

Towards a Restrained Use of Non-equivocation for Achieving Iterative Approximate Byzantine Consensus

Chuanyou Li, Michel Hurfin, Yun Wang, Lei Yu

► **To cite this version:**

Chuanyou Li, Michel Hurfin, Yun Wang, Lei Yu. Towards a Restrained Use of Non-equivocation for Achieving Iterative Approximate Byzantine Consensus. 30th IEEE International Parallel and Distributed Processing Symposium (IPDPS), May 2016, Chicago, United States. pp.10, 10.1109/IPDPS.2016.62. hal-01339477

HAL Id: hal-01339477

<https://hal.inria.fr/hal-01339477>

Submitted on 29 Jun 2016

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

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

Towards a Restrained Use of Non-equivocation for Achieving Iterative Approximate Byzantine Consensus

Chuanyou Li
INRIA, Rennes, France
chuanyou.li@gmail.com

Michel Hurfin
INRIA, Rennes, France
michel.hurfin@inria.fr

Yun Wang
Southeast University, China
yunwang@seu.edu.cn

Lei Yu
Wuhan University, China
ly.wd@whu.edu.cn

Abstract—We consider the approximate consensus problem in a partially connected network of n nodes where at most f nodes may suffer from Byzantine faults. We study under which conditions this problem can be solved using an iterative algorithm. A Byzantine node can equivocate: it may provide different values to its neighbors. To restrict the possibilities of equivocation, the 3-partial multicast primitive is considered. When a (correct or faulty) node uses this communication primitive, it provides necessarily the same value to the two identified receivers. Based on this communication primitive, a novel condition called f -resilient is proposed and proved to be necessary and sufficient to solve the approximate Byzantine consensus problem in a synchronous network. This condition takes into account two different communication primitives: unicast and 3-partial multicast. It expresses a trade-off between the two known approaches that make the problem solvable (increasing the number of neighbors or/and increasing the power of the communication primitives). The condition f -resilient does not require to eliminate all the possibilities of equivocation. Furthermore, it can be satisfied when there is just a majority of correct nodes. The relationships between the condition f -resilient and the condition h -disjoint (proposed by Alexander Jaffe et al. in 2012 to solve another problem, namely exact Byzantine consensus) are investigated. Two preliminary conclusions are obtained. When a network does not satisfy h -disjoint, it also does not satisfy f -resilient. But when a network satisfies h -disjoint, f -resilient is not necessarily satisfied. Finally, the condition is extended to cope with asynchronous networks.

Keywords-Approximate Consensus; Byzantine failure; Equivocation; Partial multicast; Iterative algorithm;

I. INTRODUCTION

The concept of approximate consensus was introduced by Dolev et al. [1]. In this problem, each node begins its participation by providing a real initial value. Then all the correct nodes must eventually obtain final values that *i*) are different from each other within a maximum gap denoted ϵ (convergence property) and *ii*) must be in the range of initial values proposed by the correct nodes (validity property).

We consider approximate Byzantine consensus in a partially connected network of n nodes where at most f nodes may suffer from Byzantine faults. We study under which conditions this problem can be solved using an iterative algorithm [2], [3]. An iterative algorithm has two main advantages: it does not rely on message-relay and it does

not assume a global knowledge. In particular, no node has to know the global network topology or the value of n . Based on iterative algorithms, [2] has proposed a necessary and sufficient condition related to the network topology which ensures the convergence property. Informally, the condition proposed in [2] imposes that each node has a sufficient number of neighbors. Moreover, in this solution, the maximal number of faulty nodes in the synchronous network is limited by the bound $n \geq 3f + 1$. This requirement on the proportion of faulty nodes is mainly due to the fact that a Byzantine node can equivocate, *e.g.*, sends different messages to different neighbors. Stronger communication primitives such as the multicast has been considered to guarantee that an identical message is received by all the neighbors of a Byzantine node. The condition on the network topology proposed by [3] relies on the fact that all the possibilities of equivocation are eliminated: the bound becomes $n \geq 2f + 1$ in synchronous networks.

In this paper, in the particular context of the iterative approximate Byzantine consensus problem, we study the interest of restricting the possibilities of equivocation without suppressing all of them. To restrict the power of equivocation, we consider a basic communication primitive called a 3-partial multicast (also named a *hyperedge* in a communication topology). A 3-partial multicast identifies one sender and two receivers. When the three nodes use this primitive, the sender is forced to send identical messages to its two neighbors. The concept of hyperedges is strongly connected to the notion of *uniform hyperedges* [4]. Yet, while a uniform hyperedge of three nodes allows any of the three nodes to act as a sender or a receiver, such a symmetry does not exist in the case of a hyperedge where the unique sender is predefined. Based on partial multicast, we propose a new condition named *f-resilient* which is proved to be necessary and sufficient for reaching iterative approximate Byzantine consensus in a synchronous network. This condition is different from the conditions proposed in [2], [3]. First, rather than eliminating any possibility of equivocation like in [3], the *f-resilient* condition just restricts equivocation and requires only that $n \geq 2f + 1$. Second, the condition takes into account two communication primitives unicast and 3-partial multicast. Thus it allows to

find a tradeoff between the two known approaches that make the problem solvable: increase the number of neighbors or/and increase the power of the communication primitives. This work has a strong connection with the *h-disjoint* condition proposed in [4] to solve another problem namely the exact Byzantine consensus problem. As the *h-disjoint* and the *f-resilient* conditions both refer to partial multicast primitives, we investigate the relationships between them. We show that when a network does not satisfy *h-disjoint*, it also does not satisfy *f-resilient*. But when a network satisfies *h-disjoint*, *f-resilient* is not always satisfied.

This paper is organized as follows. Section II introduces the model. Section III provides a formal definition of the iterative approximate consensus problem. Section IV sketches out some related works. In section V, we define the condition *f-resilient*. Then we prove its necessity and sufficiency. In section VI, we investigate the relationships between *h-disjoint* and *f-resilient*. In section VII, we claim that *f-resilient* can be extended to cope with asynchronous systems.

II. MODEL

We consider a network composed of n nodes whose identities are contained in the set $V = \{p_1, p_2, \dots, p_n\}$. Some nodes (at most f) may suffer from Byzantine faults. These Byzantine nodes may stop their computation, behave arbitrarily, and even collude together. Therefore, V is partitioned into two subsets denoted V_c and V_b . The set V_c contains the correct nodes which always follow the protocol specification. The Byzantine nodes belong to the set V_b whose composition is not known by the correct nodes. By definition, $|V_b| \leq f$.

All the nodes communicate only by exchanging messages. In this study the communication process is assumed to be synchronous and no message is lost or modified during its transfer. We consider two different communication primitives: unicast or/and partial multicast. Unicast is corresponding to a point to point communication. A set E_U contains all the unicast channels represented by ordered pairs. For example, $(p_i; p_j) \in E_U$ implies that during the iterative computation, p_i is expected to send periodically messages to p_j using the unicast primitive. When a Byzantine node uses this communication primitive, it is able to equivocate. It can send simultaneously messages with different (fake or correct) values through different unicast channels. The use of a multicast primitive is intended to reduce the power of some Byzantine nodes. In this work, we consider a particular communication primitive called the 3-partial multicast. A 3-partial multicast channel is composed of a sender and two receivers. In the triplet of three nodes identified by such a multicast channel, the sender is forced to send simultaneously the same message to the two receivers. Consequently, the power of a Byzantine node is restricted when it uses a 3-partial multicast to send a message. A set E_M contains all the triplets $(p_i; p_j, p_k)$ such that, during the

iterative computation, node p_i is expected to send periodically messages to p_j and p_k using the 3-partial multicast primitive. The elements in E_M are also called hyperedges (and not uniform hyperedges as the role of the sender is devoted to just one of the three nodes, namely the node identified on the left side of the semicolon in the triplet). Thus the fact that $(p_i; p_j, p_k) \in E_M$ does not impose (nor prevent) that $(p_j; p_i, p_k)$ or $(p_k; p_i, p_j)$ also belongs to E_M . Receivers associated to a 3-partial multicast are not ordered: the notations $(p_i; p_j, p_k)$ and $(p_i; p_k, p_j)$ represent the same element. By definition, none of the two sets E_U and E_M contains duplicated elements. Like in [5], [6], [4], each correct node knows the (unicast or/and multicast) channels it has joined (either as a sender or a receiver). Even if this particular aspect is not discussed in this paper, the partial multicast model is motivated by various environments [7] (for example, a LAN like an Ethernet bus or a token ring, a group communication primitive, ...).

