

1 Byzantine Consensus with Local Multicast 2 Channels

3 Muhammad Samir Khan 

4 Department of Computer Science, University of Illinois at Urbana-Champaign, USA

5 Nitin H. Vaidya 

6 Department of Computer Science, Georgetown University, USA

7 — Abstract —

8 Byzantine consensus is a classical problem in distributed computing. Each node in a synchronous
9 system starts with a binary input. The goal is to reach agreement in the presence of Byzantine faulty
10 nodes. We consider the setting where communication between nodes is modelled via an *undirected*
11 communication graph. In the classical *point-to-point* communication model all messages sent on an
12 edge are private between the two endpoints of the edge. This allows a faulty node to *equivocate*, i.e.,
13 lie differently to its different neighbors. Different models have been proposed in the literature that
14 weaken equivocation. In the *local broadcast* model, every message transmitted by a node is received
15 identically and correctly by all of its neighbors. In the *hypergraph* model, every message transmitted
16 by a node on a hyperedge is received identically and correctly by all nodes on the hyperedge. Tight
17 network conditions are known for each of the three cases.

18 We introduce a more general model that encompasses all three of these models. In the *local*
19 *multicast* model, each node u has one or more local multicast channels. Each channel consists of
20 multiple neighbors of u in the communication graph. When node u sends a message on a channel,
21 it is received identically by all of its neighbors on the channel. For this model, we identify tight
22 network conditions for consensus. We observe how the local multicast model reduces to each of
23 the three models above under specific conditions. In each of the three cases, we relate our network
24 condition to the corresponding known tight conditions. The local multicast model also encompasses
25 other practical network models of interest that have not been explored previously, as elaborated in
26 the paper.

27 **2012 ACM Subject Classification** Theory of computation → Distributed algorithms

28 **Keywords and phrases** Byzantine fault, distributed algorithm, consensus, broadcast, multicast

29 **Digital Object Identifier** 10.4230/LIPIcs.DISC.2021.29

30 **Funding** This research is supported in part by the National Science Foundation award 1733872. Any
31 opinions, findings, and conclusions or recommendations expressed here are those of the authors and
32 do not necessarily reflect the views of the funding agencies or the U.S. government.

33 1 Introduction

34 Byzantine consensus is a classical problem in distributed computing introduced by Lamport
35 et al [12, 14]. There are n nodes in a synchronous system. Each node starts with a binary
36 input. At most f of these nodes can be Byzantine faulty, i.e., exhibit arbitrary behavior.
37 The goal of a consensus protocol is for the non-faulty nodes to reach agreement on a single
38 output value in finite time. To exclude trivial protocols, we require that the output must be
39 an input of some non-faulty node.

40 In this paper, we study consensus under the *local multicast* model. We formalize this
41 model in Section 2. Intuitively, nodes are connected via an undirected graph G . A local
42 multicast channel is defined by a sender and a set of receivers. Each node u may potentially
43 serve as the sender on multiple local multicast channels. When node u sends a message on
44 one of its local multicast channels, it is received identically and correctly by all the receivers



© Muhammad Samir Khan and Nitin H. Vaidya;
licensed under Creative Commons License CC-BY 4.0
35th International Symposium on Distributed Computing (DISC 2021).

Editor: Seth Gilbert; Article No. 29; pp. 29:1–29:17



Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

29:2 Byzantine Consensus with Local Multicast Channels

45 in the channel. This model generalizes the following models that have been considered before
46 in the literature.

47 1. *Point-to-point communication model*: In the classical *point-to-point* communication model,
48 each edge uv in the communication graph represents a private link between the nodes u
49 and v . This model is well-studied [1, 4, 12, 13, 14]. It is well-known that $n \geq 3f + 1$ and
50 node connectivity at least $2f + 1$ are both necessary and sufficient in this model.

51 2. *Local broadcast model*: Recently, we [8] studied consensus under the *local broadcast* model
52 [2, 11], where a message sent by any node is received identically by all of its neighboring
53 nodes in the communication graph. We obtained that minimum node degree at least $2f$
54 and node connectivity at least $\lfloor 3f/2 \rfloor + 1$ are both necessary and sufficient for Byzantine
55 consensus under the local broadcast model [8].

56 3. *Hypergraph model*: A *hypergraph* is a generalization of graphs consisting of nodes and
57 hyperedges. Unlike an edge in a graph, a hyperedge can connect any number of nodes.
58 For a communication network modelled as a hypergraph, a message sent by a node u on
59 a hyperedge e (that contains u) is received identically by all nodes in the hyperedge e .
60 Communication networks modelled as hypergraphs have been studied in the literature
61 [6, 7, 15]. Ravikant et al [15] gave tight conditions for Byzantine consensus on $(2, 3)$ -
62 hypergraphs.¹ As discussed in Section 4, this result extends to general undirected
63 hypergraphs as well.

64 The classical point-to-point communication model allows a faulty node to *equivocate* [3],
65 i.e., send conflicting messages to its neighbors without this inconsistency being observed
66 by the neighbors. For example, a faulty node z may tell its neighbor u that it has input 0,
67 but tell another neighbor v that it has input 1. Since messages on each edge are private
68 between the two endpoints, so node u does not overhear the message sent to node v and vice
69 versa. The local broadcast model and the hypergraph model restrict a faulty node's ability
70 to equivocate by detecting such attempts. In the local broadcast model, a faulty node's
71 attempt to equivocate is detected by its neighboring nodes in the communication graph. In
72 the hypergraph model, a faulty node's attempt to equivocate on a hyperedge is detected
73 by the nodes in that hyperedge. In our local multicast model, a faulty node's attempt to
74 equivocate on a single multicast channel is detected by the receivers in that channel.

75 In this work, we introduce the *local multicast* model, that unifies the models identified
76 above, and make the following main contributions:

77 1. **Necessary and sufficient condition for local multicast model**: In Section 3, we
78 present a network condition and show that it is both necessary and sufficient for Byzantine
79 consensus under the local multicast model. The identified condition is inspired by the
80 network conditions for directed graphs [9, 17], where node connectivity does not adequately
81 capture the network requirements for consensus. We present a simple algorithm, inspired
82 by [8, 9, 17].

83 2. **Reductions to the existing models**: The two extremes of the local multicast model
84 are 1) each channel consists of exactly one receiver, and 2) each node has exactly one
85 multicast channel. These correspond to the point-to-point communication model and the
86 local broadcast model, respectively. In Section 4, we discuss how the network condition

¹ i.e., each hyperedge consists of either 2 or 3 nodes.

87 for the local multicast model reduces to the network requirements for the point-to-point
 88 model and the local broadcast model at the two extremes. On the other hand, if the
 89 multicast channels are induced from the hyperedges in a hypergraph, then this corresponds
 90 to the hypergraph model. In this case, the network condition reduces to the network
 91 requirements of the undirected hypergraph model given by Ravikant et al [15]. Moreover,
 92 our algorithm for the local multicast model works for all the three models identified here
 93 as well.

