread of playing honestly the real execution with the remaining so far
 440 honest players except p : $\mathcal{P} \setminus \{p \cup A(p)\}$, as if they still had input 1, and messages between
 441 them are still delivered within Δ . On the other hand, as soon as they are corrupted, players
 442 in $A(p)$ open a second thread that has the same behavior as in the *simulated* execution:
 443 messages sent to p by this thread are sent in the *real* execution (and de facto messages
 444 received from p), while messages sent to $\mathcal{P} \setminus \{p \cup A(p)\}$ are sent in the *simulated* execution.
 445 The corruptions enabling this strategy are as follows:

- 446 ■ As soon as any player q in the *real* execution sends a message to p , then corrupt this
 447 player q and arbitrarily delay the message that he sent to p
- 448 ■ As soon as p sends a message to a player q' in the *real* execution, then corrupt q' .
- 449 ■ As soon as new players are Reached By p in the *simulated* execution, then corrupt their
 450 counterparts in the *real* execution;
- 451 ■ At the beginning of each round of the real execution, \mathcal{E} checks to see which players R
 452 will send a message m to p in this round of the simulated execution under \mathcal{E}_0 . For every
 453 such R and m :
 - 454 ■ let R_2, \dots, R_k the (possibly empty) set of players which is added to the ones that p
 455 heard of so far in the simulated execution under \mathcal{E}_0 , due to the reception of m .
 - 456 ■ \mathcal{E} corrupts players R, R_2, \dots, R_k in the real execution before they could possibly sent
 457 any message in this round of the *real* execution;
 - 458 ■ by the technical assumption, the corrupted players (which by definition include the
 459 ones that p heard of so far in the simulated execution up to reception of m) are able
 460 to create all the signatures appearing on m in the simulated execution (recall that all
 461 signatures are indistinguishable from p , up to the output given by his check requests
 462 on them).
 - 463 ■ use these signatures to form the message m (which we recall is by definition a concat-
 464 enation of these signatures and bitstrings). Have R send m to p in the *real* execution.

465 This guarantees preservation of the following *Invariant*: consider the set of real executions
 466 in which the Environment can afford his strategy. Then, restricted to the honest players
 467 $\mathcal{P} \setminus \{p \cup A(p)\}$, who did not interact with p , they follow that same distribution than randomly
 468 sampled executions under Environment \mathcal{E}_1 . Symmetrically, restricted to p , they follow the
 469 same distribution than randomly sampled executions under Environment \mathcal{E}_0 . Thus, in
 470 these executions, p must then ultimately output 0. Whereas the other honest players must
 471 ultimately output 1, by indistinguishability with the distribution under \mathcal{E}_1 . So this is a
 472 consistency failure in the real execution.

473 The Environment can afford its strategy up to probability at least $1 - \eta$ because then, by
 474 Lemma 5, the corruption limit has not been achieved: $|A(p)| \leq t$.

475 **Adapting the proof for the other model [14, (2)] (Δ unknown)** We consider the same
 476 baseline optimistic Environments \mathcal{E}_0 and \mathcal{E}_1 , but in addition which both set a (small) delivery
 477 delay Δ' for messages. Then, the Environment \mathcal{E} sets $\Delta \gg \Delta'$ high enough, and makes
 478 players guess the wrong Δ' , such that, with the previous strategy, all honest players are
 479 very likely to output a value before Δ has elapsed. Concretely, in the previous strategy, the

480 Environment delivers all messages within delay Δ' , except those that he wants to arbitrarily
 481 delay. Although these messages are still guaranteed to be delivered after Δ , the strategy will
 482 still lead to the same Consistency failure because all honest players will have output by then.

483 2.4 Theorem 6

484 2.4.1 Statement and idea of proof

485 For readability we it with a *fixed* number of rounds R , as in Rabin [20, §5], and also with a
 486 η -*worst case* message complexity, as in [1, Theorem 3]. This is no more than a matter of
 487 taste: [1, Theorem 1] was stated instead with *expected* complexity (as the previous Theorem
 488 4), and, [1, Theorem 2] was stated in terms of expected number of rounds .

489 More importantly, the proof will involve two baseline Environments which, now, corrupt
 490 some players. This is the reason why we cannot anymore rule out consensus with weak
 491 unanimity. And also why we *enlarge* the message complexity measurement: we now consider
 492 messages sent by *all* so-far honest players, *including* those who may be corrupted later in
 493 the execution.

► **Theorem 6.** *Consider the partially synchronous model [14, §2.3 (3)]. Assume a standard adaptive Byzantine Environment and standalone digital signatures. Let $\eta > 0$ be small enough, $\Lambda > 0$, and consider a randomized protocol for consensus with Strong unanimity such that under every possible fixed Environment, we have that: Unanimity and Termination hold with 100% probability, and so-far honest players send a total $\leq \Lambda n$ messages after GST, except with probability $\leq \eta$. Assume furthermore that players output within R rounds after GST in every execution, where R is such that*

$$R \leq \frac{\log(t) - \log 4}{\log \Lambda + \log 4 + \log(n/t)} .$$

Then the probability of a Consistency violation is at least

$$\geq \left(1 - \frac{t}{2n}\right)(1 - 2\eta) .$$

494 The same result holds under the model [14, (2)]: in the statement, replace “after GST”
 495 by “after honest players have guessed the actual network delivery delay”.