A static hybrid graph $G = (V, E_U, E_M)$ is used to represent the communication topology. This model is hybrid because it mixes information about the exploited two communication primitives (the set E_U which models the use of the unicast primitive and the set E_M which characterizes the use of the 3-partial multicast primitive). The hybrid graph is static because no mobility is considered and the choice of the communication primitives is supposed to be defined once and prior to the agreement protocol's execution. Based on the communication topology, we now introduce the concept *source neighbor*. Roughly, each node p_i of V is associated with a set denoted N_i of *source neighbors* from which p_i can receive messages. By definition, $p_j \in N_i$ if and only if either $(p_j; p_i) \in E_U$ or $(p_j; p_i, p_k) \in E_M$ (k identifies another node p_k that is different from p_i and p_j). Note that for p_i and one of its source neighbor p_j , the set E_M may contain several multicast channels associated to them. More precisely, the number of multicast channel $(p_j; p_i, p_k)$ associated to p_j and p_i is comprised between 0 and $n - 2$. For example, it could be the case both $(p_j; p_i, p_{k_1})$ and $(p_j; p_i, p_{k_2})$ (with $k_1 \neq k_2$) belong to E_M . The network is not required to be fully connected: N_i can be a proper subset of $V - \{p_i\}$. If some source neighbors of a correct node p_i are Byzantine nodes, we assume that their power is limited: p_i is able to identify the real sender of any received message. Thus a Byzantine node cannot use a fake identity to pretend to be another node.

III. THE APPROXIMATE CONSENSUS PROBLEM

In the approximate consensus problem, each correct node has an initial real value which will be continuously updated during the computation. An iterative solution that solves this problem assumes that each node executes a sequence of *rounds*. Each round is identified by a round number r that belongs to the set $\mathfrak{R} = \{1, 2, \dots\}$. During a round, each correct node performs three operations: *sending*, *receiving*,

and *updating*. More precisely, (1) a node sends its current value to its neighbors (using either unicast or multicast primitives); (2) it receives values from different source neighbors; (3) it updates its local value by using a function that takes into account its current value and some values received during the round. The value of a correct node p_i when it begins the execution of round r is denoted as $v_i(r)$. As a Byzantine node may propose different values during the same round, the notation $v_i(r)$ is meaningless if p_i is not correct. By definition, the initial value of a correct node p_i is denoted as $v_i(1)$. When round r begins, among the set of $|V_c|$ values owned by the correct nodes, the minimum (respectively the maximum) value is denoted as $\min(V_c, r)$ (resp. $\max(V_c, r)$). More generally, at the beginning of a round r , if S is a nonempty subset of V_c , the notation $\min(S, r)$ (resp. $\max(S, r)$) represents the minimum (resp. the maximum) value among the identified set of $|S|$ values.

Definition 1. The approximate Byzantine consensus problem is formally defined by two properties:

Validity: A correct node p_i has always a valid value:

$$\forall r \in \mathfrak{R}, \forall p_i \in V_c, v_i(r) \in [\min(V_c, 1), \max(V_c, 1)].$$

Convergence: Eventually, all the correct nodes have values which are different from each other within a maximum predefined value ϵ (with $\epsilon > 0$):

$$\exists \tilde{r} \in \mathfrak{R} \text{ such that } \forall r > \tilde{r}, \max(V_c, r) - \min(V_c, r) < \epsilon.$$

IV. RELATED WORKS

Dolev et al. propose the earliest results on approximate Byzantine consensus [1]. When the network is fully connected and the total number of nodes is known, two algorithms [1] are proved correct respectively in a synchronous and an asynchronous network. [8], [9] investigate approximate Byzantine consensus in partially connected networks. However without message flooding, convergence cannot be reached. Based on the knowledge of the global topology, [10], [11] address the same problem and succeed to achieve the convergence. Without flooding and the knowledge of the global topology, [2] provides a solution based on an iterative algorithm. To ensure convergence, [2] proposes a sufficient and necessary condition to restrict the network topology. [12], [13] address the problem in a mobile environment and provide a sufficient and necessary condition which restricts the nodes' trajectories.

The strategy of some works [14], [3] consists in eliminating any possibility of equivocation. [14] improves the solution proposed in [1]: only $4f + 1$ nodes are needed in a fully connected asynchronous network. [3] considers iterative algorithms and propose a condition to ensure final convergence in partially connected networks. It requires $n \geq 2f + 1$ in synchronous networks. Rather than eliminating any possibility of equivocation, our condition only restricts them and requires $n \geq 2f + 1$ in synchronous networks.

Some works try to achieve similar goals but consider a different problem. In the case of the exact Byzantine consensus

problem, [15], [16], [5], [6], [4] have studied the power of non-equivocation. By using transferable authentication and non-equivocation, [15] succeeds to transform any protocol that works under the crash model into a protocol that tolerates Byzantine failures without increasing the number of nodes. Based on the same methodology, in [16], the resiliency bound for asynchronous multiparty computation is improved to $n \geq 2f + 1$. Recently, uniform hyperedges [5], [6], [4] have been considered to restrict equivocation. [5] proposes a Byzantine consensus algorithm by considering there is an uniform hyperedge among every three nodes. However, [6], [4] point out that having uniform hyperedges everywhere is not necessary. To solve the exact Byzantine consensus problem, [4] proposes a sufficient and necessary condition on the network called h -disjoint. As this work is strongly related to our study (even if the targeted problems are different), we investigate the relationships between the conditions h -disjoint and f -resilient in section VI.

V. A NECESSARY AND SUFFICIENT CONDITION

A. The proposed condition

To define formally our f -resilient condition, we first introduce the notion of *safe F partition* (similar to [2]). Recall that $\langle S_1, S_2, \dots, S_m \rangle$ is a partition of a set S , if and only if, $S_i \cap S_j = \emptyset$ ($i \neq j$) and $\bigcup_{1 \leq i \leq m} S_i = S$.

Definition 2. Let $G = (V, E_U, E_M)$ be an hybrid-graph, $\langle F, L, M, R \rangle$ is a F partition of G if and only if it is a partition of V , $|F| \leq f$, $|L| > 0$, $|M| \geq 0$ and $|R| > 0$.

Let $\langle F, L, M, R \rangle$ be a F partition of G . We use the notation \mathbb{L} to represent the set $L \cup M$, while \mathbb{R} corresponds to $R \cup M$. To determine if a F partition is safe or not, we focus on all the pairs of nodes (p_i, p_j) such that $p_i \in L$ and $p_j \in R$. For each pair, we consider the source neighbors of p_i which belong either to M or R (i.e., $N_i \cap \mathbb{R}$) and the source neighbors of p_j which belong either to L or M (i.e., $N_j \cap \mathbb{L}$). We also consider all the hyperedges of E_M where p_i and p_j act as the two receivers while the sender is a node p_x belonging to F . A subset F_{ij} of F is defined as follows: $F_{ij} = \{p_x | p_x \in F \wedge (p_x, p_i, p_j) \in E_M\}$.

Definition 3. $\langle F, L, M, R \rangle$ is a *safe F partition of G* if and only if one of the two following properties holds.

C_1 : $\exists p_i \in L$ such that $|N_i \cap \mathbb{R}| \geq f + 1$ or $\exists p_j \in R$ such that $|N_j \cap \mathbb{L}| \geq f + 1$;

C_2 : $\exists p_i \in L, \exists p_j \in R$, such that $1 \leq |N_i \cap \mathbb{R}| \leq f$, $1 \leq |N_j \cap \mathbb{L}| \leq f$, $|F_{ij}| + |N_i \cap \mathbb{R}| + |N_j \cap \mathbb{L}| \geq 2f + 1$.

