Tight Bounds for Asymptotic and Approximate Consensus

Matthias Függer, Thomas Nowak, Manfred Schwarz

► To cite this version:

Matthias Függer, Thomas Nowak, Manfred Schwarz. Tight Bounds for Asymptotic and Approximate Consensus. Journal of the ACM (JACM), 2021, 10.1145/3485242 . hal-03408731

HAL Id: hal-03408731

<https://hal.science/hal-03408731>

Submitted on 29 Oct 2021

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

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

Tight Bounds for Asymptotic and Approximate Consensus

MATTHIAS FÜGGER, CNRS, LMF, ENS Paris-Saclay, Université Paris-Saclay, Inria, France

THOMAS NOWAK, Université Paris-Saclay, CNRS, France

MANFRED SCHWARZ, TU Wien, Austria

Agreeing on a common value among a set of agents is a fundamental problem in distributed computing, which occurs in several variants: In contrast to exact consensus, approximate variants are studied in systems where exact agreement is not possible or required, e.g., in man-made distributed control systems and in the analysis of natural distributed systems, such as bird flocking and opinion dynamics.

We study the time complexity of two classical agreement problems: non-terminating asymptotic consensus and terminating approximate consensus. Asymptotic consensus, requires agents to repeatedly set their outputs such that the outputs converge to a common value within the convex hull of initial values; approximate consensus requires agents to eventually stop setting their outputs, which must then lie within a predefined distance of each other.

We prove tight lower bounds on the contraction ratios of asymptotic consensus algorithms subject to oblivious message adversaries, from which we deduce bounds on the time complexity of approximate consensus algorithms. In particular, the obtained bounds show optimality of asymptotic and approximate consensus algorithms presented by Charron-Bost et al. [ICALP'16] for certain systems, including the strongest oblivious message adversary in which asymptotic and approximate consensus are solvable. As a corollary we also obtain asymptotically tight bounds for asymptotic consensus in the classical asynchronous model with crashes.

Central to the lower-bound proofs is an extended notion of valency, the set of reachable limits of an asymptotic consensus algorithm starting from a given configuration. We further relate topological properties of valencies to the solvability of exact consensus, shedding some light on the relation of these three fundamental problems in dynamic networks.

CCS Concepts: • **Theory of computation** → **Distributed algorithms**.

Additional Key Words and Phrases: Asymptotic consensus; approximate consensus; dynamic networks; message adversaries; crash faults; lower bounds

1 INTRODUCTION

This work is devoted to the problem of achieving symmetry among a set of agents in a distributed system, which is studied in three variants:

In the *asymptotic consensus* problem a set of agents, each starting from an initial value in $\mathbb{R}^d$, repeatedly update their values such that all agents' values converge to a common limit within the convex hull of initial values. The problem is closely related to the *approximate consensus* problem, in which agents have to irrevocably decide on values that lie within a predefined distance $\varepsilon > 0$ of each other. Typically, initial values are constrained to be within an initial distance $\Delta \geq 0$ of each other; which we will also assume in this work. The approximate consensus problem is weaker than the *exact consensus* problem in which agents need to decide on the same value, which must be one

Authors' addresses: Matthias Függer, CNRS, LMF, ENS Paris-Saclay, Université Paris-Saclay, Inria, 4 avenue des Sciences, 91190, Gif-sur-Yvette, France, mfuegger@lsv.fr; Thomas Nowak, Université Paris-Saclay, CNRS, Laboratoire Interdisciplinaire des Sciences du Numérique, 1 rue Raimond Castaing, 91190, Gif-sur-Yvette, France, thomas.nowak@lri.fr; Manfred Schwarz, TU Wien, Embedded Computing Systems Group, Treitlstrasse 3, 1140, Vienna, Austria, mschwarz@ecs.tuwien.ac.at.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2021 Copyright held by the owner/author(s).

0004-5411/2021/10-ART

<https://doi.org/10.1145/3485242>

of the initial values. If the values are further constrained to be from $\{0, 1\}$, one speaks of binary exact consensus.

While exact consensus has traditionally received the most attention, motivated by classical distributed computing problems such as providing consistency among replicated databases, the weaker asymptotic and approximate consensus problems have gained importance in systems where exact consensus is not possible or not required. In particular, asymptotic consensus has been studied in the context of distributed control [4, 8, 10, 20, 27, 47] in different variants. Indeed, exact consensus requires quite strong assumptions on the underlying network [31, 54]. Fortunately, exact consensus is an unnecessarily strong requirement for several applications that require agreement: both the asymptotic and the approximate consensus problems have not only a variety of applications in the design of man-made control systems like sensor fusion [6], clock synchronization [42], formation control [26], rendezvous in space [43] and robot gathering [2, 13, 21], or load balancing [23], but also for analyzing natural systems like bird flocking [56], firefly synchronization [46], or opinion dynamics [36]. These problems often have to be solved under harsh environmental restrictions in which exact consensus is not achievable, or too costly to achieve: with limited computational power and local storage, under restricted communication abilities, and in presence of communication uncertainty.

In this work we study the performance of both asymptotic and approximate consensus algorithms under such harsh conditions. Specifically, we study algorithms subject to *oblivious message adversaries* [9, 14, 41, 57] with round-based computation and a dynamic communication topology whose *directed* communication graphs are chosen each round from a predefined set of communication graphs.

While this model naturally captures highly unstable communication topologies, we further show that it also allows to assess performance within classical, more stable, distributed fault models. For example, the detour via dynamic network models allows to prove an asymptotically tight lower bound on the contraction ratio and thus the time complexity of approximate agreement in distributed systems where crash-prone agents communicate by asynchronous message passing. The results in this work thus answer the question raised by Dolev et al. [25, Section 7] whether algorithms that are not of certain standard forms can produce agreement faster: our lower bounds allow algorithms that may arbitrarily use information from previous rounds.

In the following we will briefly discuss previous work on asymptotic and approximate consensus, summarize contributions of this work, and relate it to previous work. We start with a brief discussion on previous work by Charron-Bost et al. [16–18] on approximate agreement in oblivious message adversaries.

Approximate Consensus: Algorithms and Solvability. In previous work [17], Charron-Bost et al. showed that approximate consensus is solvable subject to an oblivious message adversary if and only if the adversary is *rooted*, i.e., all its communication graphs contain rooted spanning trees. These rooted spanning trees need not have any edges in common and can change from one communication round to the next.

Interestingly, solvability subject to any rooted oblivious message adversary is already provided by surprisingly simple algorithms [17]: so-called *averaging* or *convex-combination* algorithms, in which agents repeatedly broadcast their current value and update it to some weighted average of the values they received in this round, until they decide in a precomputed round. Central to the proof is a *contraction of output values* by a factor $c < 1$ every $k \geq 1$ rounds; the parameters c and k depend on the convex-combination algorithm and the message adversary. In particular,

Charron-Bost et al. [16, 18] show that

$$\Delta(t) \leq c \cdot \Delta(t - k) , \quad (1)$$

for $t \geq k$, where $\Delta(t)$ is the maximum distance among, i.e., the diameter of, agent values at the end of round $t \geq 1$ and $\Delta(0) = \Delta$ is the initial diameter.

One instance of a convex-combination algorithm for scalar input and output values, presented by Charron-Bost et al. is the *midpoint algorithm*, in which agents update their value to the midpoint of the set of received values, i.e., the average of the smallest and the largest of the received values. For input and output values of higher dimension $d > 1$, the midpoint algorithm is not well-defined. While one can show that applying it coordinate-wise in dimension $d = 2$ results again in a convex-combination algorithm, this is not anymore the case for dimension $d \geq 3$ since the new output value may lie outside of the convex hull of received values. An example for a convex-combination algorithm for arbitrary dimensions is the *MidExtremes algorithm* [32]. It generalizes the midpoint algorithm to higher dimensions by having an agent set its output onto the midpoint of the line segment spanned by two received values that have maximum distance.

Asymptotic Consensus: Algorithms and Solvability. The characterization by Charron-Bost et al. [15, 17] of solvability of approximate consensus can be immediately generalized to asymptotic consensus, showing that the problem is solvable subject to an oblivious message adversary if and only if the adversary is rooted; like approximate consensus: (i) Using the same convex-combination algorithms, except with the decision rule removed, the proved contraction of the output values by a factor of $c < 1$ every k rounds, implies that these algorithms also solve asymptotic agreement subject to any rooted oblivious message adversary. (ii) The impossibility of solving asymptotic consensus subject to any non-rooted oblivious message adversaries follows by the same partitioning argument as in the proof for approximate consensus [15, Theorem 9].

An interesting special case of rooted oblivious message adversaries are those whose graphs are *non-split*, that is, any two agents have a common incoming neighbor. Their prominent role is motivated by two properties: (i) They occur as communication graphs in benign classical distributed failure models, for example in synchronous systems with crashes, in asynchronous systems with a minority of crashes, and synchronous systems with send omissions. (ii) Charron-Bost et al. [17] showed that non-split graphs also play a central role for arbitrary rooted oblivious message adversaries: they showed that any product of $n - 1$ rooted graphs with n nodes is non-split, allowing to transform approximate and asymptotic consensus algorithms for non-split oblivious message adversaries into their *amortized* variants for rooted oblivious message adversaries.

Approximate and Asymptotic Consensus: Performance. Regarding time complexity, for scalar input and output values, the amortized midpoint algorithm subject to rooted oblivious message adversaries was shown [18] to fulfill (1) with $c = \frac{1}{2}$ and $k = n - 1$. On average, the algorithm thus leads to a contraction of the output range by a factor of $\sqrt[n-1]{\frac{1}{2}}$ per round. Stating it like this allows one to compare algorithms with different values of k . As a consequence of the above contraction of the output range, values that are initially at most Δ apart, have distance at most $\varepsilon > 0$ after $(n - 1) \lceil \log_2 \frac{\Delta}{\varepsilon} \rceil$ rounds [16, 18, Theorem 9].¹

While termination times are inherently unsuited to assess the performance of asymptotic consensus algorithms, the measure in (1) may be applied. However, measuring the contraction of the *range of outputs* is problematic for the following reason: Take any algorithm that solves asymptotic consensus and replace it by one that transiently sets an agent output value to 1 in round 1, while executing the original algorithm in successive rounds on the original initial states. The outputs rapidly

¹The result in [16, 18] is stated for the normalized case with $\Delta = 1$.

contract (temporarily) to a single value within the first round, yielding an arbitrarily fast contraction of outputs in the first round. In the second round, however, the outputs will expand again. The result is misleading in this case. Intuitively, we would expect a good measure of convergence to be invariant to such algorithmic modifications.