496 Notice that $n/t = 3$ in the maximal corruption tolerance. The basic idea is that the
 497 Environment can silent forever the players which receive more than M messages in some
 498 round (M a parameter to be defined). Then, for every player Q which is still honest after
 499 R rounds, the number of *distinct* players of which Q will have received/been forwarded
 500 signatures, is bounded by $2 \times M^R$. Indeed: a message sent in the first round contains at
 501 most the signature of 1 player, thus, by assumption, a message sent in the second round
 502 contains at most the signature of M distinct players, then we conclude by recursion that Q
 503 hears at most from the signatures from M^R distinct players. And, seen the other way, the
 504 graph of communications from Q forms a tree with at most M branches at each node, so
 505 reaches at most M^R players. We can then adapt the proof of Theorem 4.

506 We adapt the proof of Theorem 4 by modifying the baseline Environments: \mathcal{E}_0 and \mathcal{E}_1 ,
 507 which now also *corrupt* players. As before, \mathcal{E}_0 and \mathcal{E}_1 set GST from the beginning and give
 508 all players input 0, resp. 1. But in addition, they adaptively silent in every round the players
 509 that collect many signatures. Precisely, in every round, they corrupt the “collector” players
 510 who received more than $M := \frac{4\Lambda n}{t}$ messages in the previous round, and instruct them not to

511 send any message forever. When either \mathcal{E}_0 or \mathcal{E}_1 achieves the corruption quota of $t/4$, it raises
 512 a “give up” flag in his head and stops corrupting players. The following Lemma formalizes
 513 to what extent they can afford their strategy without raising the flag. For $b \in \{0, 1\}$, for
 514 simplicity, note C_b the probabilistic set of players that, *overall* in an execution under \mathcal{E}_b ,
 515 receive more than $\frac{4\Lambda n}{t}$ messages.

516 ► **Lemma 7.** *Consider the two environments \mathcal{E}_0 and \mathcal{E}_1 . Sample (independently) one*
 517 *execution of the protocol under each of them. Then, with probability greater than $1 - 2\eta$, we*
 518 *have that both cardinalities $|C_0|$ and $|C_1|$ are simultaneously bounded.*

$$519 \quad |C_b| \leq \frac{t}{4} \text{ for } b \in \{0, 1\}. \quad (7)$$

520 **Proof.** Let b be equal to 0 or 1. By definition of \mathcal{E}_b , only honest players send messages.
 521 And thus the total number of messages *received* by players is equal to the total number of
 522 messages sent by so-far honest players. Which is, by assumption, $\leq \Lambda n$ with probability
 523 $1 - \eta$.

524 Thus, under this (likely) event, there are at most $|C_b| \leq \Lambda n \frac{t}{\Lambda 4n} = \frac{t}{4}$ distinct players
 525 receiving more than $\frac{4\Lambda n}{t}$ messages overall in the execution.

526 Finally, we have probability $1 - 2\eta$ that this holds *simultaneously* in both the sampled
 527 executions. ◀

528 Next, the following Lemma 8 parallels Equation (5), and determines when our final
 529 Environment will be able to “frame” one player as before. It is a consequence of the intuition
 530 given after the statement of Theorem 6.

531 **2.4.1.1 (so far) HeardOf(Q) / ReachedBy(Q) in the most possible inclusive sense:**
 532 **players required to make the signatures (so far) received by Q / whose**
 533 **signatures (so far) received require the action of Q to be constructed**

534 We do not make anymore the “technical assumption” made for Theorem 4. Now we enlarge
 535 the definition of players HeardOf(Q) to include all those necessary to construct the messages
 536 (so far) received by Q . That is, if Q receives the sole message $s_{R_2}(s_{R_3}(m_3))$, then we now say
 537 that Q heard of both R_2 and R_3 since this message cannot be created without actions from
 538 both of them. The set ReachedBy(Q) is equally enlarged in the same inclusive sense: players
 539 who (so far) received signatures such that an action of Q is necessary to construct them.

540 ► **Lemma 8.** *We still consider two independent executions of the protocol under \mathcal{E}_0 and \mathcal{E}_1 .*
 541 *Note $A(Q) := A_0(Q) \cup A_1(Q)$. Then, under the event of (7) (of probability $1 - 2\eta$), we have*
 542 *that for every forever honest Q :*

$$543 \quad \frac{t}{2} \cong 4 \left(\frac{8\Lambda n}{t} \right)^R \geq |\text{RB}_0(Q)| + |\text{HO}_0(Q)| + |\text{RB}_1(Q)| + |\text{HO}_1(Q)| \geq |A(Q)| \quad (8)$$

544 **Proof.** The RHS upper bound is by Equation (1).