Examples of safe F partition are depicted in Figure 1. In Definition 3, the two properties C_1 and C_2 are exclusive. Both can be false but when one is true, the other one is false. Property C_1 expresses a requirement on the communication topology. At least one node in L (or in R) must have a sufficient number of source neighbors in \mathbb{R} (or in \mathbb{L}). Without the property C_2 , our condition is exactly the one

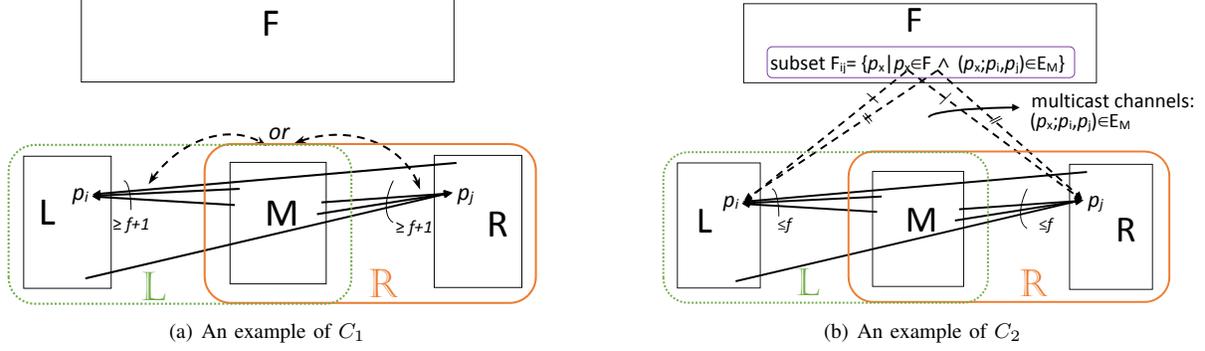


Figure 1. Examples of *safe F* partitions

Table I
EXAMPLE OF f -RESILIENT, $n = 5$, $f = 2$

sender	p_a	p_b	p_c	p_d	p_e
multicast channels	$(p_a; p_b, p_d)$	$(p_b; p_c, p_e)$	$(p_c; p_d, p_a)$	$(p_d; p_e, p_b)$	$(p_e; p_b, p_d)$
	$(p_a; p_b, p_c)$	$(p_b; p_c, p_d)$	$(p_c; p_d, p_e)$	$(p_d; p_e, p_a)$	$(p_e; p_b, p_c)$
	$(p_a; p_b, p_e)$	$(p_b; p_d, p_a)$	$(p_c; p_b, p_e)$	$(p_d; p_a, p_c)$	$(p_e; p_a, p_c)$
	$(p_a; p_c, p_e)$	$(p_b; p_d, p_e)$	$(p_c; p_e, p_a)$	$(p_d; p_c, p_b)$	$(p_e; p_c, p_d)$
	$(p_a; p_c, p_d)$	$(p_b; p_a, p_c)$	$(p_c; p_d, p_b)$	$(p_d; p_e, p_c)$	$(p_e; p_d, p_a)$

proposed in [2]. Thus, property C_2 extends the work of [2] when the topological requirement is not satisfied. In fact, property C_2 captures the interest of using a multicast primitive to prevent equivocation. The fact that the number of source neighbors is not sufficient is counterbalanced by the existence of hyperedges that link a sender p_x of F to a receiver p_i of L and a receiver p_j of R . The number of such hyperedges, namely $|F_{ij}|$, the number of source neighbors of p_i in \mathbb{R} , namely $|N_i \cap \mathbb{R}|$, and the number of source neighbors of p_j in \mathbb{L} , namely $|N_j \cap \mathbb{L}|$ are added. The sum has to be greater than or equal to $2f + 1$.

Definition 4. A hybrid-graph G is named **f -resilient** if and only if all its F partitions are safe.

The two properties that characterize a safe F partition allow to find a tradeoff between two approaches: increasing the connectivity (more source neighbors) or/and increasing the use of powerful communication primitives (more hyperedges). Through an example (See Table I), we show that an hybrid graph G can be f -resilient even if all the possibilities of equivocation are not eliminated.

In this example, G is composed of five nodes p_a, p_b, p_c, p_d, p_e among which at most two can be Byzantine nodes ($f = 2$). Any F partition such that $|F| = 2$ and $|M| = 0$ can not satisfy the first property that characterizes a safe partition: neither L nor R can contain more than two elements and thus the limit $f + 1 = 3$ imposed by the first property cannot be reached. The use of hyperedges is mandatory. Table I identifies a solution where 25 hyperedges are used (among the 30 which could be defined). In this

solution, G is f -resilient even if equivocation is restricted but not completely eradicated: if p_a is a Byzantine node, it can send different messages to p_d and p_e . This example also illustrates that the f -resilient condition requires $n \geq 2f + 1$. Obviously, when $n = 2f = 4$, any F partition with $|F| = 2, |L| = 1, |M| = 0$ and $|R| = 1$ can not be safe due to the fact that $|F| + |L| + |R| < 2f + 1$.

As the condition f -resilient requires that all F partitions are safe, an interesting future work consists in investigating the relationships between f -resilient and the connectivity of the directed communication topology. Herein we only point out that when $n = 2f + 1$, the condition implies that the communication topology is fully connected. However, when $n > 2f + 1$, this is not necessarily required.

B. Necessity

In this section, we prove the necessity of f -resilient.

Lemma 1. No iterative algorithm solves the approximate Byzantine consensus, if G is not f -resilient.

Proof: Assume that an iterative algorithm (executed endlessly by every correct node) is able to solve the approximate Byzantine consensus problem. The validity and convergence properties that characterize this problem have to be satisfied by this algorithm in any scenarios. To show that this assumption leads to a contradiction, we will identify a particular scenario denoted S in which, after a round r , at least two correct nodes (denoted p_x and p_y) definitely stop the convergence process because they have to ensure the validity property is never violated. More precisely, during

any round r' such that $r' \geq r$, p_x (resp. p_y) does not change its value $v_x = \min(V_c, r)$ (resp. $v_y = \max(V_c, r)$). Consequently the difference $v_y - v_x$ remains equal to a value greater than ϵ . In the scenario S , the node p_k (with $k = x$ or $k = y$) can receive at most $2f$ values during a round. Among these values, at most f values (that have been provided by nodes from a set A_k) are smaller or equal to its own value v_k and at most f values (that have been provided by nodes from a set B_k) are greater or equal to v_k . As $A_k \cap B_k = \emptyset$, the proof relies on the fact that the correct node p_k can not know if a node is correct or not. It cannot distinguish between two extreme situations: 1) $A_k \subseteq V_c$ and $B_k \subseteq V_b$, 2) $A_k \subseteq V_b$ and $B_k \subseteq V_c$. If the node p_k computes a new value v during a round $r' \geq r$, this value v must satisfy the property $v \leq v_k$ (because the first possible situation can be the actual scenario) and the property $v \geq v_k$ (because the second possible situation can be the actual scenario). Thus, to ensure the validity property, the algorithm must leave the values of p_x and p_y unchanged after round r . Consequently it violates the convergence property.

Above, we have depicted the general structure of the proof. Now, we demonstrate the existence of the nodes p_x and p_y and we identify a scenario S for which four sets A_x , B_x , A_y and B_y can be defined with respect to p_x and p_y . Of course, the following requirements have to be satisfied: $A_x \cap B_x = \emptyset$, $A_y \cap B_y = \emptyset$, the cardinalities of these four sets are less or equal to f and the values provided by the nodes of A_x (resp. A_y) are smaller or equal to v_x (resp. v_y) while the values provided by the nodes of B_x (resp. B_y) are greater or equal to v_x (resp. v_y).

As the hybrid-graph G is not f -resilient, there exists at least one unsafe F partition denoted $\langle F', L', M', R' \rangle$ such that: $\forall p_i \in L', \forall p_j \in R', |N_i \cap \mathbb{R}'| \leq f, |N_j \cap \mathbb{L}'| \leq f$ and $|F'_{ij}| + |N_i \cap \mathbb{R}'| + |N_j \cap \mathbb{L}'| \leq 2f$. Furthermore, we have $|F'| \leq f, |L'| > 0$ and $|R'| > 0$. The scenario S is based on this unsafe partition. We assume that the set F' is corresponding to the set of Byzantine nodes V_b . Consequently, all the nodes of $L' \cup M' \cup R'$ are correct. At the beginning of a round r (possibly the first round), we assume also that all the nodes of L' have the minimum value $\min(V_c, r)$ and all the nodes of R' have the maximum value $\max(V_c, r)$. Note that the remaining correct nodes of M' (if any) have their values within the range $[\min(V_c, r), \max(V_c, r)]$. Thereafter we consider that p_x is any node of L' while p_y is any node of R' . As $|L'| > 0$ and $|R'| > 0$, two such nodes always exist and, by definition, $v_x = \min(V_c, r)$ while $v_y = \max(V_c, r)$. Regarding the behaviors of the Byzantine nodes in scenario S , we consider that the set F' is partitioned into three subsets denoted F^1, F^2 , and F^3 . During each round, Byzantine nodes of $F^1 \cup F^2 = F'_{xy}$ use the multicast primitive and always provide the same value to p_x and p_y while Byzantine nodes of $F^3 = F'/F'_{xy}$ equivocate and provide different values. The nodes of F^1 always provide a value belonging to the range $(-\infty, v_x)$.