94 **3. Extensions to other models:** The local multicast model also captures some other
 95 models of practical interest (see Section 5). For instance, consider the scenario where
 96 nodes are connected via a wireless network. This can be modelled as local multicast
 97 over a graph G_1 . Separately, the nodes are also connected via a bluetooth network,
 98 modelled using local multicast over a graph G_2 (with the same node set as G_1). Then
 99 the union of these networks $G_1 \cup G_2$ can be captured using the local multicast model as
 100 well. As another example, consider the scenario where nodes are connected via point-
 101 to-point channels. Additionally, nodes are also connected via a wireless network with
 102 local broadcast guarantees. As before, this can also be captured using the local multicast
 103 model. Our algorithm works for these cases as well.

104 In our recent work [10], we have generalized the results in this paper to the *directed* local
 105 multicast model. The directed local multicast model corresponds to directed hypergraphs
 106 where each directed hyperedge models a multicast channel with a single sender and a non-
 107 empty set of receivers. The tight condition obtained in [10] for the directed case is a natural
 108 extension of the tight condition obtained here for the undirected case. The results and proofs
 109 in [10] are more general and encompass the results in this paper.

110 2 System Model and Problem Formulation

111 We consider a synchronous system of n nodes. Nodes communicate using *local multicast*
 112 *channels*. Each node u has a set of multicast channels ζ_u . Each multicast channel $\chi_u \in \zeta_u$
 113 is defined by the sender u and a non-empty set of receivers. For example, $\{v, w\} \in \zeta_u$ is a
 114 multicast channel of sender u with two receivers v and w . By convention used here, u is
 115 not included in the set of receivers. However, trivially, each node receives its own message
 116 transmissions as well. The communication between nodes is bidirectional so that if a node
 117 $v \in \chi_u$ for some channel $\chi_u \in \zeta_u$, then there exists a channel $\chi_v \in \zeta_v$ such that $u \in \chi_v$. A
 118 message m sent by a node u on a multicast channel χ_u is received identically and correctly
 119 by all nodes in χ_u . Moreover, each recipient $v \in \chi_u$ knows that m was sent by u on channel
 120 χ_u . We assume that each multicast channel is a FIFO communication channel.

121 The communication graph $G = (V(G), E(G))$ is an undirected graph where $V(G)$ is the
 122 set of n nodes and $uv \in E(G)$ is an edge of G if and only if there are channels χ_u and χ_v at
 123 nodes u and v , respectively, such that $u \in \chi_v$ and $v \in \chi_u$. Nodes u and v are *neighbors* in G .
 124 Observe that each multicast channel χ_u consists of a non-empty subset of the neighbors of u ,
 125 such that each neighbor of u is in at least one channel in ζ_u .

126 ■ *Neighborhood:* For a set $S \subseteq V(G)$, a node $v \in V(G) - S$ is a neighbor of S if it is
 127 a neighbor of some node $u \in S$. More generally, for two disjoint sets $A, B \subseteq V(G)$,
 128 $\Gamma_G(A, B)$ defined below is the set of neighbors of B in A .

$$129 \quad \Gamma_G(A, B) := \{u \in A \mid \exists v \in B : uv \in E(G)\}.$$

29:4 Byzantine Consensus with Local Multicast Channels

130 ■ *Adjacent*: For two disjoint sets $A, B \subseteq V(G)$, we use $A \rightarrow_G B$ (read as “adjacent”
 131 to B in G) to denote that either
 132 (i) $B = \emptyset$, or
 133 (ii) nodes in B have at least $f + 1$ neighbors in A in the graph G , i.e.,

134 $|\Gamma_G(A, B)| \geq f + 1.$

135 A *Byzantine faulty* node may exhibit arbitrary behavior. In *Byzantine consensus problem*
 136 each node starts with a binary input and must output a binary value satisfying the following
 137 constraints, in the presence of up to f Byzantine faulty nodes.

138 1. **Agreement**: All non-faulty nodes must output the same value.
 139 2. **Validity**: If a non-faulty node outputs $b \in \{0, 1\}$, then at least one non-faulty node must
 140 have input b .
 141 3. **Termination**: All non-faulty nodes must decide in finite time.

142 It is easy to show that $f < n$ is necessary for Byzantine consensus. So we assume $f < n$
 143 throughout the paper.

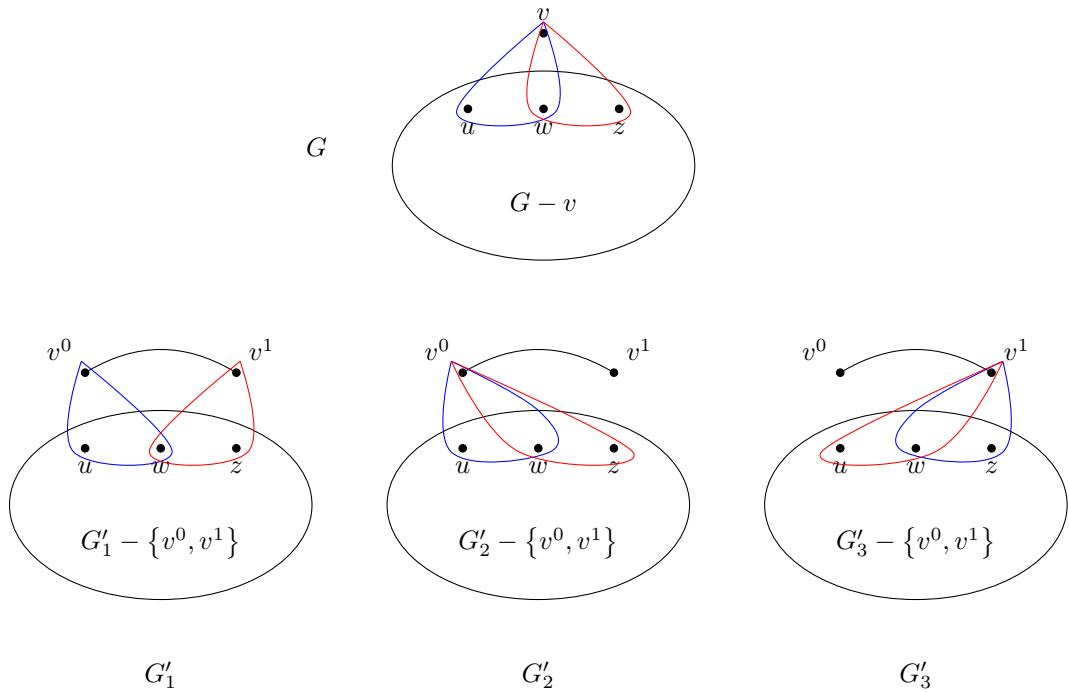
144 Node split

145 We now introduce the notion of a *node split* that is used to specify the necessary and sufficient
 146 condition under the local multicast model. As seen later, we will use the notion of node split
 147 to simulate possible equivocation by a faulty node. Intuitively, by splitting a node v , we
 148 are creating two copies of v and dividing up the channels amongst the two copies. Figure 1
 149 shows two examples of node split. Formally, splitting a node v in G creates a new graph G'
 150 as follows.