545 *First*, let us recall the intuitive fact on the signatures sent and received. Consider
 546 $b \in \{0, 1\}$, and an execution under \mathcal{E}_b . By assumption we are in the (likely) event where (7)
 547 holds. In particular, \mathcal{E}_b is able to carry his strategy without raising the “give up” flag. Hence,
 548 we have that throughout the execution, the players who send messages in a round are those
 549 who so-far received $\leq \frac{4\Lambda n}{t}$ messages in *every* previous round. To initialize the recursion: a

550 player who sends message in the second round cannot have heard of so far from more than
 551 $\frac{4\Lambda n}{t}$ distinct players, and thus cannot send or forward signatures that require more than
 552 $\frac{4\Lambda n}{t}$ distinct players to construct them. We continue by recursion until the beginning of
 553 round R , where each player did not hear of more than $\left(\frac{4\Lambda n}{t}\right)^{R-1}$ distinct players. And
 554 thus in turn, all the signatures that every single honest player will send or forward in R , are
 555 guaranteed to be constructible from a set of players (in our inclusive sense of Heard Of) of
 556 size *lower* than this number. Thus overall in such an execution, any player *who remained*
 557 *honest at the end of round R* received signatures that are constructible from a set of players
 558 of size lower than $\text{HO}_b(Q) \leq \left(\frac{4\Lambda n}{t}\right)^R$. Symmetrically, every player Q cannot reach more
 559 than $\text{RB}_b(Q) \leq \left(\frac{4\Lambda n}{t}\right)^R$ distinct players in the inclusive sense above.

560 *Summing* both previous upper bounds, then over both executions $b \in \{0, 1\}$, we deduce
 561 the upper bound in the middle of Equation (8).

562 *Taking the log* of $4\left(\frac{4\Lambda n}{t}\right)^R$ yields $\log(4) + \log(t) - \log(4)$, which is (nearly) the log of
 563 the LHS of (8). ◀

564 **Proof of Thm 6** We can now describe the Environment \mathcal{E} that will provoke consistency
 565 failures. The pattern is the same as in 4, we just change the baseline Environments \mathcal{E}_0 and \mathcal{E}_1 .
 566 \mathcal{E} selects a player p at random. It gives input 1 to all players, except to p to which it gives
 567 input 0. It initializes three sets of corrupted players, initially empty: the real ones $A(p)$ and
 568 C_1 , and the virtual one C_0 . It starts a simulation in his head where he runs environment \mathcal{E}_0 .
 569 Particularly, it corrupts in his head every player that receives more than $\frac{4\Lambda n}{t}$ messages at
 570 the end of a round. Such virtual player is added to the virtual set C_0 and virtually silenced
 571 forever. He does likewise in the real execution, adding silenced players to the real set C_1 . If
 572 by chance p gets corrupted in the simulated execution (added to C_0) or in the real execution
 573 (added to C_1), then the Environment raises an “abort” flag and stops his strategy.

574 The rest carries over unchanged: as soon as any player q in the *real* execution sends a
 575 signature to p , then \mathcal{E} corrupts this player q and arbitrarily delay the original message he
 576 sent to p . As soon as p sends a message to a real player q' , then corrupt q' . If in some round
 577 of the *simulated* execution there is a player Z which sends a message to p , which possibly
 578 requires participation from other players: Z_2, \dots, Z_k to be created (in our inclusive sense
 579 of Heard Of), then \mathcal{E} corrupts the corresponding players Z, Z_2, \dots, Z_k in the real execution
 580 at the beginning of this round, before they had time to send any message. \mathcal{E} makes them
 581 issue the signatures appearing in the simulated execution (and forge any possibly additional
 582 unsigned message forwarded by Z to p , when the case), creates from these signatures the
 583 message as in the simulated execution, then makes Z actually send it to p .

584 As soon as they are corrupted, players (in $A(p)$) are made to play the protocol with p
 585 exactly as in the simulated execution under \mathcal{E}_0 , but continue playing the real execution with
 586 the remaining honest players, as if still under \mathcal{E}_1 .

587 This guarantees preservation of the following *Invariant*: consider the set of real executions
 588 in which the Environment can afford his strategy. Then, restricted to the honest players
 589 $\mathcal{P} \setminus \{p \cup A(p) \cup C_1\}$, who did not interact with p , they follow the same distribution than
 590 randomly sampled executions under Environment \mathcal{E}_1 . Symmetrically, restricted to p , they
 591 follow the same distribution than randomly sampled executions under Environment \mathcal{E}_0 .

592 Thus, *if* p remains uncorrupted at the end of the R -th round, then it must output 0 as

593 under \mathcal{E}_0 (Weak unanimity + after GST), whereas the other honest players must ultimately
 594 output 1 at the end of the R -th round as under \mathcal{E}_1 (Weak unanimity + after GST). So this
 595 is a consistency failure in the real execution.

596 Environment \mathcal{E} can afford its strategy if those three sufficient conditions hold:

- 597 ■ p does not get corrupted in C_0 in the simulated execution, nor in C_1 in the real execution.
 598 This happens with probability at least $1 - t/(2n)$ because p was chosen at random, while
 599 the set to avoid is of size $\leq |C_0| + |C_1| \leq t/2$;
- 600 ■ \mathcal{E} can silent players in both the real and simulated executions, without reaching the
 601 corruption quota $\frac{t}{4}$ in both. By Lemma 7, this happens with probability at least $1 - 2\eta$;
- 602 ■ the size of the set $A(p)$ that \mathcal{E} also corrupts in the real execution remains smaller than
 603 $\frac{t}{2}$. But, under the two previous conditions, then notice that p is a forever honest player
 604 in *both* executions. Namely the real one, which has the same distribution than under \mathcal{E}_1
 605 for the players which did not interacted with p so far, and the simulated one, which is
 606 exactly under \mathcal{E}_0 . So we conclude by Lemma 8 that this third condition also holds.