The nodes of F^2 always provide a value belonging to the range $(v_y, +\infty)$. The nodes of F^3 always provide p_x a value belonging to the range $(-\infty, v_x)$ and always provide p_y a value belonging to the range $(v_y, +\infty)$. The scenario S is unique but nevertheless the cardinality of F^2 depends on a characteristic of the unsafe partition: if $|F'_{xy}| + |N_y \cap \mathbb{L}'| \leq f$ then $|F^2| = 0$ else $|F^2| = |F'_{xy}| + |N_y \cap \mathbb{L}'| - f$. Consequently, we have always $|F^1| + |N_y \cap \mathbb{L}'| \leq f$. Furthermore, as $|F'_{xy}| + |N_x \cap \mathbb{R}'| + |N_y \cap \mathbb{L}'| \leq 2f$, we have also $|F^2| + |N_x \cap \mathbb{R}'| \leq f$.

Let $A_x = F^1 \cup F^3, B_x = F^2 + |N_x \cap \mathbb{R}'|, B_y = F^2 \cup F^3$ and $A_y = F^1 + |N_y \cap \mathbb{L}'|$. These sets satisfy all the specified requirements and thus S is a particular scenario where the convergence property cannot not be ensured (without the risk of violating the validity property). Note that during the computation, nodes of M' can change their values but this will have no consequence for p_x and p_y . ■

C. Sufficiency

To prove that the condition f -resilient is also sufficient, we propose and describe an iterative algorithm called LIABC (for Linear Iterative Approximate Byzantine Consensus) that satisfies the validity property. Then, assuming that the f -resilient condition is satisfied, we prove that the LIABC algorithm ensures the convergence property.

Algorithm 1 Iteration r of the LIABC Algorithm

- 1: p_i sends $\langle p_i, v_i(r), r \rangle$ through all the unicast & multicast channels for which p_i acts as a sender;
 - 2: $Neb_i(r) \leftarrow$ all the $\langle p_j, v_j(r), r \rangle$ values received from an unicast or multicast channel for which p_i is a receiver;
 - 3: $Neb_i(r) \leftarrow DetectReplace(Neb_i(r))$;
 - 4: $Neb_i(r) \leftarrow RemoveDuplicates(Neb_i(r))$;
 - 5: $SNeb_i(r) \leftarrow Reduce(Neb_i(r), f)$;
 - 6: $v_i(r) \leftarrow \frac{v_i(r) + \sum_j v_j(r)}{|SNeb_i(r)| + 1}$; ($v_j \in SNeb_i(r)$)
 - 7: $r \leftarrow r + 1$;
-

1) *The LIABC algorithm and its validity:* The pseudo code in Figure 1 describes the behavior of a correct node p_i during an iteration r . During the initialization phase (not described in Figure1), r is set to the value 1 and the initial value of p_i is stored in the variable $v_i(1)$. During the sending step of an iteration r (line 1), for each channel in which p_i acts as a sender (*i.e.* unicast channels $(p_i; ?)$ and multicast channels $(p_i; ?, ?)$) p_i sends a message with its identity p_i , its current value $v_i(r)$ and the round number r . During the receiving step (line 2), p_i waits to receive messages coming from channels for which it acts as a receiver (*i.e.* unicast channels $(?; p_i)$ and multicast channels $(?; p_i, ?)$ or $(?; ?, p_i)$). The set $Neb_i(r)$ contains all the received messages. During the updating step (lines 3-6), the node p_i first calls the function *DetectReplace* to identify some Byzantine nodes among its neighbors. If a node p_b shares

a communication channel with p_i and has not sent the expected message, p_b is necessarily a Byzantine node. Indeed, the network is synchronous and no channel drops messages. Similarly, if p_i shares several channels with a same source neighbor p_b but receives different values from p_b during the same round, p_i can also conclude that p_b is a Byzantine node. Of course, only unexpected behaviors are detected and thus some Byzantine nodes remains undetected: the failure detection mechanism satisfies a strong accuracy property (no correct node is ever suspected) but no completeness property. If a Byzantine node p_b is detected during round r , all the information it has provided is removed from the set $Neb_i(r)$. Finally, for each detected node p_b , the function *DetectReplace* insert a single information, namely $\langle p_b, \perp, r \rangle$ in $Neb_i(r)$. The value \perp is such that any value provided by a correct node is strictly greater than \perp . Note that the code can be optimized to ignore the values provided by detected nodes during the remainder of computation and not only during the current iteration. Herein, we ignore this possibility. As a node (even a correct one) can send several messages to p_i (through different channels), the set $Neb_i(r)$ can contain duplicated values. The call to the function *RemoveDuplicates* (line 5) ensures that $Neb_i(r)$ contains only one value from each source neighbor. Finally, the call to the function *Reduce* aims at suppressing some values to keep only those that are not risky (*i.e.* using the remaining values during the computation of the new value of p_i cannot compromise the validity property). If $Neb_i(r)$ contains less than f values greater than $v_i(r)$, p_i suppresses all these values. Otherwise, p_i suppresses the f largest values of $Neb_i(r)$. Likewise, if $Neb_i(r)$ contains less than f values smaller than $v_i(r)$, p_i suppresses all these values. Otherwise, p_i suppresses the f smallest values. Note that at most $2f$ values are suppressed. Obviously, the \perp values corresponding to detected Byzantine nodes are suppressed at this stage. After reducing, p_i calculates the average between its current value $v_i(r)$ and the remaining values of $SNeb_i(r)$ (line 6). The round number r is increased by 1 at the end.

Theorem 1. *The algorithm LIABC satisfies validity.*

Of course, the validity property holds after the initialization phase. Let us consider that this is no more true when round $r > 1$ begins. It means that, during round $r - 1$, at least one correct node p_x has used a value beyond the range $[\min(V_c, 1), \max(V_c, 1)]$ to update its own value. The function *Reduce* ensures that this scenario is impossible: if a value v proposed by a undetected Byzantine node still remains in $SNeb_x(r - 1)$, then at least one correct neighbor of p_x (or p_x itself) has provided a value v_a and at least one correct neighbor of p_x (or p_x itself) has provided a value v_b such that $v_a \leq v \leq v_b$. Thus p_x can not compute an average (line 6) with a value beyond the range $[\min(V_c, r - 1), \max(V_c, r - 1)]$. Indeed the algorithm LIABC ensures a stronger property (Theorem 2).

Theorem 2. *During the sequence of rounds $1 \dots r \dots$, the minimum value $\min(V_c, r)$ is non-decreasing and the maximum value $\max(V_c, r)$ is non-increasing.*

2) *The convergence:* Definition 3 identifies two distinct safety properties (denoted C_1 and C_2). Therefore, when an hybrid graph G is f -resilient, two cases can be distinguished. First, all the F partitions are safe because each of them satisfies the property C_1 . Second, some (at least one) F partitions are safe because they satisfy the property C_2 . In [2], the authors have already consider the first case. So, we only focus on the proof related to the second case. Any F partition mentioned below is supposed to be *safe*. A particular set of partitions Γ is defined.

Definition 5. $\Gamma = \{F = \langle F, L, M, R \rangle \mid F = V_b\}$

Lemma 2. *Let G be a f -resilient hybrid graph and let F be a partition of Γ such that, during a round r , $\max(L, r) < \min(\mathbb{R}, r)$ and $\max(\mathbb{L}, r) < \min(R, r)$. Whatever the behaviors of the Byzantine nodes of set F , at least one of the two following properties holds:*

- 1) $\exists p_i \in L, \exists v \in SNeb_i(r)$ s.t. $v \in [\min(\mathbb{R}, r), \max(\mathbb{R}, r)]$
- 2) $\exists p_j \in R, \exists v \in SNeb_j(r)$ s.t. $v \in [\min(\mathbb{L}, r), \max(\mathbb{L}, r)]$

Proof: The proof is by contradiction. Consider a particular F partition $\langle F, L, M, R \rangle$ of the set Γ such that $\max(L, r) < \min(\mathbb{R}, r)$ and $\max(\mathbb{L}, r) < \min(R, r)$. Assume that this partition satisfies none of the dual properties.