Replacing the agent outputs of the configuration at the end of round t by the *valency of the configuration at the end of round t* , i.e., the limits that are reachable from this configuration, solves this problem. Indeed, from the perspective of valency, both of the above algorithms behave identically since the modification to the algorithm did not prune the set of reachable limits.

Instead of (1), we thus study

$$\delta(t) \leq c \cdot \delta(t - k) , \quad (2)$$

where $\delta(t)$ is the diameter of the valency of the configuration at the end of round t . To compare algorithms with different values of k , and since asymptotic consensus algorithms do not have to guarantee a regular contraction according to (2), we also study the *contraction ratio* of an algorithm subject to a message adversary as the supremum over all its executions of

$$\limsup_{t \rightarrow \infty} \sqrt[k]{\delta(t)} . \quad (3)$$

First note that, by definition, the contraction ratio is between 0 and 1. The limit superior ensures the existence of a contraction ratio in case the limit does not exist and gives a conservative bound. One also immediately observes: (i) If an asymptotic consensus algorithm guarantees a regular contraction as in (2), its contraction ratio is $\sqrt[k]{c}$, the average contraction per round. (ii) If an asymptotic consensus algorithm in fact solves exact consensus, the agent outputs will not change anymore after all agents have decided. Since from this round t on, $\delta(t) = 0$, we have that the contraction ratio is 0 for this case.

From the fact that convex-combination algorithms set their output to within the convex hull of the previously received values, it follows that $\delta(t) \leq \Delta(t)$ for all $t \geq 0$ (Lemma 3.1). Consequently, results on the contraction of outputs by convex-combination algorithms can be directly translated into upper bounds on their contraction ratio. For example, the contraction of the output range by the midpoint algorithm subject to a non-split oblivious message adversary of $\frac{1}{2}$ per round [18] implies a contraction ratio of $\frac{1}{2}$ or less. Likewise, for higher dimensions, the contraction of the diameter of the set of output values by the MidExtremes algorithm subject to a non-split oblivious message adversary, of $\sqrt{7/8}$ per round [32, Theorem 1] implies a contraction ratio of $\sqrt{7/8}$ or less.

A natural question is whether *non-convex-combination or non-memoryless algorithms*, i.e., algorithms that (i) do not necessarily set their output values to within the convex hull of previously received values or (ii) whose output is a function not only of the previously received values, allow faster contraction. In the context of classical failure models, the question for lower bounds independent of such assumptions, was raised by Dolev et al. [25]. As an example for (i), consider the algorithm where each agent sends an equal fraction of its current output value to all out-neighbors and sets its output to the sum of values received in the current round. Note that the algorithm is not a convex-combination algorithm as its output may lie outside the convex hull of the values of its in-neighbors. However, it solves asymptotic consensus for a fixed directed communication graph. Other examples of algorithms that violate (i) and (ii) are from control theory, where the usage of overshooting fast second-order controllers is common (see, e.g., [5]).

Contribution. We prove asymptotically tight lower bounds on the contraction ratio of any asymptotic consensus algorithm with scalar input and output values, regardless of the structure of the algorithm: algorithms can be full-information and agents can set their outputs outside the convex hull of received values. This, e.g., includes using higher-order filters, in contrast to the 0-order

filters of convex-combination algorithms. In particular, the following lower bounds hold for an oblivious message adversary $\mathcal{M}$ with n agents: If exact consensus is solvable subject to $\mathcal{M}$, an optimal contraction ratio of 0 can be achieved. Otherwise:

- In a system with $n = 2$ agents, the contraction ratio is lower-bounded by $1/3$ (Theorem 4.1). This is tight (algorithm presented by Charron-Bost et al. [18] and Lemma 3.1).
- For an arbitrary communication graph G , we define the set $\text{deaf}(G) = \{F_1, \dots, F_n\}$, where F_i is derived from G by making agent i *deaf* in F_i , i.e., removing the incoming edges of i in G . In a system with $n \geq 3$ agents, if $\mathcal{M}$ contains the communication graphs in $\text{deaf}(G)$, then the contraction ratio is lower-bounded by $1/2$ (Theorem 5.1). This is tight for oblivious non-split message adversaries because of the midpoint algorithm [18] and Lemma 3.1.
- We then show that if $\mathcal{M}$ contains certain rooted graphs Ψ , then the contraction ratio is lower bounded by $\sqrt[n-2]{1/2}$ (Theorem 6.1). This is asymptotically tight for oblivious rooted message adversaries because of the amortized midpoint algorithm [18] and Lemma 3.1. Calling an oblivious message adversary *stronger* than another one if its set of communication graphs is a superset of the other one's set of communication graphs, this specifically proves optimality of the amortized midpoint algorithm subject to the strongest oblivious message adversary in which asymptotic and approximate consensus is solvable: the oblivious message adversary whose set of communication graphs is the set of *all* directed rooted communication graphs. The maximality of this set follows from the fact that if it contains a non-rooted communication graph, asymptotic consensus becomes unsolvable [17].
- For arbitrary oblivious message adversaries we show that in a system with $n \geq 3$ agents, any asymptotic consensus algorithm must have a contraction ratio of at least $1/(D + 1)$, where D , the so-called α_G -diameter of $\mathcal{M}$, i.e., the smallest value which allows a connection of any pair of communication graphs that occur in $\mathcal{M}$ via an indistinguishability chain of length at most D (Theorem 7.12).
- We demonstrate how to apply the above mentioned bound to obtain new lower bounds on contraction ratios for classical failure models as an immediate corollary. Specifically, we consider asynchronous message-passing system of size n with up to $f < n/2$ crashes. For such systems, algorithms operating in *asynchronous rounds* are widely used [19, 25, 44]: each agent broadcasts its round message, waits for $n - f$ messages of the current round, updates its state based on the received messages and its previous state, and advances to the next round.
- With the above results, we show that it immediately follows that no algorithm operating in asynchronous rounds can achieve a contraction ratio better than $\frac{1}{\lceil n/f \rceil + 1}$ (Theorem 8.2). This shows that the asynchronous algorithms for systems of size $n > 5f$ with up to f Byzantine failures by Dolev et al. [25] and for systems of size $n > 2f$ with up to f crashes by Fekete [29] have asymptotically optimal contraction ratios for round-based algorithms.
- We then present an algorithm for $n > f$ that does *not* operate in asynchronous rounds and achieves a contraction ratio of 0, demonstrating a large gap between round-based and non round-based algorithms for asymptotic consensus.

Table 1 summarizes lower and upper bounds.

Our lower bounds also hold for asymptotic consensus with input and output values of arbitrary dimension. For example, for the oblivious message adversary $\mathcal{M}$ that contains $\text{deaf}(G)$ for a communication graph G with $n \geq 3$ nodes, we obtain: From the fact that the MidExtremes algorithm has a contraction ratio of $\sqrt{7/8}$ for oblivious non-split message adversaries (MidExtremes algorithm [32] and Lemma 3.1) and from the lower bound shown in this work (Theorem 5.1), the optimal contraction ratio for higher dimensions is between $1/2$ and $\sqrt{7/8}$ for the oblivious message adversary $\mathcal{M}$.

agents	oblivious message adversary			asynchronous + f crashes	
	general non-split	non-split with $\alpha_{\mathcal{G}}$ -diameter D	general rooted	round-based alg. $0 < f < \frac{n}{2}$	arbitrary alg. $0 < f < n$
$n = 2$	$\frac{1}{3}^*$	0 or $\frac{1}{3}^*$	$\frac{1}{3}^*$	N/A	0^*
$n \geq 3$	$\frac{1}{2}^*$	0 or $\left[\frac{1}{D+1}^*, \frac{1}{2}\right]$	$\left[\sqrt[n-2]{\frac{1}{2}}^*, \sqrt[n-1]{\frac{1}{2}}\right]$	$\left[\frac{1}{\lceil n/f \rceil + 1}^*, \frac{1}{\lceil n/f \rceil - 1}\right]$	

Table 1. Summary of lower and upper bounds on contraction ratios for asymptotic consensus with scalar input and output values. New bounds proved in this work are marked with an *. The three leftmost columns are worst-case contraction ratios for the case the oblivious message adversary is (i) a general non-split, (ii) a non-split message adversary with $\alpha_{\mathcal{G}}$ -diameter D , and (iii) a general rooted message adversary. Contraction ratios are 0 if and only if exact consensus is solvable. The two rightmost columns summarize the bounds for the classical model of an asynchronous system with crashes.

Finally, we extend the above results on contraction ratios to derive new lower bounds on the decision times of any approximate consensus algorithm: Let $\Delta > 0$ be the largest distance between initial values. For $n = 2$ we obtain a lower bound of $\log_3 \frac{\Delta}{\epsilon}$ (Theorem 9.2). For $n \geq 3$ and message adversaries that include deaf(G) for a communication graph G , we show a lower bound of $\log_2 \frac{\Delta}{\epsilon}$ (Theorem 9.3), and for $n \geq 4$ and message adversaries that include certain Ψ graphs, we obtain a lower bound of $(n - 2) \log_2 \frac{\Delta}{\epsilon}$ (Theorem 9.4). For arbitrary oblivious message adversaries in which exact consensus is not solvable, we show a lower bound of $\log_{D+1} \frac{\Delta}{\epsilon n}$ (Theorem 9.5). Again, the algorithms by Charron-Bost et al. [18] have matching time complexities; showing optimality of these algorithms also for solving approximate consensus with scalar input and output values.

A preliminary version of this paper was presented at the conference PODC 2018 [33].

Related Work. The problem of asymptotic consensus in dynamic networks has been extensively studied in distributed computing and control theory (see, e.g., [4, 8, 10, 20, 27, 47]). The question of guaranteed convergence speeds and decision times of the corresponding approximate consensus problems, naturally arise in this context. Algorithms with convergence times exponential in the number of agents have been proposed. In particular, Cao et al. [10, Equation (26)] proved that the Equal Neighbor algorithm, which updates its value to the unweighted average of received values, has a contraction ratio of at most $\sqrt[n-1]{1 - 1/n^{n-1}}$ subject to rooted message adversaries, which leads to an exponential upper bound on the convergence time.