607 In conclusion, those three sufficient conditions hold with probability greater than $(1 -$
 608 $\frac{t}{2n})(1 - 2\eta)$.

609 2.5 Theorem 9

610 ► **Theorem 9.** *Consider any synchronous randomized protocol for consensus, assuming*
 611 *standalone digital signatures, such that Termination + Consistency + Strong Unanimity*
 612 *hold with probability $p > 5/6$ under any standard adaptive Byzantine Environments with up*
 613 *to t corruptions; and such that honest players only multicast messages. Then, there are*
 614 *executions in which so-far honest players multicast more than t messages.*

615 We just describe the differences with the baseline argument that we set in A.2.2, and still
 616 follow the simplification discussed in 2.1. There is no more designated sender P_2 . Instead,
 617 Environment \mathcal{E}_0 gives input 0 to all players except to the corrupted P_1 , and input 1 to all
 618 players in the simulated execution. Likewise for \mathcal{E}'_0 , which gives input 0 to all the players
 619 except to player P_1 , to which it gives the input bit b , which we do not specify on purpose.
 620 Whereas in the simulated execution, \mathcal{E}'_0 gives input 1 to all players. The same modifications
 621 holds for \mathcal{E}_1 and \mathcal{E}'_1 (inputs 1 in the real executions, 0 in the simulated ones, and the same
 622 bit b to P_1 under \mathcal{E}'_1).

623 The second difference is that under \mathcal{E}_0 and \mathcal{E}_1 , now, whenever a player q multicasts a
 624 message in the simulated execution, then it is immediatly *corrupted* in the real execution.
 625 This enables the Environment to use the signature of q and carry over the same strategy:
 626 make P_1 issue valid messages in the real execution, that he would have sent in the simulated
 627 execution.

628 Denote by Q_0 the probabilistic set of forever honest players under \mathcal{E}_0 *except* P_1 . They
 629 must ultimately output 0, by Strong unanimity. Denote by Q'_0 the probabilistic set of forever
 630 honest players under \mathcal{E}'_0 *except* P_1 . The set of executions restricted to them has the same
 631 distribution than for Q_0 under \mathcal{E}_0 .¹ Thus players in Q'_0 also ultimately output 0. For the

¹ There are two ways to cope with this argument of a *probabilistic* group of players witnessing the same distributions of executions, which is also implicit e.g. in [1, Thm 1 & Thm 3]. The quick way around is just to fix once for all a q at random, and then conditioning on the events where he belongs to *both* Q_0 and Q'_0 (considering two parallel samplings of executions under \mathcal{E}_0 and \mathcal{E}'_0). The formal way consists in seeing the set of honest players as one single big probabilistic machine that interacts with the

632 same reason (indistinguishability with \mathcal{E}_1), we have that forever honest players Q'_1 under \mathcal{E}'_1
 633 must ultimately output 1. Finally, to ensure *Consistency* with Q'_0 in \mathcal{E}'_0 , then the honest
 634 player P_1 must output 0. For the same Consistency reason, he must output 1 under \mathcal{E}'_1 .

635 But \mathcal{E}'_0 and \mathcal{E}'_1 restricted to P_1 have identical distributions. So P_1 should output the same
 636 value in both, which is impossible. Notice that the value of his fixed input bit b played no
 637 role in the argument.

638 ——— References ———

- 639 1 Ittai Abraham, T-H. Hubert Chan, Danny Dolev, Kartik Nayak, Rafael Pass, Ling Ren, and
 640 Elaine Shi. Communication complexity of byzantine agreement, revisited. In *Proceedings of*
 641 *the 2019 ACM PODC*, 2019.
- 642 2 Ziv Bar-Joseph and Michael Ben-Or. A tight lower bound for randomized synchronous
 643 consensus. In *Proceedings of the Seventeenth Annual ACM Symposium on Principles of*
 644 *Distributed Computing*, 1998.
- 645 3 Michael Ben-Or. Another advantage of free choice (extended abstract): Completely asynchron-
 646 ous agreement protocols. In *Proceedings of the Second ACM PODC*, 1983.
- 647 4 Alexandra Boldyreva. Threshold signatures, multisignatures and blind signatures based on the
 648 gap-diffie-hellman-group signature scheme. In *Public Key Cryptography — PKC 2003*, 2002.
- 649 5 Elette Boyle, Ran Cohen, and Aarushi Goel. Succinctly reconstructed distributed signatures
 650 and balanced byzantine agreement, 2020. [arXiv:2002.02516](https://arxiv.org/abs/2002.02516).
- 651 6 Christian Cachin, Klaus Kursawe, Frank Petzold, and Victor Shoup. Secure and efficient
 652 asynchronous broadcast protocols. In Joe Kilian, editor, *Advances in Cryptology — CRYPTO*
 653 *2001*. Springer Berlin Heidelberg, 2001.
- 654 7 Christian Cachin, Klaus Kursawe, and Victor Shoup. Random oracles in constantinople:
 655 Practical asynchronous byzantine agreement using cryptography. *Journal of Cryptology*, 2005.
- 656 8 Miguel Castro and Barbara Liskov. Practical byzantine fault tolerance. In *Proceedings of the*
 657 *Third Symposium on Operating Systems Design and Implementation*, OSDI '99, Berkeley, CA,
 658 USA, 1999. USENIX Association.
- 659 9 B. Charron-Bost and A. Schiper. Improving fast paxos: being optimistic with no overhead. In
 660 *2006 12th Pacific Rim International Symposium on Dependable Computing (PRDC'06)*, Dec
 661 2006.
- 662 10 Jing Chen and Silvio Micali. Algorand, 2016. [arXiv:1607.01341](https://arxiv.org/abs/1607.01341).
- 663 11 Bernardo David, Peter Gazi, Aggelos Kiayias, and Alexander Russell. Ouroboros praos: An
 664 adaptively-secure, semi-synchronous proof-of-stake blockchain. In *Advances in Cryptology —*
 665 *EUROCRYPT 2018*. Springer International Publishing.
- 666 12 Danny Dolev and Rüdiger Reischuk. Bounds on information exchange for byzantine agreement.
 667 *J. ACM*, 32(1):191–204, January 1985.
- 668 13 Danny Dolev and H. Strong. Authenticated algorithms for byzantine agreement. *SIAM J.*
 669 *Comput.*, 1983.
- 670 14 Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. Consensus in the presence of partial
 671 synchrony. *J. ACM*, 35(2):288–323, April 1988.
- 672 15 Vassos Hadzilacos and Joseph Y. Halpern. Message-optimal protocols for byzantine agreement.
 673 *Mathematical systems theory*, 1993.
- 674 16 T.-H. Hubert Chan, Rafael Pass, and Elaine Shi. Consensus through herding. In *Advances in*
 675 *Cryptology — EUROCRYPT 2019*, 2019.
- 676 17 Valerie King and Jared Saia. Breaking the $o(n^2)$ bit barrier: scalable byzantine agreement
 677 with an adaptive adversary. In *PODC '10*, 2010.