As G is f -resilient, the partition F is safe. Yet the property C_1 cannot be satisfied by this partition: $\forall p_i \in L$ and $\forall p_j \in R$, neither $|N_i \cap \mathbb{R}| \geq f + 1$ nor $|N_j \cap \mathbb{L}| \geq f + 1$ is possible. As the proofs are symmetric, we just show why $|N_i \cap \mathbb{R}| \geq f + 1$ is impossible. As $\max(L, r) < \min(\mathbb{R}, r)$, if $|N_i \cap \mathbb{R}| \geq f + 1$, whatever the values received from the Byzantine nodes, p_i gathers at least $f + 1$ values from correct nodes of \mathbb{R} . Thus after reducing (line 5), $SNeb_i(r)$ contains at least a value v within the range $[\min(\mathbb{R}, r), \max(\mathbb{R}, r)]$.

Therefore, the property C_2 is necessarily satisfied by F : $\exists p_i \in L, \exists p_j \in R$, such that $1 \leq |N_i \cap \mathbb{R}| \leq f$, $1 \leq |N_j \cap \mathbb{L}| \leq f$ and $|F_{ij}| + |N_i \cap \mathbb{R}| + |N_j \cap \mathbb{L}| \geq 2f + 1$. The $|N_i \cap \mathbb{R}|$ correct nodes that belong to \mathbb{R} provide to p_i a value v such that $v \geq \min(\mathbb{R}, r)$. The node p_i can also receive values from the Byzantine nodes of F . Let us define two subsets of F denoted F^1 and F^2 such that F^1 contains all the nodes of F that send to p_i a value v^1 such that $v^1 \geq \min(\mathbb{R}, r)$ and F^2 contains all the nodes of F that send to p_i a value v^2 such that $v^2 < \min(\mathbb{R}, r)$. We must have $|F^1| + |N_i \cap \mathbb{R}| \leq f$. Otherwise, after the reducing operation, a value within the range $[\min(\mathbb{R}, r), \max(\mathbb{R}, r)]$ will still remain in the set $SNeb_i(r)$ of p_i .

Consider now the node p_j . This node receives from $|N_j \cap \mathbb{L}|$ correct nodes of \mathbb{L} a value v such that $v \leq \max(\mathbb{L}, r)$. As we have $|N_j \cap \mathbb{L}| \leq f$ and $|F_{ij}| + |N_i \cap \mathbb{R}| + |N_j \cap \mathbb{L}| \geq 2f + 1$, we can conclude that $|F_{ij}| + |N_i \cap \mathbb{R}| \geq f + 1$. We have previously shown that $|F^1| + |N_i \cap \mathbb{R}| \leq f$. Consequently,

$|F_{ij}| > |F^1|$. Byzantine nodes of F_{ij} can either provide to p_j the value they also provide to p_i or provide no value to p_j . In the latter case, the node is recognized as been a Byzantine node by p_j which will add a \perp value (such that $\perp < \min(\mathbb{L}, r) < \min(\mathbb{R}, r)$) to its set $N_{eb_j}(r)$ (line 3). Consequently, among the nodes of F^2 , at least $|F_{ij}| - |F^1|$ Byzantine nodes also send to p_j a value v^2 such that $v^2 < \min(\mathbb{R}, r)$. If M is empty, any value v^2 such that $v^2 < \min(\mathbb{R}, r)$ satisfies also $v^2 \leq (\mathbb{L}, r)$. Otherwise, when the set M is not empty, we have $\min(\mathbb{R}, r) \leq \max(\mathbb{L}, r)$. Consequently, p_j can receive $|F_{ij}| - |F^1| + |N_j \cap \mathbb{L}|$ values that are less or equal to $\max(\mathbb{L}, r)$. Again, as we have previously shown that $|F^1| + |N_i \cap \mathbb{R}| \leq f$, the inequality $|F_{ij}| + |N_i \cap \mathbb{R}| + |N_j \cap \mathbb{L}| \geq 2f + 1$ leads us to conclude that $|F_{ij}| - |F^1| + |N_j \cap \mathbb{L}| \geq f + 1$. After the call to *Reduce*, the f minimum values are suppressed. But at least one value v within the range $v \in [\min(\mathbb{L}, r), \max(\mathbb{L}, r)]$ still remains. This contradicts our assumption. ■

For each round r , we define now two particular sets of nodes as follows. The set $S_{\min}(r)$ includes all the correct nodes that have the minimum value $\min(V_c, r)$ at the beginning of round r while the set $S_{\max}(r)$ includes all the correct nodes that have the maximum value $\max(V_c, r)$.

Corollary 1. *Let G be a f -resilient hybrid graph. For any round r , at least one of the two following properties holds:*

- 1) $\exists p_i \in S_{\min}(r)$, such that $v_i(r+1) > \min(V_c, r)$
- 2) $\exists p_j \in S_{\max}(r)$, such that $v_j(r+1) < \max(V_c, r)$

Proof: During round r , consider the partition $\langle F, L, M, R \rangle$ of Γ that satisfies $L = S_{\min}(r)$ and $R = S_{\max}(r)$. The conditions expressed in Lemma 2 hold and consequently either a node p_i of $S_{\min}(r)$ computes (at line 6) an average with a value greater than $\min(V_c, r)$ or a node p_j of $S_{\max}(r)$ computes an average with a value less than $\max(V_c, r)$. Thus, in round $r+1$, either $v_i(r+1) > \min(V_c, r)$ or $v_j(r+1) < \max(V_c, r)$. ■

Corollary 2. $\forall r \in \mathfrak{R}$, there always exists a round r^* with $r^* > r$, such that $\forall r' \geq r^*$, $\max(V_c, r') - \min(V_c, r') < \max(V_c, r) - \min(V_c, r)$.

Proof: Theorem 2 guarantees $\min(V_c, r)$ is non-decreasing and $\max(V_c, r)$ is non-increasing. Remember that the *Reduce* function ensures that no correct node computes a new value with a value outside the range $[\min(V_c, r), \max(V_c, r)]$. Therefore, if a correct node holds a value within the range $(\min(V_c, r), \max(V_c, r))$ in round r , by using average calculation (line 6), its value cannot reach the bounds $\min(V_c, r)$ or $\max(V_c, r)$ in the future. It means that both $|S_{\min}(r)|$ and $|S_{\max}(r)|$ are non-increasing. Corollary 1 ensures that in round r , either a correct node that has the minimum value $\min(V_c, r)$ will increase its value or a correct node that has the maximum value will decrease its value. As both $|S_{\min}(r)|$ and $|S_{\max}(r)|$ are limited, there must exist a round r^* , such that $\forall r' \geq r^*$ either

$\min(V_c, r') > \min(V_c, r)$ or $\max(V_c, r') < \max(V_c, r)$. ■

Lemma 3. *The algorithm LIABC can solve approximate Byzantine consensus in G if G is f -resilient.*

Corollary 2 ensures the distance $\max(V_c, r) - \min(V_c, r)$ decreases continuously. Furthermore, due to Theorem 2, any decrease of this distance is irreversible. To prove lemma 3, we have still to demonstrate that, when G is f -resilient, for any given $\epsilon > 0$, there does not exist a value $\mu \geq \epsilon$, such that $\lim_{r \rightarrow +\infty} \max(V_c, r) - \min(V_c, r) = \mu$. We prove this by contradiction. First we define Hypothesis 1 and then we prove it is not true.

Hypothesis 1. *Given a value $\epsilon > 0$, there exists a value $\mu \geq \epsilon$, such that $\lim_{r \rightarrow +\infty} \max(V_c, r) - \min(V_c, r) = \mu$.*

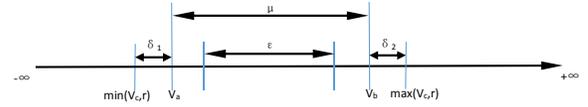


Figure 2. The example of ϵ and μ

To facilitate the proof, we rewrite Hypothesis 1 in an equivalent way. There exist two values denoted v_a and v_b with $v_b - v_a = \mu$ (see Figure 2). For any round r , we suppose that $\min(V_c, r) = v_a - \delta_1(r)$ and $\max(V_c, r) = v_b + \delta_2(r)$. The dual parameters $\delta_1(r)$ and $\delta_2(r)$ satisfy $\delta_1(r) \geq 0$ and $\delta_2(r) \geq 0$ but at least one of them is strictly greater than 0 (no violation of Corollary 2). If $\delta_1(r) > 0$ (or/and $\delta_2(r) > 0$), as r increases $\delta_1(r)$ (or/and $\delta_2(r)$) is monotone non-increasing and converges to zero. Based on these new notations, we define two sets related to r : $S_{n\min}(r) = \{p_i | v_i(r) \in [v_a - \delta_1(r), v_a] \wedge p_i \in V_c\}$ and $S_{n\max}(r) = \{p_i | v_i(r) \in [v_b, v_b + \delta_2(r)] \wedge p_i \in V_c\}$. $S_{n\min}(r)$ includes the correct nodes that have a value less than or equal to v_a in round r , while $S_{n\max}(r)$ includes the correct nodes that have a value greater than or equal to v_b in round r . Note that by assumption $S_{n\min}(r) \neq \emptyset$ and $S_{n\max}(r) \neq \emptyset$ in any round r .