151 ■ The node v is replaced by two nodes v^0 and v^1 .
 152 ■ We add an edge v^0v^1 to $E(G')$.
 153 ■ We add a multicast channel $\{v^1\}$ to v^0 and a multicast channel $\{v^0\}$ to v^1 .
 154 ■ For every multicast channel χ_v of node v in G , choose exactly one of v^0 and v^1 as node
 155 v' . Create a multicast channel $\chi'_{v'}$ of v' with $\chi'_{v'} = \{u \mid u \in \chi_v\}$, i.e., each neighbor of v
 156 in χ_v is assigned to $\chi'_{v'}$.
 157 ■ The above step adds edges to $E(G')$, of the form uv' such that $v' \in \{v^0, v^1\}$, but v' is
 158 not assigned to any multicast channel at node u . We specify these assignments as follows.
 159 Consider an edge uv' , for $v' \in \{v^0, v^1\}$, in G' . For each multicast channel χ_u of node u
 160 in G , such that $v \in \chi_u$, add v' to the corresponding multicast channel χ'_u in G' . Now
 161 each neighbor w of u in G' is part of at least one multicast channel at node u .

162 Observe that for every node $u \in V(G')$, each of its multicast channels in G' corresponds
 163 to a single multicast channel in G , except for the two channels $\{v^1\}$ and $\{v^0\}$ at nodes v^0
 164 and v^1 , respectively (where node v was split). Similarly, for every node $u \in V(G)$, each of
 165 its multicast channels in G corresponds to a single multicast channel in G' .

166 To split two nodes u and v in G , we first split u to obtain G' . We then split v to obtain
 167 G'' from G' . The order of splits does not matter. This process naturally extends to splitting
 168 multiple nodes as well. For a set $F \subseteq V(G)$, let $\Lambda_F(G)$ be the set of all graphs that can be



(a) Splitting a single node v . Only the channels in ζ_v are drawn here. There are two channels in ζ_v : $\{u, w\}$ and $\{w, z\}$, drawn with blue and red colors, respectively. There are three possible graphs in $\Lambda_{\{v\}}(G)$, other than G , corresponding to the assignment of channels when v is split into v^0 and v^1 . These are depicted as G'_1 , G'_2 , and G'_3 .



(b) Splitting two nodes u, v in a 4-node graph G . Directed edges of the same color, pointing out from the same sender node, represent a single channel. G' is obtained by splitting nodes u and v into u^0, u^1 and v^0, v^1 , respectively. The cyan channel is assigned to v^1 , the violet channel is assigned to v^0 , and the red channel is assigned to u^0 .

■ **Figure 1** Examples of the node split operation.

29:6 Byzantine Consensus with Local Multicast Channels

169 obtained from G by splitting some subset of nodes in the set F . For a graph $G' \in \Lambda_F(G)$,
170 we use F' to denote the set of nodes in G' that correspond to nodes in F in G , i.e.,

$$171 \quad F' := (V(G') \cap F) \cup (V(G') - V(G)).$$

172 Note that there are two choices in the node split operation above which give rise to all
173 the graphs in $\Lambda_F(G)$:

- 174 1. choice of which nodes in F to split, and
- 175 2. assignment of multicast channels for each split node.

176 As needed, we will occasionally clarify these choices to specify how a graph $G' \in \Lambda_F(G)$ was
177 constructed by splitting some nodes in F .

178 3 Main Result

179 The main result of this paper is a tight characterization of network requirements for Byzantine
180 consensus under the local multicast model. Consider a graph $G' \in \Lambda_F(G)$ obtained from G
181 by splitting some nodes in a set F . Recall that we use F' to denote the set of nodes in G'
182 that correspond to nodes in F in G .

183 ▶ **Theorem 1.** *Under the local multicast model, Byzantine consensus tolerating at most f
184 faulty nodes is achievable on graph G if and only for every $F \subseteq V(G)$ of size at most f ,
185 every $G' \in \Lambda_F(G)$ satisfies the following: for every partition² (L, C, R) of $V(G')$, either*

- 186 1. $L \cup C \rightarrow_{G'} R - F'$, or
- 187 2. $R \cup C \rightarrow_{G'} L - F'$.

188 While we allow a partition to have empty parts, the interesting partitions are those where
189 both L and R are non-empty, but C can be possibly empty. In Section 4, we show that when
190 the local multicast model corresponds to the point-to-point, local broadcast, or hypergraph
191 model, the above condition reduces to the corresponding known tight network conditions in
192 each of the three cases.

193 We prove the necessity of Theorem 1 in Section 6. In Section 7, we give an algorithm to
194 constructively show the sufficiency. The above condition is similar to the network condition
195 for directed graphs in the point-to-point communication model [16, 17] and in the local
196 broadcast model [9]. Note that [9, 16, 17] deal with consensus on arbitrary *directed* graphs,
197 where connectivity constraints do not adequately capture the tight network requirements. In
198 this paper, we are interested in undirected graphs. However, since the local multicast model
199 is quite general and captures various models with different connectivity requirements, it is
200 plausible that no concise network connectivity property will be able to properly characterize
201 the tight condition.

202 For convenience, we give a name to the condition in Theorem 1.

203 ▶ **Definition 2.** *A graph G satisfies condition LCR with parameter F if for every $G' \in \Lambda_F(G)$
204 and every partition (L, C, R) of $V(G')$, we have that either*

² with a slight abuse of terminology, we allow a partition of a set to have empty parts.

205 1. $L \cup C \rightarrow_{G'} R - F'$, or
 206 2. $R \cup C \rightarrow_{G'} L - F'$.

207 We say that G satisfies condition LCR, if G satisfies condition LCR with parameter F for
 208 every set $F \subseteq V(G)$ of cardinality at most f .

209 4 Reductions to Other Models

210 In this section, we discuss how condition LCR relates to the tight conditions for the classical
 211 point-to-point communication model, the local broadcast model, and the hypergraph model.

212 Point-to-point Channels

213 The classical point-to-point communication model corresponds to the case where each
 214 multicast channel in the graph consists of a single receiver node, so that the communication on
 215 an edge uv is private between the two nodes u and v . Under the point-to-point communication
 216 model, it is well known that $n \geq 3f + 1$ [5, 12, 14] and node connectivity at least $2f + 1$
 217 [4, 5] are both necessary and sufficient for consensus on arbitrary undirected graphs.

218 When G has only point-to-point channels, i.e., each multicast channel consists of a single
 219 receiver node, then G satisfies condition LCR if and only if $n \geq 3f + 1$ and G has node
 220 connectivity $\geq 2f + 1$. We prove this formally in [10]. Therefore, the two models are
 221 equivalent when only point-to-point channels are present.

222 Local Broadcast

223 The local broadcast model corresponds to the other extreme where each node in the graph
 224 has exactly one multicast channel, so that the messages transmitted by a node u are received
 225 identically and correctly by all neighbors of u . Under the local broadcast model, we [8]
 226 showed that node degree at least $2f$ and connectivity at least $\lfloor 3f/2 \rfloor + 1$ are both necessary
 227 and sufficient for consensus on arbitrary undirected graphs.

228 When G has only local broadcast channels, i.e., each node has a single multicast channel,
 229 then G satisfies condition LCR if and only if G has minimum node degree $\geq 2f$ and node
 230 connectivity $\geq \lfloor 3f/2 \rfloor + 1$. We prove this formally in [10]. Therefore, the two models are
 231 equivalent when only local broadcast channels are present.