Environment (inputs and outputs) and with the corrupted players (messages). When some player q gets corrupted, then we model this by shutting down the program of q in the machine, and creating one new message interface port with the corrupted q , for every so-far honest player remaining in the machine.

- 678 18 Eleftherios Kokoris-Kogias, Alexander Spiegelman, Dahlia Malkhi, and Ittai Abraham. Boot-
 679 strapping consensus without trusted setup: Fully asynchronous distributed key generation.
 680 *IACR Cryptology ePrint Archive*, 2019:1015, 2019.
- 681 19 Leslie Lamport, Robert Shostak, and Marshall Pease. The byzantine generals problem. *ACM*
 682 *Trans. Program. Lang. Syst.*, 1982.
- 683 20 Michael O. Rabin. Randomized byzantine generals. In *Proceedings of the 24th Annual SFCS*,
 684 1983.
- 685 21 Alin Tomescu, Robert Chen, Yiming Zheng, Ittai Abraham, Benny Pinkas, Golan Gueta, and
 686 Srinivas Devadas. Towards scalable threshold cryptosystems. In *to appear in IEEE Security &*
 687 *Privacy 2020*, 2020.
- 688 22 Vinod Vaikuntanathan. *Randomized algorithms for reliable broadcast*. PhD thesis, MIT, 2009.
- 689 23 Maofan Yin, Dahlia Malkhi, Michael K. Reiter, Guy Golan Gueta, and Ittai Abraham. Hotstuff:
 690 Bft consensus with linearity and responsiveness. In *Proceedings of the 2019 ACM Symposium*
 691 *on Principles of Distributed Computing*, PODC '19, New York, NY, USA, 2019. ACM.

692 **A** Revisiting state of the art lower bounds, their proofs and their 693 limitations

694 **A.1** Remainder: from broadcast to consensus

695 First, recall that consensus under synchrony is possible only with a *honest majority* $t < \frac{n}{2}$.
 696 Indeed, otherwise, consider four players: P_1 and P_2 honest with inputs 1 and 2. Plus P_3 and
 697 P_4 corrupted, who play honestly as if they had inputs 1 and 2. Then, to satisfy unanimity,
 698 P_1 must decide 1, because from his point of view it could be the case that only P_1 and P_3
 699 are noncorrupted. And likewise P_2 must decide 2 to satisfy unanimity with the possibly
 700 honest P_4 .

701 Second, let us recall the problem of broadcast, which is also known as “Byzantine
 702 agreement” [19]. In this problem there is a privileged player P_1 , the “sender”, and the goal
 703 is to have every player output the same value, such that (validity) if P_1 is honest, then this
 704 value is P_1 ’s input. Under synchrony *and* with a honest majority, then we have that the
 705 problem of broadcast is essentially *weaker* than consensus. Indeed, consensus then compiles
 706 itself into Broadcast with no communication overhead [22, §2.2]: have P_1 send his input
 707 s_1 to every player in the first round. Then in the second round players start a consensus
 708 protocol with input: the value that they received from P_1 , if any, or a default value otherwise
 709 —e.g. 0. This is why all the known lower bounds for synchronous broadcast, that we recall
 710 below, *automatically hold* for consensus.

711 Third, we do not consider broadcast under partial synchrony in this paper because it
 712 is *trivially impossible*: consider a malicious sender which remains silent forever but GST
 713 holds from the beginning. Then honest players must output some value in a finite time. The
 714 problem is that from the point of view of honest players, the situation is undistinguishable
 715 from a scenario where the sender P_1 is honest with an input s_1 that they ignore, while GST
 716 is never established: the Environment delays all outgoing messages from P_1 forever. Since
 717 other players output anyway some value, this value is likely to be different from s_1 , which
 718 violates validity.