Lemma 4. *Suppose Hypothesis 1 is true. If $\delta_1(r) > 0$ in any round r , the set V_c has a proper subset $V'_c \neq \emptyset$, such that $\forall p_i \in V'_c$ its value v_i converges to v_a . Symmetrically, if $\delta_2(r) > 0$ in any round r , the set V_c has a proper subset $V''_c \neq \emptyset$, such that $\forall p_j \in V''_c$ its value v_j converges to v_b .*

Proof: As the proofs for V'_c and V''_c are similar, we only consider the case of V'_c . First, if V'_c exists (nonempty), it is a proper subset of V_c . Otherwise, it means that $S_{n\max}$ will become empty during a given round and thus Hypothesis 1 is not true. The proof focuses now on the correct nodes p_i that stay in $S_{n\min}$ forever or that leave and integrate again this set infinitely often. These nodes are elements of a set \tilde{V}_c . More formally, $\tilde{V}_c \subseteq V_c$ and $\forall p_i \in \tilde{V}_c$ there exists an infinite subset $\mathfrak{R}_i \subset \mathfrak{R}$, such that $\forall r \in \mathfrak{R}_i$, $p_i \in S_{n\min}(r)$ and $\forall r \notin \mathfrak{R}_i$

\mathfrak{R}_i , $p_i \notin S_{nmin}(r)$. By Hypothesis 1, we know $S_{nmin}(r)$ is never empty. Consequently, \tilde{V}_c is also non empty. If we prove that the value $v_i(r)$ of each node p_i of \tilde{V}_c converges to v_a , we can conclude the existence of V'_c .

For any round r of \mathfrak{R} , we start by estimating an upper bound for $v_i(r)$ when $p_i \in \tilde{V}_c$. There are two cases. First, if $r \in \mathfrak{R}_i$, we have immediately $v_i(r) \leq v_a$. Second, when $r \notin \mathfrak{R}_i$, we have $v_i(r) > v_a$ but we know that after a finite number of rounds we will have again a round r' such that $r' \in \mathfrak{R}_i$ (and so $v_i(r') \leq v_a$). There exists an upper bound θ on the difference $r' - r \leq \theta < +\infty$ (θ can be any large integer but not infinite). We compute now an over estimation of the maximum decrease of the value of p_i (during θ consecutive rounds). In an extreme scenario, from round r till round $r' - 1$, p_i always uses its own value and $n - 1$ values $v_a - \delta_1(r)$ when it computes an average (line 6). Consequently, when the round r' begins, we have $\frac{v_i(r) + (n^\theta - 1)(v_a - \delta_1(r))}{n^\theta} \leq v_i(r') \leq v_a$. From this formula, we extract an upper bound for $v_i(r)$ which is equal to $v_a + (n^\theta - 1)\delta_1(r)$. Thus we have $v_i(r) \leq v_a$ if $r \in \mathfrak{R}_i$ and $v_i(r) \leq v_a + (n^\theta - 1)\delta_1(r)$ if $r \notin \mathfrak{R}_i$. We select the greater value $v_a + (n^\theta - 1)\delta_1(r)$ as the upper bound of $v_i(r)$. Note that $v_a + (n^\theta - 1)\delta_1(r)$ is also the upper bound of $v_i(\tilde{r})$ for any round $\tilde{r} > r$. Again if $\tilde{r} \in \mathfrak{R}_i$, we have $v_i(\tilde{r}) \leq v_a$. If $\tilde{r} \notin \mathfrak{R}_i$, by using the same extreme estimation, we get $v_i(\tilde{r}) \leq v_a + (n^\theta - 1)\delta_1(\tilde{r})$. As δ_1 is non-increasing, we have $\delta_1(\tilde{r}) \leq \delta_1(r)$. Thus $v_i(\tilde{r}) \leq v_a + (n^\theta - 1)\delta_1(r)$. Note that the validity property ensures that the lower bound of $v_i(r)$ is $v_a - \delta_1(r)$. Obviously, for any node p_i of \tilde{V}_c , when $\delta_1(r)$ converges to 0, the upper bound of $v_i(r)$ converges to v_a and the lower bound of $v_i(r)$ converges also to v_a . Thus, $v_i(r)$ must converge to v_a . We conclude the existence of V'_c (which is equal to \tilde{V}_c). ■

Corollary 3. *Suppose $\forall r \in \mathfrak{R}$, $\delta_1(r) > 0$. For any correct node p_k , if v_k does not converge to v_a , then $\exists r^* \in \mathfrak{R}$ and $\exists \ell_{low} > v_a$, such that $\forall r \geq r^*$, $v_k(r) \geq \ell_{low}$.*

Proof: if v_k does not converge to v_a , there exists a value α such that $\forall r \in \mathfrak{R}$, $\exists \hat{r} > r$, such that $v_k(\hat{r}) - v_a \geq \alpha > 0$. There are two cases corresponding to the non-convergence. First, since a round r , $v_k(r)$ is always greater than or equal to $v_a + \alpha$. In this case, $v_a + \alpha$ is a lower bound for v_k since the round r . In the second case, from time to time v_k is less than $v_a + \alpha$ and from time to time it is greater than or equal to $v_a + \alpha$. If this scenario happens, each time v_k is less than $v_a + \alpha$, after a limited number of rounds, it must again become greater than or equal to $v_a + \alpha$. So p_k must have infinite opportunities to compute an average with a value greater than $v_a + \alpha$. Formally, there is an infinite set $\mathfrak{R}_k \subset \mathfrak{R}$, such that $\forall r \in \mathfrak{R}_k$ node p_k computes the average with at least a value greater than $v_a + \alpha$. For any consecutive elements r, r' of \mathfrak{R}_k , they satisfy $r' - r \leq \theta < +\infty$ (θ can be any large integer but not infinite). Consider a round $r \in \mathfrak{R}$

with $r > \theta$, we know $\exists \tilde{r} \in \mathfrak{R}_k$ with $r - \tilde{r} \leq \theta$. We consider an extreme scenario to estimate the lower bound of $v_k(r)$. Assume that p_k computes an average with $v_k(\tilde{r})$, $v_a + \alpha$ and $(n - 2)(v_a - \delta_1(\tilde{r}))$ during round \tilde{r} . Then from the round $\tilde{r} + 1$ to the round $r - 1$, p_k continuously decrease its value by computing an average with its own value and $n - 1$ values equal to $v_a - \delta_1(\tilde{r})$. Thus, when the round r begins, we have $v_k(r) \geq \frac{v_k(\tilde{r}) + (v_a + \alpha) + (n^\theta - 1)(v_a - \delta_1(\tilde{r}))}{n^\theta}$. As $v_k(\tilde{r}) \geq v_a - \delta_1(\tilde{r})$, we get $v_k(r) \geq \frac{v_a + \alpha + n^\theta(v_a - \delta_1(\tilde{r}))}{n^\theta}$. As δ_1 is non-increasing and converges to 0, two properties hold: first $\frac{v_a + \alpha + n^\theta(v_a - \delta_1(\tilde{r}))}{n^\theta}$ is non-decreasing and second when $v_a + \alpha > n\delta_1(\tilde{r})$ we have $\frac{v_a + \alpha + n^\theta(v_a - \delta_1(\tilde{r}))}{n^\theta} > v_a$. Thus when the round r is large enough, $v_k(r)$ has a lower bound $\ell_{low} > v_a$. ■

The twin corollary of Corollary 3 is given below.