Olshevsky and Tsitsiklis [53], proposed an algorithm with polynomial convergence time in bidirectional networks with certain stability assumptions on the occurring communication graphs. The bounds on convergence times were later on refined by Nedic et al. [49]. Chazelle [20] proposed a convex-combination algorithm with polynomial convergence time, which works for any bidirectional connected message adversary.

To speed up convergence times, algorithms where agents set their output based on values that have been received in rounds prior to the previous round have also been considered in literature: Olshevsky [52] proposed a linear convergence time algorithm that uses messages from two rounds, restricted to a fixed bidirectional communication graph, however. Yuan et al. [58] proposed a linear convergence-time algorithm for a possibly non-bidirectional fixed topology. It requires storing all received values. Charron-Bost et al. [18] presented the midpoint algorithm, which has constant convergence time for non-split message adversaries and the amortized midpoint algorithm with linear convergence time subject to rooted message adversaries.

To the best of our knowledge, the only study of lower bounds in dynamic networks has been done by Cao et al. [11]: the authors identified $1 - 1/n$ as the worst-case scrambling constant of the Equal Neighbor algorithm in non-split communication graphs, and $1 - 1/n^{n-1}$ for that of the product of $n - 1$ rooted communication graphs. While scrambling constants can be used to prove upper bounds on the contraction ratio of asymptotic consensus algorithms, a lower bound on the scrambling constant does not in general imply a similar lower bound on the contraction ratio.

In the context of classical distributed computing failure scenarios, Dolev et al. [25] studied the approximate consensus problem: they considered fully-connected synchronous distributed systems with up to f Byzantine agents, and its asynchronous variant. The two presented algorithms require $n \geq 3f + 1$ for the synchronous and $n \geq 5f + 1$ for the asynchronous distributed system, the first of which is optimal in terms of resilience [30]. The latter result was improved to $n \geq 3f + 1$ by Abraham et al. [1]. Both papers also address the question of optimal contraction ratio in such systems. Since, however, in synchronous systems with $n \geq 3f + 1$ exact consensus is solvable, leading to a contraction ratio of 0, the authors consider bounds for round-by-round contraction ratios. Dolev et al. [25] showed that the achieved round-by-round contraction ratio of $\frac{1}{2}$ is actually tight for a certain class of algorithms that repeatedly set their output to the image of a so-called cautious function applied to the multiset of received values. A lower bound for arbitrary algorithms, however, remained an open problem. In higher dimensions, i.e., for any $d \geq 1$, Mendes et al. [45] proposed algorithms with decision time of $d \cdot \lceil \log_2 \frac{\sqrt{d}\Delta}{\epsilon} \rceil$ under the optimal resiliency condition $n \geq f \cdot (d + 2) + 1$. By using the MidExtremes algorithm, Függer and Nowak [32] showed that the problem is solvable within time $O(\log \frac{\Delta}{\epsilon})$. Central to this result is the property that the MidExtremes algorithm subject to non-split oblivious message adversaries guarantees a contraction of output values by a factor $c = \sqrt{7/8}$ every $k = 1$ rounds for arbitrary dimensions [32, Theorem 1].

Fekete [28] also studied round-by-round contraction ratios for several failure scenarios in which exact consensus is solvable. He proved asymptotically tight lower bounds for synchronous distributed systems in presence of crashes, omission, and Byzantine agents. The bounds hold for approximate consensus algorithms that potentially take into account information from all previous rounds. Fekete [29] later presented an algorithm for asynchronous message-passing systems with a minority of crashes, also proving a tight lower bound on the contraction ratio of any algorithm operating in asynchronous rounds for such systems. In the nonuniform iterated immediate snapshot (NIIS) model, Hoest and Shavit [39, Theorem 5.2] proved tight upper and lower bounds on the round complexity for the approximate consensus problem.

Herlihy et al. [38] studied generalizations of the approximate consensus problem in terms of combinatorial topology with round-based communication objects.

Determining the solvability of distributed problems, in particular decision problems, has a long tradition. One of the early general characterizations of solvability is by Biran et al. [7] who determined the class of decision problems solvable with one crash-faulty process. A much more general approach has been taken by the combinatorial-topology approach, which has characterized decision-problem solvability in a variety of models [37]. Sharp solvability characterizations are available, in particular in immediate-snapshot models [37, Theorem 5.2.7]. Coulouma et al. [22] characterized the oblivious message adversaries that allow solvability of exact consensus. A general characterization of message adversaries that allow exact consensus to be solved was recently established by Nowak et al. [51]. Other recent research efforts in this direction include the investigation of communication complexity [24] and of distributed network algorithms [12].

Paper Organization. The rest of the paper is organized as follows. Section 2 describes the system model. In Section 3, we formally define the notions of valency and contraction ratio and deduce some useful lemmas. We prove a tight lower bound on the contraction ratio for $n = 2$ agents in

Section 4. In Section 5, we give a tight lower bound for the non-split case with $n \geq 3$ agents, which we extend to an almost-tight lower bound in the rooted case in Section 6. We relate the problems of asymptotic and exact consensus in Section 7. In Section 8, we deduce an almost-tight lower bound for round-based algorithms in asynchronous message passing with crash faults. Section 9 translates our lower bounds for asymptotic consensus to tight and almost-tight time-complexity lower bounds for approximate consensus. We give concluding remarks in Section 10. Appendices A and B contain auxiliary lemmas and proofs.

2 SYSTEM MODEL

We consider a set $[n] = \{1, \dots, n\}$ of n agents. We assume a distributed, round-based computational model in the spirit of the Heard-Of model [19] and similar approaches [34, 40, 54, 55]. Computation proceeds in *rounds*: In every round, each agent sends a message to its outgoing neighbors, receives messages from its incoming neighbors, and finally updates its state according to a deterministic local algorithm, i.e., a transition function that maps the collection of incoming messages to a new state. Rounds are communication closed in the sense that no agent receives messages in round t that are sent in a round different from t . Communication starts at round $t = 1$; round $t = 0$ only holds the agents' initial states.

Communications that occur in a round are modeled by a *directed graph* with a node for each agent. An edge (i, j) is present in the communication graph of a round if agent i 's message was successfully received by agent j in that round. Since an agent can obviously communicate with itself instantaneously, every communication graph contains a self-loop at each node. In the following, we use the *product* of two communication graphs G and H , denoted $G \circ H$, which is the directed graph with an edge from i to j if there exists k such that (i, k) and (k, j) are edges in G and H , respectively.

We call an infinite sequence of communication graphs a *communication sequence*. In each communication sequence, the communication graph at round t is denoted by G_t , and $\text{In}_i(t) = \text{In}_i(G_t)$ and $\text{Out}_i(t) = \text{Out}_i(G_t)$ are the sets of incoming and outgoing neighbors (in-neighbors and out-neighbors for short) of agent i in G_t .

A *message adversary* is a set of communication sequences. An *oblivious* message adversary is defined by a set $\mathcal{G}$ of communication graphs; its communication sequences are all sequences in which all communication graphs are chosen in $\mathcal{G}$. We use the notation $\mathcal{M}(\mathcal{G})$ for the oblivious message adversary defined by the set $\mathcal{G}$ of communication graphs and we say that a communication graph G *occurs in* oblivious message adversary $\mathcal{M}(\mathcal{G})$ if $G \in \mathcal{G}$.

Let us fix an algorithm $\mathcal{A}$. A *configuration* is a collection of n agent states, one per agent. Since agents are deterministic, given some configuration C and some communication graph G , the algorithm $\mathcal{A}$ uniquely determines a new configuration, which we simply denote $C.G$ if no confusion can arise. The *execution* E of $\mathcal{A}$ starting from the initial configuration C_0 with the communication sequence $(G_t)_{t \geq 1}$ is the sequence $C_0, G_1, \dots, C_{t-1}, G_t, C_t, \dots$ of alternating configurations and communication graphs such that for each round t , we have $C_t = C_{t-1}.G_t$. We denote the set of executions with communication sequences in $\mathcal{M}$, starting from any initial configuration, by $\mathcal{E}_{\mathcal{M}, \mathcal{A}}$. We equip this set with the distance $\text{dist}(E, E') = 1/2^\theta$, where θ is the first index at which the configurations of E and E' differ. This is a compact metric space if the algorithm has a finite number of initial configurations: Compactness of the execution space for models described by safety properties [3] was shown for various special cases (e.g., [35, Lemma 5.1] or [51, remark after Lemma 4.9]). A proof for the case of oblivious message adversaries can be found in Lemma A.4 in the appendix.

Finally, any configuration that occurs in some execution with a communication sequence in $\mathcal{M}$ starting at initial configuration C_0 is said to be *reachable from C_0 by $\mathcal{A}$ subject to $\mathcal{M}$* . In the sequel, the algorithm and the message adversary are omitted if no confusion can arise.

2.1 Asymptotic Consensus

We assume that the local state of agent i includes a variable y^i in Euclidean d -space, and we let $y_E^i(t) \in \mathbb{R}^d$ denote the value of y^i at the end of round t in execution E . The value $y_E^i(0)$ is the initial value of agent i in execution E . We set $y_E(t) = (y_E^1(t), \dots, y_E^n(t))$. We write

$$\text{diam}(A) = \sup_{x, y \in A} \|x - y\|$$

for the diameter of $A \subseteq \mathbb{R}^d$ and $\Delta(y_E(t)) = \text{diam}\{y_E^1(t), \dots, y_E^n(t)\}$ for the diameter of the set of values at the end of round t .

We say an algorithm *solves the asymptotic consensus problem* subject to a message adversary $\mathcal{M}$ if the following holds for every execution E with a communication sequence in $\mathcal{M}$:

- *Convergence.* Each sequence $y_E^i(t)$ converges as $t \rightarrow \infty$.
- *Agreement.* If $y_E^i(t)$ and $y_E^j(t)$ converge, then their limits are equal.
- *Validity.* If $y_E^i(t)$ converges, then its limit lies in the convex hull of the set of initial values $y_E^1(0), \dots, y_E^n(0)$.

The Validity condition's requirement is motivated on one hand to rule out trivial solutions: ones that always output the same value, irrespective of the initial values. Also, staying in the range of initial values is required for applications such as sensor fusion or clock synchronization.