232 Hypergraphs

233 The last model we consider in this section is the hypergraph model. In a hypergraph
 234 $H = (V(H), E(H))$, each hyperedge $e \in E(H)$ is a subset of nodes $e \subseteq V(H)$. A hyperedge
 235 $e \in E(H)$ is called an $|e|$ -hyperedge. Each hyperedge is effectively a multicast channel, i.e.,
 236 a message sent by a node u on an edge $e \supseteq \{u\}$ is received identically and correctly by all
 237 nodes $v \in e$. However, any node on a hyperedge can act as a sender for this channel. In
 238 our local multicast model with communication graph G , this corresponds to the case where,
 239 for every pair of nodes u, v and multicast channel χ_u of u such that $v \in \chi_u$, there exists a
 240 channel χ_v of v such that $\chi_v = (\chi_u \cup \{u\}) - v$.

241 Ravikant el. al. [15] obtained tight conditions for the hypergraph model. We observe that
 242 while the conditions were presented as a tight characterization for $(2, 3)$ -hypergraphs³ in [15],

³ H is a $(2, 3)$ -hypergraph if each hyperedge is either a 2-hyperedge or a 3-hyperedge.

243 they also hold for general hypergraphs. In our local multicast model, when the communication
 244 graph G and its local multicast channels correspond to an undirected hypergraph, then
 245 condition LCR reduces to the tight conditions for hypergraphs given in [15]. The formal
 246 proof is given in [10].

247 5 Application to New Models

248 As mentioned in Section 1, the local multicast model also encompasses some other network
 249 models of practical interest that, to the best of our knowledge, have not been considered
 250 before in the literature. Suppose the n nodes are connected via a local multicast network
 251 with graph G_1 . For example, network connectivity in G_1 can be via point-to-point links or
 252 via wireless channels modelled as local broadcast. Additionally, the n nodes are connected
 253 via another local multicast network with graph G_2 . For example, G_2 may correspond to
 254 a wireless network with different frequencies and/or technologies. The complete system,
 255 where nodes can communicate on channels in G_1 as well as on channels in G_2 , can also be
 256 characterized by the local multicast model. We omit details for brevity, but this corresponds
 257 to the natural union of G_1 and G_2 , with each node now having access to its multicast channels
 258 in G_1 as well as its multicast channels in G_2 .

259 6 Necessity of Condition in Theorem 1

260 Intuitively, consider a set $F \subseteq V(G)$ of size at most f , such that the graph G violates
 261 condition LCR with parameter F . With F as a candidate faulty set, the splitting of nodes
 262 in F captures possible equivocation by nodes in F : a faulty node can behave as if it has
 263 input 0 on some of its multicast channels and behave as if it has input 1 on the other
 264 multicast channels. Let $G' \in \Lambda_F(G)$ be a graph obtained by splitting nodes in F . We use
 265 F' to denote the nodes in G' that correspond to nodes in F in G . Suppose (L, C, R) is a
 266 partition of G' . Now consider the execution where non-faulty nodes in L have input 0. Since
 267 $R \cup C \not\rightarrow_{G'} L - F'$, nodes in $L - F'$ can not distinguish between F and its neighbors in
 268 $R \cup C$, i.e., $\Gamma_{G'}(R \cup C, L - F')$ as the set of faulty nodes. So non-faulty nodes in L are stuck
 269 with outputting 0 in this case. Similarly if non-faulty nodes in R have input 1, then they
 270 have no choice but to output 1, creating the desired contradiction.

271 A formal necessity proof is given in [10] for the *directed* local multicast model, which
 272 generalizes the *undirected* local multicast model considered in this paper. It follows the
 273 standard state machine based approach [1, 4, 5], similar to [9, 17]. Suppose there exists a set
 274 $F \subseteq V(G)$, of size at most f , such that G does not satisfy condition LCR with parameter F ,
 275 but there exists an algorithm \mathcal{A} that solves consensus on G . Algorithm \mathcal{A} outlines a procedure
 276 \mathcal{A}_u for each node u that describes u 's state transitions, as well as messages transmitted on
 277 each channel of u in each round. Now there exists a graph $G' \in \Lambda_F(G)$ and a partition
 278 of $V(G')$ that does not satisfy the requirements of condition LCR. To create the required
 279 contradiction, we work with an algorithm for G' instead of \mathcal{A} . To see why this works, observe
 280 that an algorithm \mathcal{A} on graph G can be adapted to create an algorithm \mathcal{A}' for a graph
 281 $G' \in \Lambda_F(G)$ as follows. Each round i in the algorithm \mathcal{A} is now split into two sub-rounds $i(a)$
 282 and $i(b)$ in \mathcal{A}' . We consider each of these rounds separately and specify the corresponding
 283 steps for each node in G' for the algorithm \mathcal{A}' .

284 ■ *Round $i(a)$:* Each node $v \in V(G') \cap V(G)$ that was not split runs \mathcal{A}_v as specified for
 285 round i . For a node $v \in V(G') - V(G)$ that was split into $v^0, v^1 \in V(G')$, both v^0 and

286 v^1 run \mathcal{A}_v for round i with the following modification. Consider a multicast channel
 287 $\chi_v \in \zeta_v$ of node v in G . Let χ'_{v^0} (resp. χ'_{v^1}) be the corresponding multicast channel in G'
 288 at node v^0 (resp. v^1). If the algorithm \mathcal{A}_v wants to transmit a message on χ_v , then v^0
 289 (resp. v^1) sends the message on χ'_{v^0} (resp. χ'_{v^1}), while v^1 (resp. v^0) ignores this message
 290 transmission. Observe that, for any node $u \in \chi_v$, u receives messages on the channel
 291 from exactly one of v^0 and v^1 .

292 ■ *Round $i(b)$:* This round is reserved for the split nodes. Consider a node $v \in F - V(G')$
 293 that was split into $v^0, v^1 \in V(G')$. Node v^0 forwards all messages it received in round
 294 $i(a)$ to v^1 and v^1 forwards all messages it received in round $i(a)$ to v^0 . This allows both
 295 v^0 and v^1 to run \mathcal{A}_v in the next round.

296 Now, \mathcal{A}' might not solve consensus on G' , or may not even terminate. However, as long
 297 as care is taken with regards to which nodes are allowed to be faulty in G' and the input
 298 of the split nodes, the guarantees for \mathcal{A} will imply that \mathcal{A}' does indeed terminate and solve
 299 consensus on G' . In particular, we want that

300 1. the faulty nodes in G' correspond to at most f nodes in G ,
 301 2. for each node $v \in F - V'$ that was split into $v^0, v^1 \in V'$, either
 302 a. both v^0 and v^1 have the same input, or
 303 b. at least one of v^0 and v^1 is faulty.

304 So for necessity, it is enough to show that no algorithm exists for a hypergraph $G' \in \Lambda_F(G)$,
 305 under the two conditions identified above. We formalize this property and use it in the
 306 formal necessity proof in [10].

307 7 Algorithm for the Local Multicast Model

308 To prove the sufficiency portion of Theorem 1, we work with a different network condition,
 309 which we will be equivalent to condition LCR. We first introduce some notation that is used
 310 in the algorithm. For a set of nodes $U \subseteq V(G)$, we use $G[U]$ to denote the subgraph induced
 311 by the nodes in U . The multicast channels in $G[U]$ are obtained from the multicast channels
 312 in G by removing nodes in $V(G) - U$ from each channel, with some channels possibly being
 313 deleted entirely. We use $G - U$ to denote the subgraph $G[V(G) - U]$.