Corollary 4. *Suppose $\forall r \in \mathfrak{R}$, $\delta_2(r) > 0$. For any correct node p_k , if v_k does not converge to v_b , then $\exists r^* \in \mathfrak{R}$ and $\exists \ell_{up} < v_b$, such that $\forall r \geq r^*$, $v_k(r) \leq \ell_{up}$.*

Note that Hypothesis 1 includes the following situation: $\exists r' \in \mathfrak{R}$ such that $\forall r \geq r'$, either $\delta_1(r) = 0$ or $\delta_2(r) = 0$. In the following, we simply write $\delta_1(r') = 0$ or $\delta_2(r') = 0$ when the value equal to 0 since a round r' . We write $\delta_1(r) > 0$ or $\delta_2(r) > 0$ when the value is greater than 0 in any round r . There are three possible cases: (1) $\delta_1(r) > 0$ and $\delta_2(r') = 0$; (2) $\delta_1(r') = 0$ and $\delta_2(r) > 0$; (3) $\delta_1(r) > 0$ and $\delta_2(r) > 0$. In each case, we consider a partition of V_c into three subsets. We use the same notation $\langle V_c^1, V_c^2, V_c^3 \rangle$ to refer to this partition. The definitions of the sets are different in each case. If $\delta_1(r) > 0$, then based on Lemma 4, $V_c^1 = \{p_i | p_i \in V_c \wedge v_i \text{ converges to } v_a\}$. If $\delta_1(r') = 0$, then $V_c^1 = \{p_i | p_i \in V_c \wedge \forall r \geq r', v_i(r) = v_a\}$. Symmetrically, If $\delta_2(r) > 0$, $V_c^2 = \{p_j | p_j \in V_c \wedge v_j \text{ converges to } v_b\}$. If $\delta_2(r') = 0$, $V_c^2 = \{p_j | p_j \in V_c \wedge \forall r \geq r', v_j(r) = v_b\}$. Whatever the situation, $V_c^3 = V_c / (V_c^1 \cup V_c^2)$: the third set contains the remaining correct nodes.

Lemma 5. *Hypothesis 1 is false.*

Proof: Clearly, $\langle F, L = V_c^1, M = V_c^3, R = V_c^2 \rangle$ forms a F partition and belongs to the set Γ . The first conclusion is $\exists r^* \in \mathfrak{R}$, such that $\forall r \geq r^*$ both $\max(L, r) < \min(\mathbb{R}, r)$ and $\min(R, r) > \max(\mathbb{L}, r)$ are satisfied. Consider the set L . There are two cases. First, $\delta_1(r') = 0$. The definition of V_c^1 in this case implies $\exists r^* \geq r'$, such that $\forall r \geq r^*$ only the correct nodes in L still hold the minimum value v_a . Consequently, $\forall r \geq r^*$, $\max(L, r) < \min(\mathbb{R}, r)$. Second, $\delta_1(r) > 0$. In this case we know $\forall p_i \in L$, v_i converges to v_a (Lemma 4). Furthermore, $\forall p_k \in M \cup R$, v_k does not converge to v_a determines it has a lower bound $\ell_{low} > v_a$ (Corollary 3). Therefore, $\exists r^* \in \mathfrak{R}$, such that $\forall r \geq r^*$, $\max(L, r) < \min(\mathbb{R}, r)$. By considering the set R , we can conclude $\exists r^* \in \mathfrak{R}$, then $\forall r \geq r^*$, $\min(R, r) > \max(\mathbb{L}, r)$. The proof is symmetric.

The above conclusion ensures Lemma 2 can be applied in each round $r \geq r^*$. Although different values of δ_1 and δ_2 leads to three different cases, whatever the case, based on Lemma 2 and the fact that the number of nodes is finite, we can infer that at least one node on one side ($\exists p_i \in L$ and/or $\exists p_j \in R$) has infinite opportunities to update its value with a value from (v_a, v_b) . Without loss of generality, suppose since a round r^* , $\exists p_i \in L$, such that p_i has infinite opportunities to update its value with a value that has a lower bound greater than v_a . Formally, there exists an infinite set $\mathfrak{R}_i \subset \mathfrak{R}$, such that $\forall r \in \mathfrak{R}_i$ node p_i computes the average with at least a value with a lower bound greater than v_a . The infinite property of \mathfrak{R}_i ensures that for any consecutive elements r, r' of \mathfrak{R}_i , they satisfy $r' - r \leq \theta < +\infty$ (θ can be any large integer but not infinite). The way we have proved Corollary 3 also works here. By using an extreme estimation, we can conclude there exists a round and from then on v_i has lower bound greater than v_a . This contradicts the fact that $p_i \in L$ ($L = V_c^1$) which means that either v_i converges to v_a or is always equal to v_a . Hypothesis 1 is false. ■

Theorem 3. *Approximate Byzantine consensus can be solved in G if and only if G is f -resilient.*

VI. h -DISJOINT AND f -RESILIENT

As mentioned in Section IV, the use of the primitive 3-partial multicast has been investigated to solve exact Byzantine consensus. Based on uniform hyperedges, the property h -disjoint proposed in [4] is proved to be necessary and sufficient for reaching exact Byzantine consensus. The properties h -disjoint and f -resilient provide different (but strongly related) requirements on the communication topology. An interesting future work would be to determine if there exists a generalized property that copes with both exact Byzantine consensus and approximate Byzantine consensus. The comparison of h -disjoint and f -resilient done in this paper is a first step in this direction.

We summarize the property h -disjoint first. It requires that for all disjoint nodes subsets A, B, C , they satisfy the follow condition, $\exists p'_i \in A, \exists p'_j \in B, \exists p'_k \in C$, such that (p'_i, p'_j, p'_k) is a uniform hyperedge and $|A|, |B|, |C| \geq h$, $|A| + |B| + |C| \geq \frac{n+3h}{2}$. For a uniform hypergraph H with $2f + 1 \leq n \leq 3f$, the exact Byzantine consensus is solvable if and only if H is $(n - 2f)$ -disjoint.

To do a comparison, a standard model is needed to cover both the uniform hypergraph H and the hybrid graph G . We expand H and define unicast channels: $H = (\mathcal{V}, \mathcal{E}_U, \mathcal{E}_M)$. There is a one-one mapping between the two sets \mathcal{V} and V . We use the notation $p'_i \asymp p_i$ ($p'_i \in \mathcal{V}, p_i \in V$) to express it. Suppose $|\mathcal{V}| = |V| = n$ and $2f + 1 \leq n \leq 3f$. Regarding the unicast channels, the dual sets \mathcal{E}_U and E_U are equivalent: $(p'_i; p'_j) \in \mathcal{E}_U$ if and only if $(p_i; p_j) \in E_U$. According to h -disjoint, the set \mathcal{E}_M only contains uniform hyperedges. We assume that $(p'_i, p'_j, p'_k) \in \mathcal{E}_M$ if and only if

$(p_i; p_j, p_k), (p_j; p_i, p_k), (p_k; p_i, p_j) \in E_M$. The connectivity of both communication graphs H and G is at least $2f + 1$, which implies that each node has at least $2f + 1$ source neighbors (an exception is $n = 2f + 1$, if so, consider each node has $2f$ source neighbors).

Lemma 6. *Let $n = 2f + h$ ($1 \leq h \leq f$). If H is not h -disjoint, then G is not f -resilient.*

Proof: When H is not h -disjoint, there exists three subsets $A, B, C \subseteq \mathcal{V}$ such that $|A|, |B|, |C| \geq h$, $|A| + |B| + |C| \geq (n + 3h)/2$, and $\forall p'_i \in A, \forall p'_j \in B, \forall p'_k \in C, (p'_i, p'_j, p'_k) \notin \mathcal{E}_3$. Due to $|A|, |B|, |C| \geq h$, we can conclude that at most one set among A, B and C can have a cardinality greater than f (otherwise $|A| + |B| + |C| > n$). The proof below includes two cases. First, among A, B, C only one set has a cardinality greater than f and second, all the three sets have a cardinality less than or equal to f . Due to $|A| + |B| + |C|$ may be less than n , suppose $D = \mathcal{V}/(A \cup B \cup C)$ (if $A \cup B \cup C = \mathcal{V}, D = \emptyset$).

Consider the first case: among A, B, C only one set has a cardinality greater than f . Without loss of generality, suppose $|A| > f$. Let $A' \subset A$ and $|A| - |A'| = f$. Due to $|A| > f$, we know $A' \neq \emptyset$. Consider the F partition of V with $F \asymp A/A', L \asymp B, M \asymp D \cup A'$ and $R \asymp C$. Note that no uniform hyperedges among A, B, C implies no 3-partial multicast channel among F, L and R . Due to $|F| = f, |L| = |B| \geq h, |R| = |C| \geq h$ and $n = 2f + h$, we have $\forall p_i \in L, \forall p_j \in R, |N_i \cap \mathbb{R}| \leq |\mathbb{R}| \leq f$ and $|N_j \cap \mathbb{L}| \leq |\mathbb{L}| \leq f$. Consequently, $|F_{ij}| + |N_i \cap \mathbb{R}| + |N_j \cap \mathbb{L}| \leq 0 + f + f < 2f + 1$, G is not f -resilient.