The *consensus function* defined by $y^*: \mathcal{E}_{\mathcal{M}, \mathcal{A}} \rightarrow \mathbb{R}^d$, where $y^*(E)$ is the common limit of the n sequences $(y_E^i(t))_{t \geq 0}$, is a priori not continuous when using the metric “dist”, which we defined earlier, on the set $\mathcal{E}$ of executions. Indeed, there exist asymptotic consensus algorithms whose consensus functions are not continuous — as shown in the next section.

2.2 Solvability of Asymptotic Consensus with Convex-Combination Algorithms

In a previous paper [17], Charron-Bost et al. proved the following characterization of oblivious message adversaries for which approximate consensus is solvable. With much of the same arguments, this can be extended to asymptotic consensus.

THEOREM 2.1 (ADAPTED FROM [17]). *The asymptotic consensus problem is solvable subject to an oblivious message adversary $\mathcal{M}$ if and only if each graph that occurs in $\mathcal{M}$ has a rooted spanning tree.*

For the proof of the sufficient condition, Charron-Bost et al. focused on convex-combination algorithms. In particular, they showed [17] that, while not all, a large class of convex-combination algorithms solve asymptotic consensus subject to rooted message adversaries. Such algorithms are memoryless, require little computational overhead and, more importantly, have the benefit of working in anonymous networks. Interestingly, their consensus function y^* is continuous, as we show in the next theorem.

Continuity of the consensus function of exact consensus is known (e.g., [50, Lemma 4.4]). We included a proof for message adversaries in Lemma A.5 in the appendix. The proof heavily relies on the existence of a decision value, i.e., the common decision value remains fixed from a finite time on. Asymptotic consensus, on the other hand, does not have such a finite time. A different argument is thus needed.

THEOREM 2.2. *The consensus function of every convex-combination algorithm that solves asymptotic consensus is continuous.*

Fig. 1. The communication graphs T_{bi} and T_{uni} for $n = 3$

PROOF. Let $(E_s)_{s \geq 0}$ be a sequence of executions that converges to E , i.e., $\text{dist}(E_s, E) \rightarrow 0$ as $s \rightarrow \infty$. By definition of the distance on the execution space, this means that

$$\forall t \geq 0 \exists s_t \forall s \geq s_t: y_s(0) = y(0), y_s(1) = y(1), \dots, y_s(t) = y(t) \quad (4)$$

where $y_s(t)$ and $y(t)$ denote the vectors $y_{E_s}(t)$ and $y_E(t)$ of the round- t output values of agents in execution E_s and E , respectively.

Let $\varepsilon > 0$. By definition of the limit y^* of execution E , there exists some t such that

$$\forall i \in [n]: \|y^i(t) - y^*\| \leq \varepsilon/3.$$

By (4), there is an s_t such that

$$\forall s \geq s_t \forall i \in [n]: \|y_s^i(t) - y^*\| \leq \varepsilon/3.$$

By the triangle inequality, this means

$$\forall s \geq s_t \forall i, j \in [n]: \|y_s^i(t) - y_s^j(t)\| \leq 2\varepsilon/3.$$

Because the algorithm is a convex-combination algorithm, the limit y_s^* lies in the convex hull of the points $y_s^1(t), \dots, y_s^n(t)$. That is, denoting by y_s^* the common limit in execution E_s , we have

$$\forall s \geq s_t \forall i \in [n]: \|y_s^i(t) - y_s^*\| \leq 2\varepsilon/3.$$

Combining these inequalities gives

$$\forall s \geq s_t: \|y_s^* - y^*\| \leq \|y_s^* - y_s^i(t)\| + \|y_s^i(t) - y^*\| \leq \frac{2\varepsilon}{3} + \frac{\varepsilon}{3} = \varepsilon$$

where i is any agent. This proves $\lim_{s \rightarrow \infty} y_s^* = y^*$ as required. $\square$

As mentioned previously, there exist asymptotic consensus algorithms with non-continuous consensus functions:

Example 2.3. Take the oblivious message adversary $\mathcal{M} = \mathcal{M}(\{T_{bi}, T_{uni}\})$ with $n = 3$ agents where communication graph T_{bi} is the bi-directional triangle and T_{uni} is a uni-directional cyclic triangle as depicted in Figure 1. Algorithm $\mathcal{A}$ relays all initial values and sets its output value to the minimal initial value if all of the communication graphs up to the current round were bi-directional, and to the maximal initial value else.

Since all communication graphs in $\mathcal{M}$ are strongly connected, all agents know the set of initial values at the end of the second round. In particular, algorithm $\mathcal{A}$ solves asymptotic consensus subject to the oblivious message adversary $\mathcal{M}$. Now consider the executions $E^{(t)}$ for every $t \geq 2$ where agents start with values 0, 1, and 2, and where all communication graphs are equal to T_{bi} except in round t in which the communication graph is T_{uni} . Since every agent can locally tell the difference between T_{bi} and T_{uni} , all agents output value 2 from round t on. In the limit execution $E^{(\infty)}$, however, all communication graphs are equal to T_{bi} without exception and hence all agents output value 0 from round 2 on. We thus have $E^{(t)} \rightarrow E^{(\infty)}$ but $y^*(E^{(t)}) \not\rightarrow y^*(E^{(\infty)})$ as $t \rightarrow \infty$. Algorithm $\mathcal{A}$'s

Fig. 2. The contraction of the valency along an execution $E = C_0, G_1, C_1, G_2, \dots$ with an average contraction of $\sqrt{0.8 \cdot 0.75} = \sqrt{0.6} \approx 0.774$ per round for the first two rounds.

consensus function is hence not continuous. The example easily generalizes to arbitrary $n \geq 3$ by considering bidirectional and unidirectional ring graphs. $\square$

3 VALENCY AND CONTRACTION RATIO

We now extend the notion of valency, which has been developed for exact consensus [31, 48], to asymptotic consensus. We fix an algorithm $\mathcal{A}$ that solves asymptotic consensus subject to a certain oblivious message adversary $\mathcal{M}$ with $n \geq 2$ agents. Let C be a configuration reachable by $\mathcal{A}$ subject to $\mathcal{M}$. Then we define the *valency* of C by

$$Y_{\mathcal{M}, \mathcal{A}}^*(C) = \{y^*(E) \in \mathbb{R}^d \mid C \text{ occurs in } E \in \mathcal{E}_{\mathcal{M}, \mathcal{A}}\}.$$

Observe that if $\mathcal{A}$ is a convex-combination algorithm, then the valency of a configuration C is a compact set of $\mathbb{R}^d$ since the consensus function is continuous and the set of executions in which C occurs is a compact set. Set $\delta_{\mathcal{M}, \mathcal{A}}(C) = \text{diam}(Y_{\mathcal{M}, \mathcal{A}}^*(C))$ the diameter of the set of reachable limits whose executions include configuration C . We have $\delta_{\mathcal{M}, \mathcal{A}}(C_t) \rightarrow 0$ in any execution $E = C_0, G_1, C_1, G_2, \dots$ by Convergence and Agreement.

Figure 2 depicts the evolution of the valency along an execution E . In the shown example valencies are intervals. However, in general, this is not the case.

To study the speed of convergence of an algorithm $\mathcal{A}$ subject a message adversary $\mathcal{M}$, bounding the evolution of $\delta_{\mathcal{M}, \mathcal{A}}(C_t)$ along any execution $E = C_0, G_1, C_1, G_2, \dots$ is a natural choice. While we prove lower bounds on $\delta_{\mathcal{M}, \mathcal{A}}(C_t)$ for arbitrary $t \geq 0$ in this work, we also introduce the *contraction ratio* of algorithm $\mathcal{A}$ subject to message adversary $\mathcal{M}$ as

$$\sup_{E \in \mathcal{E}_{\mathcal{M}, \mathcal{A}}} \limsup_{t \rightarrow \infty} \sqrt[t]{\delta_{\mathcal{M}, \mathcal{A}}(C_t)}$$

where we write $E = C_0, G_1, C_1, G_2, \dots$. Observe that the contraction ratio, by definition, is between 0 and 1. If exact consensus is solvable subject to $\mathcal{M}$, the optimal contraction ratio 0 can be achieved.

For convex-combination algorithms we obtain the following:

LEMMA 3.1. *Let C be a configuration of convex-combination algorithm $\mathcal{A}$ that solves asymptotic consensus subject to message adversary $\mathcal{M}$, and denote by $\Delta(C)$ the diameter of the set of outputs in C . Then $\delta_{\mathcal{M}, \mathcal{A}}(C) \leq \Delta(C)$.*

In particular, for a convex-combination algorithm that guarantees contraction of output values by a factor $c < 1$ every $k \geq 1$ rounds, i.e., fulfills (1), the contraction ratio is less or equal to $\sqrt[k]{c}$.

PROOF. Denote by $Y(C)$ the set of output values in configuration C , and by $\text{hull}(X)$ the convex hull of a set X .

Fig. 3. The rooted communication graphs H_0 , H_1 , and H_2 for $n = 2$.

The first statement follows from the fact that convex-combination algorithms by definition set their outputs within the convex hull of previously received values. Any limit that is reachable in an execution that contains configuration C by algorithm $\mathcal{A}$ subject to $\mathcal{M}$ is thus necessarily within $\text{hull}(Y(C))$. The statement now follows from $\Delta(C) = \text{diam}(Y(C)) = \text{diam}(\text{hull}(Y(C))) \geq \text{diam}(Y_{\mathcal{M}, \mathcal{A}}^*(C)) = \delta_{\mathcal{M}, \mathcal{A}}(C)$.

The second statement follows from the first statement and the definition of the contraction ratio. $\square$

This allows us to directly compare different message adversaries and algorithms with each other. The measure is motivated by the observation that a large class of algorithms lead to a round-wise contraction of $\delta_{\mathcal{M}, \mathcal{A}}(C_t)$ by a ratio $\gamma < 1$: they guarantee that $\delta_{\mathcal{M}, \mathcal{A}}(C_t) \leq \gamma^t \delta_{\mathcal{M}, \mathcal{A}}(C_0)$ for all $t \geq 0$. Any such algorithm has a contraction ratio of at most γ . In general, however, a round-wise contraction ratio may be too conservative to assume: algorithms may contract quickly in general, but rarely have bad round-wise contraction or no contraction at all. The contraction ratio accounts for this by accumulating over rounds; see, e.g., Figure 2 where the average contraction factor over two rounds is less than 0.774 while the first round contracts by 0.8.