719 **A.2** Revisiting communication lower bounds and their proofs for 720 *synchronous* broadcast

721 Let us revisit some synchronous bounds for authenticated broadcast and their proofs, from
 722 Dolev-Reischuk [12] (expanded from PODC’82) and Abraham et al [1] (PODC’19). Apart

723 from these papers, let us mention that Hadzilacos-Halpern [15, Figure 1] also consider lower
 724 bounds for the communication complexity of authenticated broadcast. They measure the
 725 complexity over *optimistic* executions. That is, those in which the Environment corrupts no
 726 player. Notice that in Thm 4 we will *also* consider a complexity on average on two optimistic
 727 Environments (the one where all players have input 0, the other where all players have input
 728 1). The difference being that [15, Figure 1] deals with the *message* complexity (and hold
 729 only in the deterministic setting). This measurement is, as we stressed above, *smaller* than
 730 the number of signatures complexity that we consider instead in Thm 4. Hence, they prove
 731 that the message complexity is linear under these optimistic environments.

732 Notice that in deterministic consensus, in which both players and the environment are
 733 deterministic, then an execution is *fully determined* by the Environment. This is why in this
 734 context we sometimes identify a given Environment \mathcal{E} with the corresponding execution.

735 **A.2.1 [12, Theorem 1]: *deterministic* broadcast has a quadratic $\frac{nt}{4}$
 736 worst case complexity in the number of *signatures* (sent or
 737 forwarded by honest players)**

738 Let us briefly recall the argument, that we trivially adapt to consensus here, because we will
 739 build on this version for Theorem 4. Let us consider the two executions following from the
 740 two “optimistic” Environments \mathcal{E}_0 and \mathcal{E}_1 which: leave all players honest, and give them all
 741 input 0 —resp. 1. Suppose by contradiction that, in *each* execution/Environment \mathcal{E}_0 and
 742 \mathcal{E}_1 , honest players send $\leq \frac{nt}{4}$ signatures. For any player Q , note $A(Q)$ the set of players,
 743 cumulated over both executions/Environments \mathcal{E}_0 and \mathcal{E}_1 , which: either receive a message
 744 with the signature of Q , or such that Q receives a message with the signature of one of them.
 745 This is by definition the set of players with which Q interacts. Then by assumption, there
 746 must exist at least one player p such that the cardinality $|A(p)| \leq t$. [Indeed notice that,
 747 since all players are honest, summing over every honest player p' the signatures *received* is
 748 then equal to summing over every honest player p' the total number of signatures sent]. Now,
 749 consider an Environment $\mathcal{E}_{01}^{A(p)}$ which: assigns input 0 to p and leaves him honest, corrupts
 750 all players in $A(p)$, and assigns 1 to the remaining honest players $U := \mathcal{P} - \{p\} - A(p)$.
 751 Then the Environment makes, on the one hand, players in $A(p)$ behave towards p as in the
 752 optimistic execution \mathcal{E}_0 where everyone has input 0. Thus p must ultimately output 1. And
 753 on the other hand, the Environment makes $A(p)$ behave towards U as if in the execution
 754 \mathcal{E}_1 where everyone has input 1. Thus players in U ultimately output 1. Notice that the
 755 Environment $\mathcal{E}_{01}^{A(p)}$ is able to carry out this strategy, since, by construction of the set $A(p)$,
 756 players in U never see the signature of p in \mathcal{E}_1 , nor does p see any signature from U in \mathcal{E}_0 .

757 **A.2.2 [1, Theorem 3]: assuming *no digital signatures*, and that players
 758 *only multicast* messages, then *randomized* broadcast has a
 759 quadratic $\Omega(nt)$ worst case complexity in the number of
 760 messages sent by so-far honest players**

761 Let us reformulate the argument for two reasons. First because we will build on it for our
 762 Theorem 4. Second, to emphasize that the existence of digital signatures (=standalone
 763 PKI) would break the argument. Which is something that we will repair in our synchronous
 764 Theorem 9 —at the price of considering *consensus* but not broadcast anymore. For simplicity
 765 we do not consider probabilities in the argument.

766 The idea of the proof is that honest player P_1 will be made “schizophrenic” and receive
 767 messages from two executions in parallel. He will not be able to distinguish what execution
 768 are playing the corrupted players and which one are playing the honest ones (depending
 769 whether he is under \mathcal{E}'_0 or \mathcal{E}'_1 , to be defined). We consider a protocol for broadcast with
 770 sender P_2 . Let us consider two baseline Environments \mathcal{E}_0 and \mathcal{E}_1 , that have the following
 771 strategy in each execution. \mathcal{E}_0 gives input 0 to P_2 , corrupts P_1 and simulates an execution in
 772 its head where P_2 would have input 1 instead. P_1 sends messages in the real execution as if
 773 he received *both* messages from the real execution, and from the simulated one. \mathcal{E}_1 follows a
 774 symmetric strategy: it gives input 0 to P_2 , corrupts P_1 and simulates an execution in its head
 775 where P_2 would have input 0 instead. As before, P_1 sends messages in the real execution as
 776 if he received *both* messages from the real execution, and from the simulated one. Since the
 777 sender P_2 is honest under \mathcal{E}_0 , by validity of broadcast, honest players ultimately output 0 in
 778 every execution. Likewise, since the sender P_2 is honest under \mathcal{E}_1 , by validity of broadcast,
 779 honest players ultimately output 1 in every execution.