314 A *path* is a sequence of distinct nodes such that if u precedes v in the sequence, then v is
 315 a neighbor of u in G (i.e., uv is an edge). For a path P and node z , we use $z \cdot P$ to denote
 316 the path obtained by prefixing the node z to P .

317 ■ *uv-paths:* For two nodes $u, v \in V(G)$, a *uv-path* P_{uv} is a path from u to v . u is called
 318 the *source* and v the *terminal* of P_{uv} . Any other node in P_{uv} is called an *internal* node
 319 of P_{uv} . Two *uv*-paths are *node-disjoint* if they do not share a common internal node.

320 ■ *Uv-paths:* For a set $U \subset V(G)$ and a node $v \notin U$, a *Uv-path* is a *uv-path* for some node
 321 $u \in U$. All *Uv*-paths have v as terminal. Two *Uv*-paths are *node-disjoint* if they do not
 322 have any nodes in common except the terminal node v . In particular, two node-disjoint
 323 *Uv*-paths have different source nodes. By definition, the number of disjoint *Uv*-paths
 324 is upper bounded by the size of the set U . Note the difference in definition between
 325 node-disjoint *uv*-paths and node-disjoint *Uv*-paths when $U = \{u\}$ is a singleton set. The
 326 former requires only internal nodes to be different, while the latter needs to have different
 327 source nodes as well. For the former, there can be more than one such node-disjoint path,
 328 while for the latter, there is at most one.

29:10 Byzantine Consensus with Local Multicast Channels

329 ■ *Propagate*: For two disjoint node sets $A, B \subseteq V(G)$, we use $A \rightsquigarrow_G B$ (read as A “propagates” to B in G) to denote that either
 330 (i) $B = \emptyset$, or
 332 (ii) for every $v \in B$, there exist at least $f + 1$ node-disjoint Av -paths in the graph $G[A \cup B]$.

333 We now give a different network condition which is equivalent to condition LCR, but will
 334 be useful for specifying an algorithm for the local multicast model and proving its correctness.
 335 Recall that we use F' to denote the set of nodes in G' corresponding to nodes in F in G .

336 ► **Definition 3.** A graph G satisfies condition AB with parameter F if for every $G' \in \Lambda_F(G)$
 337 and every partition (A, B) of $V(G')$, we have that either

338 1. $A \rightsquigarrow_{G'} B - F'$, or
 339 2. $B \rightsquigarrow_{G'} A - F'$.

340 We say that G satisfies condition AB, if G satisfies condition AB with parameter F for every
 341 set $F \subseteq V$ of cardinality at most f .

342 ► **Theorem 4.** A graph G satisfies condition LCR if and only if G satisfies condition AB.

343 We skip the proof of Theorem 4, which is (almost) identical to proof of Theorem 5.2 in
 344 [9]. We show the sufficiency of condition AB (and hence condition LCR) constructively. For
 345 the rest of this section, we assume that G satisfies condition AB. The proposed algorithm
 346 is given in Algorithm 1. It draws inspiration from algorithms in [8, 9, 17]. Each node v
 347 maintains a binary state variable γ_v , which we call v 's γ value. Each node v initializes γ_v to
 348 be its input value.

349 The nodes use “flooding” to communicate with the rest of the nodes. We refer the reader
 350 to [8, 9] for details about the flooding primitive. Briefly, when a node u wants to flood a
 351 binary value $b \in \{0, 1\}$, it transmits b to all of its neighbors, who forward it to their neighbors,
 352 and so forth. If a node u receives a message on channel χ , then u appends the channel id of
 353 χ when forwarding the message to its neighbors. By adding some simple sanity checks, one
 354 can assume that even a faulty node v does indeed transmit some value when it is v 's turn
 355 to forward a message. In at most n synchronous rounds, the value b will be “flooded” in G .
 356 However, faulty nodes may tamper messages when forwarding, so some nodes may receive a
 357 value $\bar{b} \neq b$ along paths that contain faulty nodes.

358 The algorithm proceeds in phases. Every iteration of the `for` loop (starting at line 3) is
 359 a phase numbered $1, \dots, 2^f$. Let F^* denote the actual set of faulty nodes. Each iteration
 360 of the `for` loop, i.e. phase > 0 , considers a candidate faulty set F . In this iteration, nodes
 361 attempt to reach consensus, by updating their γ state variables, assuming the candidate set
 362 F is indeed faulty. Let Z and N be the set of nodes in $G - F$ that have their state variable
 363 set to 0 and 1, respectively, at the beginning of the iteration. Each iteration has three steps.

364 ■ In step (a), each node v floods its γ_v value.
 365 ■ In step (b), based on the values received during flooding, each node v creates its estimate
 366 of the sets Z and N , by ignoring all paths that pass through the candidate faulty set F ,
 367 i.e., paths that have internal nodes from F . This estimate is created in a manner so that
 368 1. when $F \neq F^*$, this estimate may be incorrect, but
 369 2. when $F = F^*$, this estimate is indeed correct.

■ **Algorithm 1** Proposed algorithm for Byzantine consensus under the local multicast model: Steps performed by node v are shown here.

- 1 Each node has a binary input value in $\{0, 1\}$.
- 2 Each node v maintains a binary state $\gamma_v \in \{0, 1\}$, initialized to the input value of v .
- 3 **For each** $F \subseteq V$ **such that** $|F| \leq f$:

- 4 **Step (a):** Flood value γ_v .
- 5 **if** $v \in F$ **then** skip steps (b) and (c)
- 6 **Step (b):**
- 7 Create a graph G'_v by splitting all nodes in F as follows. Set

$$F' := \{u^0 \mid u \in F\} \cup \{u^1 \mid u \in F\} \quad \text{and} \quad V(G'_v) := (V(G) - F) \cup F'.$$

The edges and channels of G'_v are as determined by the split operation, with the following choices: **For every node** $z \in F$ **and a multicast channel** $\chi_z \in \zeta_z$:

- 8 **if** $\exists w \in \chi_z$ **such that** $w \in V(G) - F$ **then**
 - 9 identify a single wv -path P_{wv} in $G - F$ (Lemma 7).
 - 10 **if** v received 0 from z along the path $z \cdot P_{wv}$ in **step (a)**, such that the initial message was sent by node z on channel χ_z **then** assign χ_z to z^0 .
 - 11 **else** assign χ_z to z^1 .
 - 12 **else**
 - 13 assign χ_z to z^1 .

14 For each node $u \in V - F$, identify a single uv -path P_{uv} in $G - F$ (Lemma 7).

Note that path P_{vv} trivially exists (P_{vv} contains only v). Initialize a partition (Z_v, N_v) of $V(G'_v)$ as follows,

$$Z_v := \{u^0 \mid u \in F\} \cup \{u \in V - F \mid v \text{ received value 0 along } P_{uv} \text{ in step (a)}\},$$

$$N_v := \{u^1 \mid u \in F\} \cup (V - F - Z_v).$$

Step (c):

- 15 **if** $Z_v \rightsquigarrow_{G'_v} N_v - F$ **then** set $A_v := Z_v$ and $B_v := N_v$
- 16 **else** set $A_v := N_v$ and $B_v := Z_v$
- 17 **if** $v \in B_v - F$ **then**
 - 18 // by construction, the paths of interest in G naturally correspond to paths in G'_v .
 - 19 **if** in **step (a)**, v received a value $\delta \in \{0, 1\}$ identically along any $f + 1$ node-disjoint A_vv -paths in the graph $G'_v[A \cup (B - F)] = G'_v - (B_v \cap F)$ **then**
 - 20 set $\gamma_v := \delta$

20 Output γ_v

29:12 Byzantine Consensus with Local Multicast Channels

370 ■ In **step (c)**, based on the estimates created in **step (b)**, a node v may update its γ_v
371 value. The rules for updates ensure that

372 1. when $F \neq F^*$, for each non-faulty node v , its state γ_v at the end of the iteration equals
373 the γ value of some non-faulty node at the beginning of the iteration (Lemma 5).