Example 3.2. Consider the oblivious message adversary $\mathcal{M} = \mathcal{M}(\{T_{\text{bi}}, T_{\text{uni}}\})$ and the algorithm $\mathcal{A}$ from Example 2.3. Also let executions $E^{(t)}$, for $t \geq 2$, and $E^{(\infty)} = C_0, G_1, C_1, G_2, \dots$ be as defined in that example. Then, for all rounds $t \geq 1$:

- $0 \in Y_{\mathcal{M}, \mathcal{A}}^*(C_t)$ since algorithm $\mathcal{A}$ results in all agent outputs eventually being 0 in execution $E^{(\infty)}$, i.e., $y^*(E^{(\infty)}) = 0$
- $2 \in Y_{\mathcal{M}, \mathcal{A}}^*(C_t)$ since algorithm $\mathcal{A}$ results in all agent outputs eventually being 2 in execution $E^{(t+1)}$, i.e., $y^*(E^{(t+1)}) = 2$

It follows that, for all rounds $t \geq 1$, the valency $Y_{\mathcal{M}, \mathcal{A}}^*(C_t)$ of configuration C_t at round t contains 0 and 2. In fact, $Y_{\mathcal{M}, \mathcal{A}}^*(C_t) = \{0, 2\}$. Thus, $\delta_{\mathcal{M}, \mathcal{A}}(C_t) = \delta_{\mathcal{M}, \mathcal{A}}(C_0)$ for all $t \geq 1$; the valency does not shrink with increasing number of rounds and the contraction ratio of algorithm $\mathcal{A}$ subject to oblivious message adversary $\mathcal{M}$ is 1. $\square$

Example 3.2 shows that there exist algorithms for which the valency does not shrink at all and which have a worst-case contraction ratio of 1. The following example demonstrates that there exist algorithms with contraction factor less than 1, though.

Example 3.3. Consider the midpoint algorithm for $n = 2$ and the oblivious message adversary $\mathcal{M} = \mathcal{M}(\{H_0, H_1, H_2\})$ with the communication graphs H_0 , H_1 , and H_2 shown in Figure 3. The midpoint algorithm for two agents takes the equally weighted average of its own and the remote agent, if it receives a message from the remote agent, and does not update its value otherwise. Assume that agent 1 starts with initial value 0 and agent 2 with value 1.

Consider an executions $E = C_0, G_1, C_1, G_2, \dots$. First observe that both initial values are in $Y_{\mathcal{M}, \mathcal{A}}^*(C_0)$: for the execution with only graphs H_1 the agent outputs converge to 0, and for the execution with only graphs H_2 , the outputs converge to 1. Thus $\delta_{\mathcal{M}, \mathcal{A}}(C_0) = 1$.

Distinguishing between the three possible cases for communication graph G_{t+1} , we obtain:

- (1) If $G_{t+1} = H_0$, then both $y^1(t+1) = y^2(t+1) = \frac{y^1(t)+y^2(t)}{2}$, and thus $Y_{\mathcal{M}, \mathcal{A}}^*(C_{t+1}) = \left\{ \frac{y^1(t)+y^2(t)}{2} \right\}$. Consequently $\delta_{\mathcal{M}, \mathcal{A}}(C_{t+1}) = 0$.
- (2) If $G_{t+1} = H_1$, then $y^1(t+1) = y^1(t)$ and $y^2(t+1) = \frac{y^1(t)+y^2(t)}{2}$. By similar arguments as for C_0 , we have $\left\{ y^1(t), \frac{y^1(t)+y^2(t)}{2} \right\} \subseteq Y_{\mathcal{M}, \mathcal{A}}^*(C_1)$. One also observes that output values less than $y^1(t+1)$ and greater than $y^2(t+1)$ are not reachable by the midpoint algorithm. Consequently $\delta_{\mathcal{M}, \mathcal{A}}(C_{t+1}) = \frac{1}{2} \cdot \delta_{\mathcal{M}, \mathcal{A}}(C_t)$.
- (3) For $G_{t+1} = H_2$, by symmetry to the previous case, $\delta_{\mathcal{M}, \mathcal{A}}(C_{t+1}) = \frac{1}{2} \cdot \delta_{\mathcal{M}, \mathcal{A}}(C_t)$.

Consequently, in any execution the valency shrinks by a factor of at least $\frac{1}{2}$ per round; the contraction ratio of the midpoint algorithm subject to oblivious message adversary $\mathcal{M}$ is thus $\frac{1}{2}$. $\square$

Motivated by the two previous examples, we begin a general analysis of valencies along executions with a lemma on subsets of message adversaries:

LEMMA 3.4. *Let $\mathcal{M}, \mathcal{M}'$ be two message adversaries with $\mathcal{M}' \subseteq \mathcal{M}$. If $\mathcal{A}$ is an algorithm that solves asymptotic consensus subject to $\mathcal{M}$, then (i) it also solves asymptotic consensus subject to $\mathcal{M}'$, (ii) for every configuration C reachable by $\mathcal{A}$ in $\mathcal{M}'$, we have $Y_{\mathcal{M}', \mathcal{A}}^*(C) \subseteq Y_{\mathcal{M}, \mathcal{A}}^*(C)$, (iii) $\delta_{\mathcal{M}', \mathcal{A}}(C) \leq \delta_{\mathcal{M}, \mathcal{A}}(C)$, and (iv) the contraction ratio subject to $\mathcal{M}'$ is less or equal to the contraction ratio subject to $\mathcal{M}$.*

PROOF. Statements (i), (ii), and (iii) immediately follow from the definition of valency. It remains to show statement (iv). From $\mathcal{E}_{\mathcal{M}', \mathcal{A}} \subseteq \mathcal{E}_{\mathcal{M}, \mathcal{A}}$ and (iii), we deduce

$$\sup_{E \in \mathcal{E}_{\mathcal{M}', \mathcal{A}}} \limsup_{t \rightarrow \infty} \sqrt[t]{\delta_{\mathcal{M}', \mathcal{A}}(C_t)} \leq \sup_{E \in \mathcal{E}_{\mathcal{M}, \mathcal{A}}} \limsup_{t \rightarrow \infty} \sqrt[t]{\delta_{\mathcal{M}, \mathcal{A}}(C_t)},$$

which concludes the proof. $\square$

We next establish two branching properties of valency of configurations in execution trees for oblivious message adversaries.

LEMMA 3.5. *Let C be a configuration reachable by algorithm $\mathcal{A}$ subject to oblivious message adversary $\mathcal{M} = \mathcal{M}(\mathcal{G})$. Then*

$$Y_{\mathcal{M}, \mathcal{A}}^*(C) = \bigcup_{G \in \mathcal{G}} Y_{\mathcal{M}, \mathcal{A}}^*(C.G).$$

PROOF. First let $y^* \in Y_{\mathcal{M}, \mathcal{A}}^*(C)$. By definition of the valency $Y_{\mathcal{M}, \mathcal{A}}^*(C)$, there exists an execution $E = C_0, G_1, C_1, G_2, \dots$ in $\mathcal{E}_{\mathcal{M}, \mathcal{A}}$ and a $t \geq 0$ such that $y^* = y^*(E)$ and $C = C_t$. Set $G = G_{t+1}$. Hence we have $C_{t+1} = C.G$. But this shows that $y^* \in Y_{\mathcal{M}, \mathcal{A}}^*(C.G)$ since $C.G$ occurs in execution E whose limit is y^* . This shows inclusion of the left-hand side in the right-hand side.

Now let $G \in \mathcal{G}$ and $y^* \in Y_{\mathcal{M}, \mathcal{A}}^*(C.G)$. Then there is an execution $E = C_0, G_1, C_1, G_2, \dots$ in $\mathcal{E}_{\mathcal{M}, \mathcal{A}}$ and a $t \geq 1$ such that $y^* = y^*(E)$ and $C.G = C_t$. Since C is a reachable configuration, there exists an execution $E' = C'_0, G'_1, C'_1, G'_2, \dots$ in $\mathcal{E}_{\mathcal{M}, \mathcal{A}}$ and an $s \geq 0$ such that $C'_s = C$. Then the sequence

$$E'' = C'_0, G'_1, \dots, C'_s, G, C_t, G_{t+1}, \dots$$

is an execution in $\mathcal{E}_{\mathcal{M}, \mathcal{A}}$ with $y^*(E'') = y^*(E) = y^*$. Hence $y^* \in Y_{\mathcal{M}, \mathcal{A}}^*(C)$ because C occurs in E'' . This shows inclusion of the right-hand side in the left-hand side and concludes the proof. $\square$

LEMMA 3.6. *Let C be a configuration reachable by algorithm $\mathcal{A}$ subject to oblivious message adversary $\mathcal{M} = \mathcal{M}(\mathcal{G})$. Then there exist $G, H \in \mathcal{G}$ such that*

$$\delta_{\mathcal{M}, \mathcal{A}}(C) = \text{diam} (Y_{\mathcal{M}, \mathcal{A}}^*(C.G) \cup Y_{\mathcal{M}, \mathcal{A}}^*(C.H)).$$

PROOF. By Lemma 3.5 it is $Y_{\mathcal{M}, \mathcal{A}}^*(C) = \bigcup_{G \in \mathcal{G}} Y_{\mathcal{M}, \mathcal{A}}^*(C.G)$. The claimed equality now follows from Lemma A.1. $\square$

Two configurations C and C' are called *indistinguishable for agent i* , denoted $C \sim_i C'$, if i is in the same state in C and in C' . As an immediate consequence of this definition, since every agent is an in-neighbor of itself, we obtain:

LEMMA 3.7. *Let C and C' be two reachable configurations and let G and G' be communication graphs that occur in an oblivious message adversary $\mathcal{M}$. If some agent i has the same in-neighbors in G and G' and if $C \sim_j C'$ for each of i 's in-neighbors j , then $C.G \sim_i C'.G'$.*

An agent i is said to be *deaf in a communication graph G* if i has only one in-neighbor in G , namely i itself. We can relate valencies of successor configurations.

LEMMA 3.8. *Let G and G' be two communication graphs that occur in the oblivious message adversary $\mathcal{M}$, and let i be an agent that has the same in-neighbors in G and G' . Further let C and C' be two configurations such that $C \sim_j C'$ for all in-neighbors j of i in G and G' . Then, if there exists a communication graph that occurs in $\mathcal{M}$ in which i is deaf, we have $Y_{\mathcal{M}, \mathcal{A}}^*(C.G) \cap Y_{\mathcal{M}, \mathcal{A}}^*(C'.G') \neq \emptyset$.*

PROOF. From Lemma 3.7, we have $C.G \sim_i C'.G'$.