780 Now, consider Environments \mathcal{E}'_0 and \mathcal{E}'_1 as follows. \mathcal{E}'_0 initially gives input 0 to P_2 then
 781 simulates an execution in its head where P_2 would have input 1 instead. Whenever a player
 782 different from P_1 is going to multicast a message in the simulated execution, it adaptively
 783 corrupts it. Let C'_0 be the set of corrupted players so far. Players in C'_0 are instructed to send,
 784 in the real execution, the messages that they would have sent in the simulated execution, but,
 785 instead of multicasting them, they send it only to P_1 . Apart from this, players in C'_0 also
 786 continue to play the real execution. \mathcal{E}'_1 has the symmetric strategy: initially gives input 1 to
 787 P_2 then simulates an execution in its head where P_2 would have input 0 instead. Whenever
 788 a player different from P_1 is going to multicast a message in the simulated execution, it
 789 adaptively corrupts it. Let C'_1 be the set of corrupted players so far. Players in C'_1 are
 790 instructed to send, in the real execution, the messages that they would have sent in the
 791 simulated execution, but, instead of multicasting them, they send it only to P_1 . Apart from
 792 this, players in C'_1 also continue to play the real execution.

793 On the one side, consider honest players apart from P_1 under Environment E'_0 . The set
 794 of executions restricted to them has an identical distribution to the set of executions under
 795 E_0 (they see a player P_1 which looks “schizophrenic”, since his messages are also triggered
 796 from what he receives in the simulated execution, resp., from players in C'_0). Thus they
 797 must ultimately output the same value: 0. Likewise, honest players apart from P_1 under
 798 Environment E'_1 , must ultimately output 1. In the middle, both Environments E'_0 and E'_1
 799 provide the honest P_1 with equally distributed executions. But, under E'_0 , P_1 must, on the
 800 one side, output 0 to be consistent with the other honest players. And on the one side, under
 801 E'_1 , he must output 1 to be consistent with the other honest players. Thus, although going
 802 through equally distributed executions, he would have to output two different values, which
 803 is impossible.

804 A.2.2.1 Two limits of the previous argument

805 First, *assuming now that messages are digitally signed*, then Environments \mathcal{E}_0 and \mathcal{E}_1 would
 806 be unable to carry their strategy. Indeed, although P_1 is the only corrupted player, he is
 807 supposed to issue messages compatible with what he received in the simulated executions.
 808 For instance, the protocol, as played in the *simulated* execution under \mathcal{E}_0 , is likely to require
 809 P_1 to multicast a message m' , that would likely contain the *signed input value* $b = 1$ that
 810 he would have received from the sender P_2 in the simulated execution. Then, recall that
 811 the strategy of \mathcal{E}_0 is to make P_1 multicast message m' in the real execution. But actually,
 812 since P_2 did *not* produce his signature on the value $b = 1$ in the real execution, and is *not*

813 corrupted, then \mathcal{E}_0 is unable to forge P_2 's signature on $b = 1$, and so would *fail* to make P_1
814 send the desired valid message m' in the real execution.

815 Second, one could consider a fix for this and have the Environment \mathcal{E}_0 corrupt the sender
816 P_2 , in order to force him to sign the messages that P_1 needs to forward in the real execution
817 (and likewise corrupt any other player whose signature is needed for P_1). The problem is
818 that, since now the sender is *corrupted*, we would then *not* be able to conclude anymore that
819 honest players need to output the specific value 0 under \mathcal{E}_0 . The same problem occurs for \mathcal{E}_1 .
820 This issue is actually stressed in [1, §6] (“so we need to show a violation of consistency, not
821 validity”).

822 A.2.3 [12, Theorem 2]: *deterministic* broadcast has a quadratic $\frac{t^2}{4}$ 823 worst case complexity in the number of *messages* (sent by 824 honest players)

825 A.2.3.1 Highlights

826 The situation of [12, Theorem 2] is tricky in many aspects. (i) On the one hand, the only
827 argument in the litterature by which we could be convinced is the one of... the *randomized*
828 lower bound of [1, Theorem 1]. We adapt it to the deterministic synchronous setting, and
829 highlight its specificities in brackets at two places of the revisited proof below. (ii) Then,
830 we briefly discuss other proofs for [12, Theorem 2] in the litterature by which we could not
831 be rigorously convinced. First, the original argument of [12, Theorem 2]. Second, the one
832 sketched in the warmup of [1, §3]. We explain why in our opinion this alternative argument
833 is incomplete, then provide a possible fix, at the cost of a suboptimal bound. (iii) In any
834 case, none of the aforementioned arguments that we revisit carries over partial synchrony.
835 Indeed, a player receiving no message *before* GST, is not supposed to output, even if he
836 actually does not know when GST will happen.