374 2. when $F = F^*$, all non-faulty nodes have identical γ values at the end of this iteration
375 (Lemma 6).

376 At the end, after all iterations of the main **for** loop, each output node v outputs its γ_v
377 value.

378 The correctness of Algorithm 1 relies on the following two key lemmas, which are proven
379 in Section A. Recall that we use F^* to denote the actual set of faulty nodes.

380 ▶ **Lemma 5.** *For a non-faulty node $v \in V - F^*$, its state γ_v at the end of any given phase
381 of Algorithm 1 equals the state of some non-faulty node at the start of that phase.*

382 ▶ **Lemma 6.** *Consider a phase > 0 of Algorithm 1 wherein $F = F^*$. At the end of this
383 phase, every pair of non-faulty nodes $u, v \in V - F^*$ have identical state, i.e., $\gamma_u = \gamma_v$.*

384 Lemma 5 ensures validity, i.e., that the output of each non-faulty node is an input of
385 some non-faulty node. It also ensures that agreement among non-faulty nodes, once achieved,
386 is not lost. Lemma 6 ensures that agreement is reached in at least one phase of the algorithm.
387 These two lemmas imply correctness of Algorithm 1 as shown in Section A.

388 8 Conclusion

389 In this paper, we introduced the local multicast model which, to the best of our knowledge,
390 has not been studied before in the literature. The local multicast model encompasses the
391 point-to-point, local broadcast, and hypergraph communication models, as well as some new
392 models which have not been considered before. We identified a tight network condition for
393 Byzantine consensus under the local multicast model, along the lines of [9, 17], and proved
394 its necessity and sufficiency. When the local multicast model represents one of point-to-point,
395 local broadcast, or hypergraph communication models, we showed how the identified network
396 condition reduces to the known tight requirements for the corresponding case.

397 A natural extension to complete the local multicast model is to consider a *directed*
398 communication graph, which corresponds to *directed* hypergraphs, and generalizes the
399 directed cases of point-to-point and local broadcast models. In our recent work [10], we have
400 extended the results in this paper to the directed setting. The natural extension of condition
401 LCR to the directed case is the tight network condition for directed local multicast. We refer
402 the reader to [10] for more details.

403 References

404 1 Hagit Attiya and Jennifer Welch. *Distributed Computing: Fundamentals, Simulations and*
405 *Advanced Topics*. John Wiley & Sons, Inc., USA, 2004.

406 2 Vartika Bhandari and Nitin H. Vaidya. On reliable broadcast in a radio network. In *Proceedings*
407 *of the Twenty-fourth Annual ACM Symposium on Principles of Distributed Computing*, PODC
408 '05, pages 138–147, New York, NY, USA, 2005. ACM. URL: <http://doi.acm.org/10.1145/1073814.1073841>, doi:10.1145/1073814.1073841.

410 3 Byung-Gon Chun, Petros Maniatis, Scott Shenker, and John Kubiatowicz. Attested append-
411 only memory: Making adversaries stick to their word. *SIGOPS Oper. Syst. Rev.*, 41(6):189–
412 204, October 2007. URL: <http://doi.acm.org/10.1145/1323293.1294280>, doi:10.1145/1323293.1294280.

413

414 4 Danny Dolev. The byzantine generals strike again. *Journal of Algorithms*, 3(1):14 –
415 30, 1982. URL: <http://www.sciencedirect.com/science/article/pii/0196677482900049>,
416 doi:[https://doi.org/10.1016/0196-6774\(82\)90004-9](https://doi.org/10.1016/0196-6774(82)90004-9).

417 5 Michael J. Fischer, Nancy A. Lynch, and Michael Merritt. Easy impossibility proofs for
418 distributed consensus problems. *Distributed Computing*, 1(1):26–39, Mar 1986. doi:10.1007/
419 BF01843568.

420 6 Mattias Fitzi and Ueli Maurer. From partial consistency to global broadcast. In *Proceedings*
421 of the *Thirty-second Annual ACM Symposium on Theory of Computing*, STOC '00, pages 494–
422 503, New York, NY, USA, 2000. ACM. URL: <http://doi.acm.org/10.1145/335305.335363>,
423 doi:10.1145/335305.335363.

424 7 Alexander Jaffe, Thomas Moscibroda, and Siddhartha Sen. On the price of equivocation
425 in byzantine agreement. In *Proceedings of the 2012 ACM Symposium on Principles of*
426 *Distributed Computing*, PODC '12, pages 309–318, New York, NY, USA, 2012. ACM. URL:
427 <http://doi.acm.org/10.1145/2332432.2332491>, doi:10.1145/2332432.2332491.

428 8 Muhammad Samir Khan, Syed Shalan Naqvi, and Nitin H. Vaidya. Exact Byzantine Consensus
429 on Undirected Graphs under Local Broadcast Model. In *Proceedings of the 2019 ACM*
430 *Symposium on Principles of Distributed Computing*, PODC '19, page 327–336, New York, NY,
431 USA, 2019. Association for Computing Machinery. doi:10.1145/3293611.3331619.

432 9 Muhammad Samir Khan, Lewis Tseng, and Nitin H. Vaidya. Exact Byzantine Consensus
433 on Arbitrary Directed Graphs Under Local Broadcast Model. In *23rd International*
434 *Conference on Principles of Distributed Systems (OPODIS 2019)*, volume 153 of *Leibniz*
435 *International Proceedings in Informatics (LIPIcs)*, pages 30:1–30:16, Dagstuhl, Germany, 2020.
436 Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. URL: <https://drops.dagstuhl.de/opus/volltexte/2020/11816>, doi:10.4230/LIPIcs.OPODIS.2019.30.

437

438 10 Muhammad Samir Khan and Nitin H. Vaidya. Byzantine consensus under directed hypergraphs.
439 arXiv report under preparation, 2021.

440 11 Chiu-Yuen Koo. Broadcast in radio networks tolerating byzantine adversarial behavior.
441 In *Proceedings of the Twenty-third Annual ACM Symposium on Principles of Distributed*
442 *Computing*, PODC '04, pages 275–282, New York, NY, USA, 2004. ACM. URL: <http://doi.acm.org/10.1145/1011767.1011807>, doi:10.1145/1011767.1011807.

443

444 12 Leslie Lamport, Robert Shostak, and Marshall Pease. The byzantine generals problem. *ACM*
445 *Trans. Program. Lang. Syst.*, 4(3):382–401, July 1982. URL: <http://doi.acm.org/10.1145/357172.357176>,
446 doi:10.1145/357172.357176.