Let D_i be a communication graph that occurs in $\mathcal{M}$ in which the agent i is deaf. Then we consider an execution E in which C occurs at some round $t_0 - 1$, G is the communication graph at round t_0 , and from there on all communication graphs are equal to D_i . Analogously, let E' be an execution identical to E except that the communication graph at round t_0 is G' instead of G . By inductive application of Lemma 3.7, we show that for all $t \geq t_0$, we have $C_t \sim_i C'_t$. In particular, we obtain $y_E^i(t) = y_{E'}^i(t)$. Thus $y^*(E) = y^*(E')$, which shows that $Y_{\mathcal{M}, \mathcal{A}}^*(C.G)$ and $Y_{\mathcal{M}, \mathcal{A}}^*(C'.G')$ intersect. $\square$

From Lemma 3.8 we determine the valency of any initial configuration when certain communication graphs occur in the oblivious message adversary. If every agent is deaf in some communication graph that occurs in the oblivious message adversary $\mathcal{M}$, then the next lemma shows that the diameter of the valency of any initial configuration is equal to the diameter of the set of its initial values. The proof is similar to classical proofs for the existence of bivalent initial configurations for exact consensus [31, Lemma 2] [48, Theorem 5.8]; except that it adds a quantitative perspective on the initial degree of non-univalence.

LEMMA 3.9. *If, for every agent i , there is a communication graph that occurs in the oblivious message adversary $\mathcal{M}$ in which i is deaf, then each initial configuration C_0 satisfies $\delta_{\mathcal{M}, \mathcal{A}}(C_0) = \Delta(C_0)$. In particular, there is an initial configuration for which $\delta_{\mathcal{M}, \mathcal{A}}(C_0) > 0$.*

PROOF. Since $Y_{\mathcal{M}, \mathcal{A}}^*(C_0)$ is a subset of the convex hull of the set of points $y^1(0), \dots, y^n(0)$ by the Validity property of asymptotic consensus and since the diameter of the convex hull of the set $\{y^1(0), \dots, y^n(0)\}$ is equal to $\Delta(C_0)$, we have the inequality $\delta_{\mathcal{M}, \mathcal{A}}(C_0) \leq \Delta(C_0)$.

To show the converse inequality, let i and j be two agents such that $\|y^i(0) - y^j(0)\| = \Delta(C_0)$. Let E be the execution with initial configuration C_0 and a constant communication graph in which agent i is deaf. Now consider $C_0^{(i)}$, an initial configuration such that all initial values are set to $y^i(0)$, and the execution $E^{(i)}$ from $C_0^{(i)}$ with the same communication sequence as in E .

By a repeated application of Lemma 3.7, we see that at each round t , we have $C_t \sim_i C_t^{(i)}$. Hence, $y^*(E) = y^*(E^{(i)})$.

From the Validity condition, we deduce that an agent that is deaf during the whole execution has to converge to its own initial value, that is, $y^*(E^{(i)}) = y^i(0)$. It then follows that $y^i(0) \in Y_{\mathcal{M}, \mathcal{A}}^*(C_0)$.

By a similar argument, we see $y^j(0) \in Y_{\mathcal{M}, \mathcal{A}}^*(C_0)$. Hence

$$\delta_{\mathcal{M}, \mathcal{A}}(C_0) \geq \|y^i(0) - y^j(0)\| = \Delta(C_0) ,$$

which concludes the proof of the first part of the lemma.

The second part follows from the first by picking any initial configuration whose initial values are not all equal, that is, $\Delta(C_0) > 0$. $\square$

4 TIGHT BOUND FOR TWO AGENTS

In this section, we prove a lower bound of $1/3$ on the contraction ratio of algorithms that solve asymptotic consensus subject to the oblivious message adversary of all rooted (and here also non-split) communication graphs with two agents. Combined with Algorithm 1, which achieves this lower bound [18], we have indeed identified a tight bound on the contraction ratio for $n = 2$. Moreover, the algorithm also shows that the lower bound is achieved by a simple convex-combination algorithm. In the NIIS model, an analogous bound for $n = 2$ was proved by Hoest and Shavit [39].

ALGORITHM 1: Algorithm with contraction ratio $1/3$ for $n = 2$

Initially:

| $y^i \in \mathbb{R}^d$;

In round $t \geq 1$ do

| send y^i to other agent;

| **if y^j was received from other agent then**

| | $y^i \leftarrow y^i/3 + 2y^j/3$;

| **end**

end

For $n = 2$, there are three possible rooted communication graphs that may occur, all of which are non-split; see Figure 3: (i) H_0 in which all messages are received, (ii) H_1 in which agent 2 receives agent 1's message but not vice versa, and (iii) H_2 in which agent 1 receives agent 2's message but not vice versa.

A straightforward analysis of Algorithm 1 shows that its contraction ratio is equal to $1/3$: Indeed, in all three cases, H_0 , H_1 , and H_2 for the communication graph G_t in round t , we have $\Delta(t) \leq \Delta(t-1)/3$. That is, we can apply Lemma 3.1 with $c = 1/3$ and $k = 1$.

THEOREM 4.1. *The contraction ratio of any asymptotic consensus algorithm for $n = 2$ agents subject to the oblivious message adversary with the three communication graphs H_0 , H_1 , and H_2 is greater or equal to $1/3$.*

PROOF. We show the stronger statement that for every initial configuration C_0 there is an execution $E = C_0, G_1, C_1, G_2, \dots$ starting from C_0 such that

$$\delta_{\mathcal{M}, \mathcal{A}}(C_t) \geq \frac{1}{3^t} \cdot \delta_{\mathcal{M}, \mathcal{A}}(C_0) \tag{5}$$

for $t \geq 0$. This, applied to an initial configuration with $\delta_{\mathcal{M}, \mathcal{A}}(C_0) > 0$, which exists by Lemma 3.9, then shows the theorem.

The proof is by inductive construction of an execution $E = C_0, G_1, C_1, G_2, \dots$ whose configurations C_t satisfy (5). Equation (5) is trivial for $t = 0$.

Now assume $t \geq 0$ and that Equation (5) holds for t . There are three possible successor configurations of C_t , one for each of the communication graphs H_0 , H_1 , and H_2 in $\mathcal{M}$. Set $C_{t+1}^k = C_t.H_k$. Further let $Y = Y_{\mathcal{M}, \mathcal{A}}^*(C_t)$, and $Y_k = Y_{\mathcal{M}, \mathcal{A}}^*(C_{t+1}^k)$.

We will show that there is some $\hat{k} \in \{0, 1, 2\}$ with $\text{diam}(Y_{\hat{k}}) \geq \text{diam}(Y)/3$. We then define $G_{t+1} = H_{\hat{k}}$ and $C_{t+1} = C_{t+1}^{\hat{k}}$. By the induction hypothesis, we then have

$$\delta_{\mathcal{M}, \mathcal{A}}(C_{t+1}) \geq \delta_{\mathcal{M}, \mathcal{A}}(C_t)/3 \geq \delta_{\mathcal{M}, \mathcal{A}}(C_0)/3^{t+1} ,$$

i.e., Equation (5) holds for $t + 1$.

Assume by contradiction that $\text{diam}(Y_k) < \text{diam}(Y)/3$ for all $k \in \{0, 1, 2\}$. From Lemma 3.5 we have $Y = Y_0 \cup Y_1 \cup Y_2$. Noting that agent 1 is deaf in H_1 and agent 2 has the same incoming edges as in H_0 , and that agent 2 is deaf in H_2 and agent 1 has the same incoming edges as in H_0 , we obtain from Lemma 3.8 that

$$Y_0 \cap Y_1 \neq \emptyset \quad \text{and} \quad Y_0 \cap Y_2 \neq \emptyset .$$

The sets Y_0 and Y_1 intersecting means

$$\text{diam}(Y_0 \cup Y_1) \leq \text{diam}(Y_0) + \text{diam}(Y_1) < \frac{2}{3} \text{diam}(Y) .$$

Further, the sets $Y_0 \cup Y_1$ and Y_2 intersecting means

$$\begin{aligned} \text{diam}(Y) &= \text{diam}(Y_0 \cup Y_1 \cup Y_2) \\ &\leq \text{diam}(Y_0 \cup Y_1) + \text{diam}(Y_2) < \text{diam}(Y) , \end{aligned}$$

a contradiction. This concludes the proof. $\square$

5 TIGHT BOUND FOR NON-SPLIT MESSAGE ADVERSARIES: CONTRACTION IN PRESENCE OF DEAF GRAPHS

In this section, we prove a lower bound of $1/2$ on the contraction ratio of asymptotic consensus algorithms for $n \geq 3$ agents subject to an oblivious message adversary in which graphs derived from a communication graph G occur, in which agents are made deaf. As a special case, this includes the oblivious message adversary of all non-split communication graphs. Formally, a communication graph G is *non-split* if for all nodes i and j , there exists a node k such that k is a common incoming neighbor, i.e., both (k, i) and (k, j) are edges of G . Charron-Bost et al. [18] presented the midpoint algorithm (given in Algorithm 2) for dimension one with contraction ratio $1/2$ for non-split communication graphs. Together, this shows tightness of our lower bound in dimension one.

ALGORITHM 2: Midpoint algorithm

Initially:

$y^i \in \mathbb{R};$

In round $t \geq 1$ do

 send y^i to all agents;

$m^i \leftarrow \min \{y^j \mid j \in \text{In}_i(t)\};$

$M^i \leftarrow \max \{y^j \mid j \in \text{In}_i(t)\};$

$y^i \leftarrow (m^i + M^i)/2;$

end

Let G be an arbitrary communication graph. Consider a system with $n \geq 3$ agents, and the n communication graphs $F_1, \dots, F_n$ where F_i is obtained by making i deaf in G , i.e., by removing all the edges towards i except the self-loop (i, i) .

With a proof similar to that of Theorem 4.1 but noting that the valencies of all pairs of successor configurations intersect, we get:

THEOREM 5.1. *The contraction ratio of any asymptotic consensus algorithm for $n \geq 3$ agents subject to an oblivious message adversary that includes $\text{deaf}(G)$ is greater or equal to $1/2$.*

PROOF. We show the stronger statement that for every initial configuration C_0 there is an execution $E = C_0, G_1, C_1, G_2, \dots$ starting at C_0 such that

$$\delta_{\mathcal{M}, \mathcal{A}}(C_t) \geq \frac{1}{2^t} \delta_{\mathcal{M}, \mathcal{A}}(C_0) \quad (6)$$

for all $t \geq 0$.