837 [For *any* deterministic binary broadcast protocol Π , then *at least* one of the following
838 two statements (9) and (10) is true:

There exists a majority of players q such that:

$$839 \quad C_0(\Pi, q) : \text{If } \{q \text{ receives no message}\}, \text{ then } \{q \text{ ultimately outputs } 0\} \quad (9)$$

840

There exists a majority of players q such that:

$$841 \quad C_1(\Pi, q) : \text{If } \{q \text{ receives no message}\}, \text{ then } \{q \text{ ultimately outputs } 1\} \quad (10)$$

842 This is trivial since every q is supposed to ultimately output some value, 0 or 1, in *any*
843 execution, by the liveness condition of *synchronous* broadcast. Let us insist anyway on
844 a point: in —many— broadcast protocols Π , we may well have that some player q , or
845 even all players, *always* receives messages, in every execution. In which case the condition
846 “If $\{q \text{ receives no message}\}$ ” is *empty* for this Π and q . And thus *both predicates* $C_0(\Pi, q)$ and
847 $C_1(\Pi, q)$ are then trivially true.

848 Let us fix now a protocol Π , and assume without loss of generality that the statement (9)
849 above holds. By the previous point, there a fortiori exists a set V of $\frac{t}{2}$ players, *not* containing
850 the sender, such that: every honest $q \in V$ ultimately outputs 0 in executions —if any— in
851 which he receives no message.]

852 Let us consider first the Environment \mathcal{E} which gives input 1 to the sender and corrupts
 853 V . He instructs every player in V to behave honestly, except that they: ignore the first
 854 $\frac{t}{2}$ messages sent to each of them —i.e. they behave as if they did not receive them— and
 855 they never communicate with each other. By definition of broadcast, all honest players $\mathcal{P} \setminus V$
 856 ultimately output the honest sender’s value 1.

857 Let us now assume by contradiction that, for every Environment, the number of messages
 858 sent by honest players in Π is $\leq \frac{t^2}{4}$. In particular, under our specific \mathcal{E} , honest players would
 859 send $\leq \frac{t^2}{4}$ messages to V . And thus there exists some fixed player $p \in V$ which receives $\leq \frac{t}{2}$
 860 messages from honest players. Note $A(p)$ the set of honest players sending messages to p
 861 under \mathcal{E} , possibly containing the sender, of cardinality at most $|A(p)| \leq \frac{t}{2}$. We can now
 862 consider the Environment \mathcal{E}' that behaves the same than \mathcal{E} with the following differences:
 863 (i) it corrupts all players in V *except* p (ii) and also corrupts $A(p)$, to which he instructs
 864 *not* to send any message to p , but otherwise to behave honestly. Now: on the one hand, [p
 865 ultimately outputs 0 because he belongs to V and does not receive any message]. On the
 866 other hand, consider the remaining honest players $U := \mathcal{P} - \{V \cup A(p)\}$. Their joint history
 867 under \mathcal{E}' is undistinguishable from the one under \mathcal{E} , because, under \mathcal{E} , p *also* ignored the
 868 first $\frac{t}{2}$ messages that he received, and *did not* receive any other message by assumption. So
 869 players in U still ultimately output 0 under \mathcal{E}' .

870 A.2.3.2 Discussing other arguments in the literature for this bound

871 To start with, consider the very end of the proof of [12, Theorem 2] in the original paper.
 872 The authors claim that p does not output at the end of the execution. We do not see why,
 873 since in our opinion, by definition of broadcast, the honest player p *should* anyway ultimately
 874 output some value. We now consider the alternative argument given in warmup of [1, §3].
 875 At the very end the authors mention a “symmetric” scenario, where the sender sends 1
 876 instead, in which the same player p would output 0. But we do not see actually why the
 877 Environment would be able to isolate the *same* player p in this symmetric scenario. One
 878 can manage to keep their overall argumentation by just fixing this point as follows, at the
 879 price of a suboptimal bound. Precisely, assume that only less than $\frac{t^2}{8}$ messages are sent in
 880 the worst case by honest players, instead of the previous $\frac{t^2}{4}$. Consider the two symmetric
 881 executions \mathcal{E}_0 and \mathcal{E}_1 in which the Environment corrupts the same arbitrary set V as before,
 882 and which differ only by the input value of the honest sender: 0 or 1. Then, as in our
 883 counting argument for [12, Theorem 1] above, there exists necessarily a player $p \in V$ which
 884 receives $\leq \frac{t}{2}$ messages in *both* those executions. We can now consider the variant executions
 885 \mathcal{E}'_0 and \mathcal{E}'_1 of \mathcal{E}_0 and \mathcal{E}_1 , where, as before, the environment leaves p uncorrupted. The *same* p
 886 receives no message in both executions, so should ultimately output the *same* value b . We
 887 see that, whatever the value of b , this leads to a safety violation either in \mathcal{E}'_0 or in \mathcal{E}'_1 .

888 A.2.4 [1, Theorem 1]: when the Environment is *strongly rushing*, then 889 the previous [12, Theorem 2] extends to *randomized broadcast*

890 Strongly rushing means that, here, the Environment has the additional power to *withdraw*
 891 messages that were sent by honest players just before it corrupted them. They obtain that
 892 to achieve a probability of failure smaller than $1/2 + \epsilon$, then honest players need sending in

:22 Lower bounds, limits of standalone digital signatures

⁸⁹³ expectation at least $(\epsilon t/2)^2$ messages.