447 13 Nancy A. Lynch. *Distributed Algorithms*. Morgan Kaufmann Publishers Inc., San Francisco,
448 CA, USA, 1996.

449 14 M. Pease, R. Shostak, and L. Lamport. Reaching agreement in the presence of faults.
450 *J. ACM*, 27(2):228–234, April 1980. URL: <http://doi.acm.org/10.1145/322186.322188>,
451 doi:10.1145/322186.322188.

452 15 D. V. S. Ravikant, V. Muthuramakrishnan, V. Srikanth, K. Srinathan, and C. Pandu Rangan.
453 On byzantine agreement over (2,3)-uniform hypergraphs. In *Distributed Computing*, pages
454 450–464, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.

455 16 Lewis Tseng and Nitin Vaidya. Exact byzantine consensus in directed graphs. *arXiv preprint*
456 *arXiv:1208.5075*, 2014. arXiv:1208.5075.

457 17 Lewis Tseng and Nitin H. Vaidya. Fault-tolerant consensus in directed graphs. In *Proceedings*
458 of the *2015 ACM Symposium on Principles of Distributed Computing*, PODC '15, page 451–460,
459 New York, NY, USA, 2015. Association for Computing Machinery. doi:10.1145/2767386.
460 2767399.

461 **A Proof of Correctness of Algorithm 1**

462 In this section, we show correctness of Algorithm 1 when the communication graph G
 463 satisfies condition AB. For the rest of this section, we assume that G satisfies condition AB.
 464 Throughout this section, we use F^* to denote the actual set of faulty nodes. We first prove
 465 Lemma 5.

466 **Proof of Lemma 5.** Fix a phase > 0 . Note that a node updates its state only in **step (c)**.
 467 Suppose a node v updates its state γ_v to α . Then, as per the update rules in **step (c)**,
 468 v must have received the value α identically along $f + 1$ node-disjoint A_vv -paths in **step (a)**.
 469 Since there are at most f faulty nodes, so at least one of these paths, say P , must
 470 have neither any faulty internal node nor a faulty source node. Since α was received along
 471 P , which has only non-faulty internal nodes, so the source node of P , say u , flooded α in
 472 **step (a)** of this phase. Since u is non-faulty, so γ_u had value α at the start of this phase.
 473 Therefore, the state of node v at the end of this phase equals the state of a non-faulty node
 474 u at the start of this phase. \blacktriangleleft

475 Before proving Lemma 9, we need some intermediate results. We first show that in every
 476 iteration of the main **for** loop, the paths in **step (b)** do exist.

477 **► Lemma 7.** *In any phase > 0 of the algorithm with a candidate faulty set F , for any two
 478 nodes $u, v \in V(G) - F$, there exists a uv -path in $G - F$.*

479 **Proof.** Suppose for the sake of contradiction that there exist two nodes $u, v \in V(G) - F$
 480 such that there is no uv -path in $G - F$. Let A be the set of nodes that are reachable by node
 481 u in $G - F$, and let $B = V - A$. Note that

- 482 (i) $|F| \leq f$,
- 483 (ii) $u \in A = A - F$ so that $A - F \neq \emptyset$, and
- 484 (iii) $v \in B - F$ so that $B - F \neq \emptyset$.

485 Now, there are no edges between A and $B - F$. Since $|F| \leq f$, so there are at most f
 486 node-disjoint Av -paths and at most f node-disjoint Bu -paths in graph G . Therefore, we
 487 have

- 488 1. $A \not\sim_G B - F$, and
- 489 2. $B \not\sim_G A - F$.

490 Since $G \in \Lambda_F(G)$, so condition AB is violated, a contradiction. \blacktriangleleft

491 When a non-faulty node wants to flood a value $b \in \{0, 1\}$, it sends a single value b on
 492 all of its multicast channels. But a faulty node might send different messages on different
 493 channels. Note however, that even a faulty node must send the exact same value on a single
 494 multicast channel.

495 **► Lemma 8.** *Consider a phase > 0 of Algorithm 1 wherein $F = F^*$. For any two non-faulty
 496 nodes $u, v \in V(G) - F^*$, we have $G'_u = G'_v$ in **step (b)** of this phase. Furthermore, if in
 497 **step (a)** of this phase, faulty node $z \in F^*$ transmitted 0 (resp. 1) on one of its channels
 498 $\chi_z \in \zeta_z$, such that $\chi_z - F^*$ is non-empty, then in **step (b)** of this phase χ_z is assigned to
 499 z^0 (resp. z^1) in $G'_u = G'_v$.*

500 **Proof.** Consider the phase where $F = F^*$ and any two non-faulty nodes $u, v \in V(G) - F^*$.
 501 Observe that the node set of the two graphs G'_u and G'_v are the same. For the edges and
 502 channels, by construction, it is sufficient to show that for any $z \in F^*$, the assignment of
 503 multicast channels to z^0 and z^1 in the split operation is the same in G'_u as in G'_v . Consider
 504 an arbitrary node $z \in F^*$ and a multicast channel $\chi_z \in \zeta_z$ at node z . There are two cases to
 505 consider:

506 ■ Case 1: There exists a node $w \in \chi_z$ such that $w \in V(G) - F^*$.
 507 Let w be any arbitrary such node. By Lemma 7, there exists a wu -path (resp. wv -path)
 508 in $G - F^*$. Let P_{wu} (resp. P_{wv}) be any arbitrary wu -path (resp. wv -path) identified
 509 by u (resp. v) in line 9. Note that P_{wu} (resp. P_{wv}) does not contain *any* faulty nodes.
 510 Therefore, a message transmitted by z on χ_z , is received by u (resp. v) along $z \cdot P_{wu}$
 511 (resp. $z \cdot P_{wv}$) untampered. Therefore, in **step (a)**, if z transmitted 0 on channel χ_z ,
 512 then u (resp. v) received value 0 from z along $z \cdot P_{wu}$ (resp. $z \cdot P_{wv}$). So, in line 11,
 513 node u (resp. node v) assigns χ_z to z^0 in G'_u (resp. G'_v). Similarly, if z transmitted 1 on
 514 channel χ_z in **step (a)**, then both u and v assign χ_z to z^1 in G'_u and G'_v , respectively.

515 ■ Case 2: There does not exist any node $w \in \chi_z$ such that $w \in V(G) - F^*$.
 516 In this case, in line 13, both u and v assign χ_z to z^1 in G'_u and G'_v , respectively.

517 In both cases, we have that the multicast channel χ_z was assigned identically by both u and
 518 v . As shown in Case 1, if z transmitted 0 (resp. 1) on χ_z and $\chi_z - F^*$ is non-empty, then
 519 χ_z was assigned to z^0 (resp. z^1) by both u and v , as required. ◀

520 ► **Lemma 9.** *Consider a phase > 0 of Algorithm 1 wherein $F = F^*$. Let*

$$521 \quad Z := \{u^0 \mid u \in F\} \cup \{w \in V(G) - F^* \mid w \text{ flooded value 0 in step (a) of this phase}\}$$

$$522 \quad N := \{u^1 \mid u \in F\} \cup \{w \in V(G) - F^* \mid w \text{ flooded value 1 in step (a) of this phase}\}.$$

524 For any two non-faulty nodes $u, v \in V(G) - F^*$, we have $Z_u = Z_v$ and $N_u = N_v$ in **step (b)** of this phase.