It suffices to show (6) for the specific oblivious message adversary $\mathcal{M}' = \mathcal{M}(\text{deaf}(G))$ because $\delta_{\mathcal{M}, \mathcal{A}}(C_t) \geq \delta_{\mathcal{M}', \mathcal{A}}(C_t)$ by Lemma 3.4 and $\delta_{\mathcal{M}', \mathcal{A}}(C_0) = \delta_{\mathcal{M}, \mathcal{A}}(C_0)$ by Lemma 3.9 whenever $\mathcal{M} \supseteq \mathcal{M}'$. We hence suppose $\mathcal{M} = \mathcal{M}'$ in the rest of the proof.

The proof is by inductive construction of an execution $E = C_0, G_1, C_1, G_2, \dots$ whose configurations C_t satisfy (6). This, applied to an initial configuration with $\delta_{\mathcal{M}, \mathcal{A}}(C_0) > 0$, which exists by Lemma 3.9, then shows the theorem. For $t = 0$ the inequality holds trivially.

Now assume $t \geq 0$ and that Equation (6) holds for t . There are n possible successor configurations based on the applicable communication graphs $F_1, \dots, F_n$. We denote $C_{t+1}^k = C_t \cdot F_k$, for any agent k . Further let $Y = Y_{\mathcal{M}, \mathcal{A}}^*(C_t)$, and $Y_k = Y_{\mathcal{M}}^*(C_{t+1}^k)$.

We will show that there exists some agent $\hat{k} \in [n]$ such that

$$\text{diam}(Y_{\hat{k}}) \geq \text{diam}(Y)/2 \quad (7)$$

We then define $G_{t+1} = F_{\hat{k}}$ and $C_{t+1} = C_{t+1}^{\hat{k}}$. By (7) and the induction hypothesis, we have

$$\delta_{\mathcal{M}, \mathcal{A}}(C_{t+1}) \geq \delta_{\mathcal{M}, \mathcal{A}}(C_t)/2 \geq \delta_{\mathcal{M}, \mathcal{A}}(C_0)/2^{t+1} \quad (8)$$

i.e., Equation (6) holds for $t + 1$.

Assume by contradiction that $\text{diam}(Y_k) < \text{diam}(Y)/2$ for all $k \in [n]$. Recall that agent i is deaf in F_i and has the same in-neighbors in all the communication graphs F_j with $j \neq i$. Since $n \geq 3$, for any pair of agents i, j we may select an agent ℓ different from i and j such that ℓ has the same in-neighbors in F_i as in F_j . Lemma 3.8 with the assumption that F_ℓ is in $\mathcal{M}$ shows that for any pair of agents i, j , we have

$$Y_i \cap Y_j \neq \emptyset \quad (9)$$

By Lemma 3.6, there exist $k, k' \in [n]$ such that $\text{diam}(Y_k \cup Y_{k'}) = \text{diam}(Y)$. In particular, we can choose $i = k$ and $j = k'$, which implies that

$$\begin{aligned} \text{diam}(Y) &= \text{diam}(Y_k \cup Y_{k'}) \leq \text{diam}(Y_k) + \text{diam}(Y_{k'}) \\ &< \text{diam}(Y) \end{aligned} \quad (10)$$

which is a contradiction and concludes the proof. $\square$

The oblivious message adversary $\mathcal{M}(\text{deaf}(K_n))$, where K_n is the complete digraph on n nodes, is a subset of the oblivious message adversary in which all non-split communication graphs occur. Hence the lower bound holds and, since Algorithm 2 is applicable and achieves a contraction ratio of $1/2$, a tight bound for one-dimensional values follows.

In fact, it would even be sufficient to reduce $\text{deaf}(G)$ to the graphs F_i, F_j, F_ℓ for three agents $i, j, \ell \in [n]$. With the same proof as Theorem 5.1 we get:

Fig. 4. Rooted communication graph Ψ_i for $n = 6$. Self-loops are omitted.

THEOREM 5.2. *The contraction ratio of any asymptotic consensus algorithm subject to an oblivious message adversary that includes three different communication graphs from $\text{deaf}(G)$ is greater or equal to $1/2$.*

The number of three different communication graphs from $\text{deaf}(G)$ is best possible. In fact, the oblivious message adversary with only two communication graphs F_i and F_j from $\text{deaf}(K_n)$ allows to solve exact consensus, which means that a contraction ratio of 0 for asymptotic consensus is achievable.

6 ALMOST-TIGHT BOUND FOR ROOTED MESSAGE ADVERSARIES: CONTRACTION IN PRESENCE OF Ψ GRAPHS

We next prove a lower bound of $\sqrt[n-2]{1/2}$ on the contraction ratio of asymptotic consensus algorithms for $n \geq 4$ agents.

For $i \in \{1, 2, 3\}$, let Ψ_i (see Figure 4) be the communication graph where agents $4 \leq j \leq n-1$ form a path with edges from j to $j+1$, agents $\{1, 2, 3\} \setminus i$ have n as their in-neighbor and 4 as their out-neighbor, and i has 4 as its out-neighbor. For $i \in \{1, 2, 3\}$, let σ_i be the finite communication sequence of $n-2$ repetitions of graph Ψ_i .

First observe that any communication sequence arising from the concatenation of σ_i sequences necessarily is a communication sequence of the oblivious message adversary that contains the graphs Ψ_1 , Ψ_2 , and Ψ_3 , which are rooted. However, the analysis of the set of these communication sequences necessitates the use of non-oblivious message adversaries. By such an analysis we obtain:

THEOREM 6.1. *The contraction ratio of any asymptotic consensus algorithm in an oblivious message adversary that includes the graphs Ψ_1 , Ψ_2 , and Ψ_3 is greater or equal to $\sqrt[n-2]{1/2}$.*

Since the amortized midpoint algorithm guarantees a contraction of $\sqrt[n-1]{1/2}$ for rooted message adversaries [18], Theorem 6.1 shows that this is asymptotically optimal for oblivious rooted message adversaries.

The proof of Theorem 6.1 is similar to that of Theorems 4.1 and 5.1. One difference is that we construct the lower-bound execution by inductively extending it by $n-2$ communication graphs (that is, some σ_i) at a time instead of a single communication graph. The key indistinguishability argument is that the third agent among 1, 2, and 3 does not see a difference between appending σ_i and appending σ_j .

6.1 Non-Oblivious Message Adversaries

To prove Theorem 6.1, we start by generalizing some of the basic lemmas we proved for the specific case of oblivious message adversaries. Towards that purpose, we denote for two finite sequences u and v , their concatenation by $u \cdot v$. We say a configuration C is reachable by algorithm $\mathcal{A}$ subject to message adversary $\mathcal{M}$ via the finite communication sequence $\pi = G_1, G_2, \dots, G_t$, with $t \geq 0$,

if π is the prefix of a communication sequence in $\mathcal{M}$ and $E = C_0, G_1, C_1, G_2, \dots, C_{t-1}, G_t, C, \dots \in \mathcal{E}_{\mathcal{M}, \mathcal{A}}$. For a finite prefix π of a communication sequence in $\mathcal{M}$, we define $\Sigma(\pi)$ to be the set of communication graphs G such that $\pi \cdot G$ is also a prefix of a communication sequence in $\mathcal{M}$.

The following lemmas are generalizations of Lemmas 3.5, 3.6, 3.8, and 3.9. Their proofs are in large parts analogous to the proof of the versions for oblivious message adversaries; we state them for completeness in the appendix.

LEMMA 6.2. *Let C be a configuration reachable by algorithm $\mathcal{A}$ subject to message adversary $\mathcal{M}$ via the finite communication sequence π . Then*

$$Y_{\mathcal{M}, \mathcal{A}}^*(C) = \bigcup_{G \in \Sigma(\pi)} Y_{\mathcal{M}, \mathcal{A}}^*(C.G) .$$

LEMMA 6.3. *Let C be a configuration reachable by algorithm $\mathcal{A}$ subject to message adversary $\mathcal{M}$ via the finite communication sequence π . Then there exist $G, H \in \Sigma(\pi)$ such that*

$$\delta_{\mathcal{M}, \mathcal{A}}(C) = \text{diam} \left(Y_{\mathcal{M}, \mathcal{A}}^*(C.G) \cup Y_{\mathcal{M}, \mathcal{A}}^*(C.H) \right) .$$

LEMMA 6.4. *Let C and C' be two reachable configurations subject to message adversary $\mathcal{M}$ via the finite communication sequences π and π' , respectively. If $C \sim_i C'$ and there exist communication sequence α and α' such that $\pi \cdot \alpha \in \mathcal{M}$, $\pi' \cdot \alpha' \in \mathcal{M}$, and i is deaf in all communication graphs in α and α' , then $Y_{\mathcal{M}, \mathcal{A}}^*(C) \cap Y_{\mathcal{M}, \mathcal{A}}^*(C') \neq \emptyset$.*

LEMMA 6.5. *Let $\Delta \geq 0$. If there exist agents $i \neq j$ and communication sequences $\alpha^{(i)}, \alpha^{(j)} \in \mathcal{M}$ such that agent i is deaf in $\alpha^{(i)}$ and agent j is deaf in $\alpha^{(j)}$, then there is an initial configuration C_0 with $\delta_{\mathcal{M}, \mathcal{A}}(C_0) = \Delta$. In particular, there is an initial configuration for which $\delta_{\mathcal{M}, \mathcal{A}}(C_0) > 0$.*

6.2 Proof of Theorem 6.1

The proof of Theorem 6.1 is by means of induction. The following relation is at the heart of the inductive step:

LEMMA 6.6. *For $i, j, \ell \in \{1, 2, 3\}$ with $\ell \neq i, j$ we have $C_t.\sigma_i \sim_\ell C_t.\sigma_j$.*

PROOF. We inductively show the following stronger statement. Let σ_i^k be the sequence of repetitions of graph Ψ_i of length $k \in [n-2]$. For agents $i, j, \ell \in \{1, 2, 3\}$ with $\ell \neq i, j$, and $m \in \{k+3, \dots, n\}$, we have $C_t.\sigma_i^k \sim_\ell C_t.\sigma_j^k$ and $C_t.\sigma_i^k \sim_m C_t.\sigma_j^k$.