Now consider the second case: among A, B, C no set has a cardinality greater than f . There are two subcases depending on $|A| + |D| \leq f$ or $|A| + |D| > f$. Consider $|A| + |D| \leq f$ first. Let $B' \subset B, C' \subset C$ with $|A| + |B'| + |C'| + |D| = f$ (if $|A| + |D| = f$, then $B' = C' = \emptyset$). Consider the following F partition of V with $F \asymp A \cup B' \cup C' \cup D, L \asymp B/B', M \asymp \emptyset$ and $R \asymp C/C'$. Note that $|A| \geq h$ and no uniform hyperedge is among A, B, C , such that $\forall p_i \in L, \forall p_j \in R$, we have $|F_{ij}| \leq f - h$. Moreover, $M = \emptyset$ implies $|\mathbb{L}| + |\mathbb{R}| = 2f + h - f = f + h$. Consequently, $|F_{ij}| + |N_i \cap \mathbb{R}| + |N_j \cap \mathbb{L}| \leq f - h + f + h = 2f < 2f + 1$, G is not f -resilient in this subcase. Now consider $|A| + |D| > f$. To facilitate the presentation, we include four parameters a, b, c, d to express $|A|, |B|, |C| \geq h$: $|A| = h + a, |B| = h + b, |C| = h + c$ and $|D| = d$. Due to $|A| + |B| + |C| \geq (n + 3h)/2 = f + 2h$, thus $d \leq n - (f + 2h) = f - h$. $|A| < f$ and $|A| + |D| > f$ in this case implies $\exists D' \subset D$, such that $|D'| + |A| = f$. Let $|D'| = d'$, thus $d' + a = f - h$. Note that we also have $a + b + c + d = 2f + h - 3h = 2f - 2h$, such that $b + c = 2f - 2h - (d' + a + d - d')$. Due to $d' + a = f - h$, thus $b + c = f - h - (d - d')$. Consider the F partition of V with $F \asymp A \cup D', L \asymp B, M \asymp D/D'$ and $R \asymp C$, such that $|\mathbb{L}| = h + b + d - d'$ and $|\mathbb{R}| = h + c + d - d'$. No uniform hyperedge is among A, B, C , thus $\forall p_i \in L$

and $\forall p_j \in R$, we have $|F_{ij}| \leq |D'| = d'$. Consequently, $|F_{ij}| + |N_i \cap \mathbb{R}| + |N_j \cap \mathbb{L}| \leq |F_{ij}| + |\mathbb{R}| + |\mathbb{L}| \leq d' + (h + b + d - d') + (h + c + d - d') = 2h + 2d - d' + (b + c)$. We already get $d \leq f - h$ and $b + c = f - h - (d - d')$, thus $2h + 2d - d' + (b + c) = h + f + d \leq h + f + f - h < 2f + 1$, G is not f -resilient. ■

In another aspect, H is h -disjoint does not always mean G is f -resilient. We find a counterexample in [6] (page 452). A 1-disjoint hypergraph for $n = 5$ and $f = 2$ is provided, but this example cannot satisfy 2-resilient.

VII. ASYNCHRONOUS NETWORKS

In asynchronous networks, the message transmission and processing delay is arbitrary but finite. Under the model asynchrony, no correct node can differentiate if another node is just slow or stops working. Therefore, in each algorithm round, a correct node p_i can only expect to receive the values from its $|N_i| - f$ sources neighbors. Solving approximate Byzantine consensus in asynchronous networks also benefits from the 3-partial multicast primitive. We claim that the idea behind f -resilient can be extended straightforward into asynchronous networks by increasing the requirement of cardinalities (see Definition 6).

Definition 6. $\langle F, L, M, R \rangle$ is an **asy-safe** F partition of G if and only if one of the following properties holds:

- (1) $\exists p_i \in L$ such that $|N_i \cap \mathbb{R}| \geq 2f + 1$ or $\exists p_j \in R$ such that or $|N_j \cap \mathbb{L}| \geq 2f + 1$;
- (2) $\exists p_i \in L$ and $\exists p_j \in R$, such that $f + 1 \leq |N_i \cap \mathbb{R}| \leq 2f$, $f + 1 \leq |N_j \cap \mathbb{L}| \leq 2f$, $|F_{ij}| + |N_i \cap \mathbb{R}| + |N_j \cap \mathbb{L}| \geq 4f + 1$.

VIII. CONCLUSION

We propose a novel condition f -resilient, which takes into account both unicast and partial multicast communication primitives. f -resilient is proved to be necessary and sufficient for reaching iterative approximate consensus in synchronous networks. The relationships between h -disjoint and f -resilient are also investigated. Finally, the idea behind f -resilient is extended to cope with asynchronous networks.

ACKNOWLEDGMENT

This work is partially supported by National 973 Hi-Tech Program, China & by a Mayenne's military project NoPEC.

REFERENCES

- [1] D. Dolev, A. N. Lynch, S. Pinter, W. E. Stark, and E. W. Weihl. Reaching approximate agreement in the presence of faults. In *proc. of the IEEE symp. on reliability in distributed software and database systems*, pages 145–154, 1983.
- [2] N. Vaidya, L. Tseng, and G. Liang. Iterative approximate byzantine consensus in arbitrary directed graphs. In *proc. of the ACM symp. on principles of distributed computing*, pages 365–374, 2012.
- [3] H. LeBlanc, H. Zhang, S. Sundaram, and X. Koutsoukos. Consensus of multi-agent networks in the presence of adversaries using only local information. In *proc. of the Int. conf. on high confidence networked systems*, pages 1–10, 2012.
- [4] A. Jaffe, T. Moscibroda, and S. Sen. On the price of equivocation in byzantine agreement. In *proc. of the ACM symp. on principles of distributed computing*, pages 309–318, 2012.
- [5] M. Fitzi and U. Maurer. From partial consistency to global broadcast. In *proc. of the ACM symp. on theory of computing*, pages 494–503, 2000.
- [6] D.V.S. Ravikant, V. Muthuramakrishnan, V. Srikanth, K. Sri-nathan, and C. P. Rangan. On byzantine agreement over $(2,3)$ -uniform hypergraphs. In *proc. of the Int. symp. on distributed computing*, pages 450–464, 2005.
- [7] M. Franklin and M. Yung. Secure hypergraphs: Privacy from partial broadcast. In *proc. of the ACM symp. on the theory of computing*, pages 36–44, 1995.
- [8] M. H. Azadmanesh and H. Bajwa. Global convergence in partially fully connected networks (PFCN) with limited relays. In *proc. of the IEEE conf. on industrial electronics society*, pages 2022–2025, 2001.
- [9] M. H. Azadmanesh and A. W. Krings. A step toward global convergence in partially connected networks. In *proc. of the Int. conf. on parallel and distributed computing systems*, pages 234–241, 1997.
- [10] S. Sundaram and C. N. Hadjicostis. Distributed function calculation via linear iterations in presence of malicious agents - part i: Attacking the networks. In *proc. of the American control conf.*, pages 1350–1355, 2008.
- [11] S. Sundaram and C. N. Hadjicostis. Distributed function calculation via linear iterations in presence of malicious agents - part ii: Overcoming malicious behavior. In *proc. of the American control conf.*, pages 1356–1361, 2008.
- [12] C. Li, M. Hurfin, and Y. Wang. Approximate byzantine consensus in sparse, mobile ad-hoc networks. *Journal of Parallel and Distributed Computing, Elsevier*, 74:2860–2871, June 2014.
- [13] B. Charron-Bost, M. Függer, and T. Nowak. Approximate consensus in highly dynamic networks: The role of averaging algorithms. In *proc. of the Int. Conf. on Automata, Languages, and Programming - ICALP*, pages 528–539, 2015.
- [14] I. Abraham, Y. Amit, and D. Dolev. Optimal resilience asynchronous approximate agreement. In *proc. of the Int. conf. on principles of distributed Systems, LNCS 3544*, pages 229–239, 2005.
- [15] A. Clement, F. Junqueira, A. Kate, and R. Rodrigues. On the (limited) power of non-equivocation. In *proc. of the ACM symp. on principles of distributed computing*, pages 301–308, 2012.
- [16] M. Backes, F. Bendun, A. Choudhury, and A. Kate. Asynchronous MPC with a strict honest majority using non-equivocation. In *proc. of the ACM symp. on principles of distributed computing*, pages 10–19, 2014.