526 **Proof.** Consider the phase where $F = F^*$ and any two non-faulty nodes $u, v \in V(G) - F^*$.
 527 We show that $Z \subseteq Z_v$ and $N \subseteq N_v$ (resp. $Z \subseteq Z_u$ and $N \subseteq N_u$). Since $Z \cup N = Z_u \cup N_u =$
 528 $Z_v \cup N_v$, it follows that $Z = Z_u = Z_v$ and $N = N_u = N_v$. For a node $w \in F^*$, the two split
 529 nodes w^0 and w^1 are assigned identically by both u and v . So consider an arbitrary node
 530 $w \in V(G) - F^* = (Z \cup N) - \{u^0, u^1 \mid u \in F^*\}$. Recall that we are considering the phase
 531 > 0 of the algorithm where $F = F^*$ is the actual set of faulty nodes. There are two cases to
 532 consider:

533 ■ Case 1: $w \in Z - \{u^0 \mid u \in F^*\}$, i.e., $w \notin F^*$ flooded 0 in **step (a)** of this phase.
 534 Let P_{wv} be the wv -path identified by v in **step (b)**. Note that P_{wv} is contained entirely
 535 in $G - F^*$ so that P_{wv} does not have any faulty nodes. It follows that, in **step (a)**, since
 536 w flooded value 0 so v received value 0 along P_{wv} . Therefore, in **step (b)**, v puts w in
 537 the set Z_v .
 538 ■ Case 2: $w \in N - \{u^1 \mid u \in F^*\}$, i.e., $w \notin F^*$ flooded 1 in **step (a)** of this phase.
 539 Let P_{wv} be the wv -path identified by v in **step (b)**. Note that P_{wv} is contained entirely
 540 in $G - F^*$ so that P_{wv} does not have any faulty nodes. It follows that, in **step (a)**, since
 541 w flooded value 1 so v received value 1 along P_{wv} . Therefore, in **step (b)**, v puts w in
 542 the set N_v .

543 So we have that $Z \subseteq Z_v$ and $N \subseteq N_v$, as required. A symmetric argument gives $Z \subseteq Z_u$
544 and $N \subseteq N_u$. As argued before, this implies that $Z = Z_u = Z_v$ and $N = N_u = N_v$. \blacktriangleleft

545 We are now ready to prove Lemma 6.

546 **Proof of Lemma 6.** Consider the phase where $F = F^*$. Suppose $u, v \in V(G) - F^*$ are any
547 two non-faulty nodes. By Lemma 9, we have $Z = Z_u = Z_v$ and $N = N_u = N_v$, where Z and
548 N are as in the statement of Lemma 9. By Lemma 8, we have $G'_u = G'_v$. Let $G' = G'_u = G'_v$.
549 We use F' to denote the set of nodes in G' corresponding to nodes in F^* in G .

550 We now show that all non-faulty nodes in $V(G) - F^*$ have identical state at the end of
551 this phase. Consider **step (c)** of this phase. If either $Z - F^*$ or $N - F^*$ is empty, then
552 all non-faulty nodes have identical state at the start of the phase and they do not update
553 their state in **step (c)**. So suppose that both $Z - F^*$ and $N - F^*$ are non-empty. Observe
554 that, at the start of **step (c)**, all nodes in $Z - F^*$ have identical state of 0, while all nodes
555 in $N - F^*$ have identical state of 1. We show that in **step (c)** either all nodes in $Z - F^*$
556 update their state to 1, or all nodes in $N - F^*$ update their state to 0.

557 Note that $G' \in \Lambda_{F^*}(G)$. By condition AB, either $Z \rightsquigarrow_{G'} N - F'$ or $N \rightsquigarrow_{G'} Z - F'$. We
558 consider each case as follows.

559 ■ Case 1: $Z \rightsquigarrow_{G'} N - F'$.

560 Consider an arbitrary node $v \in (Z \cup N) - F'$. In **step (c)**, v sets $A_v = Z$ and $B_v = N$.
561 If $v \in A_v - F' = Z - F'$, then v has state 0 at the start of this phase and does not update
562 it in **step (c)**. So suppose that $v \in B_v - F' = N - F'$. Now, if in **step (a)** v received
563 the value 0 identically along some $f + 1$ node-disjoint Zv -paths in $G' - (N \cap F')$, then
564 v sets $\gamma_v = 0$ in **step (c)**. We show that such $f + 1$ node-disjoint Zv -paths do indeed
565 exist. Since $Z \rightsquigarrow_{G'} N - F'$, so there exist $f + 1$ node-disjoint Zv -paths in $G' - (N \cap F')$.
566 Without loss of generality, only the source nodes on these paths are from Z . For each
567 such path, observe that only the source node, say $z \in Z$, can be faulty. If the source
568 node z is faulty, then by Lemma 8, and construction of G' and Z , z sent the value 0 on
569 the first channel on this path in **step (a)**. If z is non-faulty, then by construction of
570 Z , z flooded value 0 in **step (a)**. Now all other nodes on the path are non-faulty, so v
571 received value 0 along this path in **step (a)**. Therefore, v received value 0 identically
572 along the $f + 1$ node-disjoint Zv -paths in **step (a)**, as required.

573 ■ Case 2: $Z \not\rightsquigarrow_{G'} N - F'$ so that $N \rightsquigarrow_{G'} Z - F'$ by condition AB.

574 Consider an arbitrary node $v \in (Z \cup N) - F'$. In **step (c)**, v sets $A_v = N$ and $B_v = Z$.
575 If $v \in A_v - F' = N - F'$, then v has state 1 at the start of this phase and does not
576 update it in **step (c)**. So suppose that $v \in B_v - F' = Z - F'$. As in Case 1, since
577 $N \rightsquigarrow_{G'} Z - F'$, so there exist $f + 1$ node-disjoint Nv -paths in $G' - (Z \cap F')$ such that v
578 received the value 1 identically along these paths in **step (a)**. Therefore, v sets $\gamma_v = 1$
579 in **step (c)**, as required.

580 In both of the cases, all non-faulty nodes have identical state at the end of this phase, as
581 required. \blacktriangleleft

582 Using Lemmas 5 and 6, we can now prove the sufficiency of condition AB. Recall that
583 by Theorem 4, condition AB is equivalent to condition LCR. Thus this shows the reverse
584 direction of Theorem 1.

585 **Proof of Theorem 1 (\Leftarrow direction).** Algorithm 1 satisfies the *termination* condition because
586 it terminates in finite time.

587 In one of the iterations of the main `for` loop, we have $F = F^*$, i.e., F is the actual set of
588 faulty nodes. By Lemma 6, all non-faulty nodes have the same state at the end of this phase.
589 By Lemma 5, these states remain unchanged in any subsequent phases. Therefore, all nodes
590 output an identical state. So the algorithm satisfies the *agreement* condition.

591 At the start of phase 1, the state of each non-faulty node equals its own input. By
592 inductively applying Lemma 5, we have that the state of a non-faulty node always equals
593 the *input* of some non-faulty node, including in the last phase of the algorithm. So the
594 output of each non-faulty node is an input of some non-faulty node, satisfying the *validity*
595 condition. \blacktriangleleft