Observe that agents ℓ and $\{4, \dots, n\}$ have the same in-neighbors in Ψ_i and Ψ_j . The base case $k = 1$ follows from the observation and Lemma 3.7. For the inductive step $k \mapsto k+1$, observe that agent ℓ and $\{k+4, \dots, n\}$ have only incoming edges from agents ℓ and $\{k+3, \dots, n\}$. From the hypothesis and Lemma 3.7, the inductive step follows. $\square$

We are now in the position to show Theorem 6.1 by induction. Let message adversary $\mathcal{M}_{\text{seq}}$ contain any communication sequence arising from the concatenation of σ_i sequences defined at the start of the section.

We show the stronger statement that for every initial configuration C_0 there is an execution $E = C_0, G_1, C_1, G_2, \dots$ starting from C_0 such that

$$\delta_{\mathcal{M}, \mathcal{A}}(C_t) \geq \frac{1}{2^{\lceil \frac{t}{n-2} \rceil}} \delta_{\mathcal{M}, \mathcal{A}}(C_0) \tag{11}$$

for all $t \geq 0$. It suffices to show (11) for $\mathcal{M}_{\text{seq}}$ because $\delta_{\mathcal{M}, \mathcal{A}}(C_t) \geq \delta_{\mathcal{M}_{\text{seq}}, \mathcal{A}}(C_t)$ by Lemma 3.4 and $\delta_{\mathcal{M}_{\text{seq}}, \mathcal{A}}(C_0) = \delta_{\mathcal{M}, \mathcal{A}}(C_0)$ by Lemma 6.5 whenever $\mathcal{M} \supseteq \mathcal{M}_{\text{seq}}$. We hence assume $\mathcal{M} = \mathcal{M}_{\text{seq}}$ in the rest of the proof.

The proof is by inductive construction of an execution $E = C_0, G_1, C_1, G_2, \dots$ whose configurations C_t satisfy (11). This, applied to an initial configuration with $\delta_{\mathcal{M}, \mathcal{A}}(C_0) > 0$, which exists by Lemma 6.5, then shows the theorem. The base case $t = 0$ is trivially fulfilled.

For the inductive step $t = (n - 2)k \mapsto t \leq (n - 2)(k + 1)$ assume that Equation (11) holds for $t = (n - 2)k$. First observe that, by definition of $\mathcal{M}_{\text{seq}}$, there are three possible sequences leading from round t to round $t + n - 2$, namely σ_1 , σ_2 , and σ_3 . We thus have $Y_{\mathcal{M}, \mathcal{A}}^*(C_t) = Y_{\mathcal{M}, \mathcal{A}}^*(C_{t+1}^1) \cup Y_{\mathcal{M}, \mathcal{A}}^*(C_{t+1}^2) \cup Y_{\mathcal{M}, \mathcal{A}}^*(C_{t+1}^3) = \dots = Y_{\mathcal{M}, \mathcal{A}}^*(C_{t+n-2}^1) \cup Y_{\mathcal{M}, \mathcal{A}}^*(C_{t+n-2}^2) \cup Y_{\mathcal{M}, \mathcal{A}}^*(C_{t+n-2}^3)$, where $C_{t+n-2}^u = C_t \cdot \sigma_u$ for agent $u \in \{1, 2, 3\}$.

Abbreviate $Y = Y_{\mathcal{M}, \mathcal{A}}^*(C_t)$ and $Y_u = Y_{\mathcal{M}, \mathcal{A}}^*(C_{t+n-2}^u)$. We will show that there exists a $\hat{u} \in \{1, 2, 3\}$ with

$$\text{diam}(Y_{\hat{u}}) \geq \text{diam}(Y)/2 . \quad (12)$$

We then define $C_{t+n-2} = C_{t+n-2}^{\hat{u}}$. By (12) and the induction hypothesis, we then have

$$\delta_{\mathcal{M}, \mathcal{A}}(C_{t+n-2}) = \dots = \delta_{\mathcal{M}, \mathcal{A}}(C_{t+1}) \geq \frac{\delta_{\mathcal{M}, \mathcal{A}}(C_t)}{2} \geq \frac{1}{2^{\lceil \frac{t}{n-2} \rceil + 1}} \delta_{\mathcal{M}, \mathcal{A}}(C_0) ,$$

i.e., Equation (11) holds up to round $t + n - 2$.

Assume by contradiction that for all $u \in \{1, 2, 3\}$ $\text{diam}(Y_u) < \text{diam}(Y)/2$. Since $n \geq 3$ and $C_t \cdot \sigma_i \sim_\ell C_t \cdot \sigma_j$ by Lemma 6.6, together with Lemma 6.4, we conclude that, for any pair $i, j \in \{1, 2, 3\}$ we have

$$Y_i \cap Y_j \neq \emptyset .$$

By Lemma 6.3, there exist $u, u' \in \{1, 2, 3\}$ such that $\text{diam}(Y_u \cup Y_{u'}) = \text{diam}(Y)$. In particular, we can choose $i = u$ and $j = u'$, which implies that

$$\begin{aligned} \text{diam}(Y) &= \text{diam}(Y_u \cup Y_{u'}) \leq \text{diam}(Y_u) + \text{diam}(Y_{u'}) \\ &< \text{diam}(Y) \end{aligned}$$

which is a contradiction and concludes the proof.

7 RELATION TO EXACT CONSENSUS AND GENERALIZED BOUNDS

In the *exact consensus* problem, the local state of an agent i is augmented with a variable d^i initialized to $\perp$. Agent i is allowed to set d^i to some value $v \neq \perp$ only once, in which case we say that i *decides* v . An algorithm *solves exact consensus* subject to $\mathcal{M}$ if each execution E with a communication sequence in $\mathcal{M}$ satisfies:

- *Termination.* Each agent eventually decides.
- *Exact agreement.* If agents i and j have decided, then the decision values are equal: $d^i = d^j$
- *Validity.* If agent i decides v , then v is one of the initial values $y_E^1(0), \dots, y_E^n(0)$.

Coulouma et al. [22] characterized the oblivious message adversaries in which exact consensus is solvable. While their characterization was stated for binary consensus in which initial values and outputs are in $\{0, 1\}$, their proofs also hold for the exact consensus as defined above. Charron-Bost et al. [17] showed that approximate consensus is solvable in a significantly broader class: it is solvable subject to an oblivious message adversary if and only if it is rooted. As discussed in Section 1, the same characterization holds for asymptotic consensus. In this section we aim to shed light on the deeper relation between exact consensus and asymptotic consensus by studying valencies and contraction ratios. Our main results are a characterization of the topological structure of valencies with respect to solvability of exact consensus (Theorem 7.9) and nontrivial lower bounds on the contraction ratios whenever exact consensus is not solvable (Theorem 7.12 and Corollary 7.13). We

Fig. 5. Graphs H_0 , H_1 , and H_2 from Example 7.2 with roots colored in red. The figure also shows the set of nodes $\text{In}_{R(H_1)}(G) = \text{In}_{\{1\}}(G)$ for all graphs $G \in \mathcal{G} = \{H_0, H_1, H_2\}$ in blue. The relation α_{H_1} holds among graphs with the same set $\text{In}_{R(H_1)}$, i.e., the same blue area.

also prove that a contraction ratio of 0 can be achieved if and only if exact consensus is solvable (Theorem 7.14).

We start by recalling some definitions from Coulouma et al. [22]. In the following, we denote by $R(G)$ the set of roots of a communication graph G , i.e., the set of agents that have a directed path to all other agents in G . For a set $S \subseteq [n]$, let $\text{In}_S(G) = \bigcup_{j \in S} \text{In}_j(G)$. The set $\text{Out}_S(G)$ is defined analogously.

Definition 7.1 ([22, Definition 4.7]). Let $\mathcal{G}$ be a set of communication graphs. Given $G, H, K \in \mathcal{G}$, we define $G\alpha_K H$ if $\text{In}_{R(K)}(G) = \text{In}_{R(K)}(H)$. The relation $\alpha_{\mathcal{G}}^*$ is the transitive closure of the union of the relations α_K where K varies in $\mathcal{G}$.

Example 7.2. Let $\mathcal{G} = \{H_0, H_1, H_2\}$ with the graphs H_0 , H_1 , and H_2 as previously defined. The sets of roots of the graphs are given by $R(H_0) = \{1, 2\}$, $R(H_1) = \{1\}$, and $R(H_2) = \{2\}$; see Figure 5. Since $\text{In}_{R(H_1)}(H_0) = \text{In}_{\{1\}}(H_0) = \text{In}_{\{1\}}(H_2) = \{1, 2\}$, it is $H_0\alpha_{H_1}H_2$. However, neither $H_0\alpha_{H_1}H_1$, nor $H_1\alpha_{H_1}H_2$. $\square$

Definition 7.3 ([22, Definition 4.8]). Let $\mathcal{G}$ be a set of communication graphs. We define $\beta_{\mathcal{G}}$ to be the coarsest equivalence relation included in $\alpha_{\mathcal{G}}^*$ such that for all G, H holds:

(Closure Property) If $G\beta_{\mathcal{G}}H$, then there exists a nonnegative integer q and communication graphs $H_0, \dots, H_q \in \mathcal{G}$ and $K_1, \dots, K_q \in \mathcal{G}$ such that

- (i) $G = H_0$ and $H = H_q$
- (ii) $\forall r \in [q]: H_r\beta_{\mathcal{G}}G$ and $K_r\beta_{\mathcal{G}}G$
- (iii) $\forall r \in [q]: H_{r-1}\alpha_{K_r}H_r$

We next show properties of subsets of $\mathcal{G}$ that are $\beta_{\mathcal{G}}$ -classes.

LEMMA 7.4. Let $\mathcal{G}$ be a set of communication graphs and let $\mathcal{G}' \subseteq \mathcal{G}$ be a $\beta_{\mathcal{G}}$ -class. Then $G\alpha_{\mathcal{G}'}^*H$ and $G\beta_{\mathcal{G}'}H$ for all $G, H \in \mathcal{G}'$.

PROOF. Let $G, H \in \mathcal{G}'$. Since $G\beta_{\mathcal{G}}H$, there is a q and $H_0, \dots, H_q \in \mathcal{G}$ and $K_1, \dots, K_q \in \mathcal{G}$ such that (i) $G = H_0$ and $H = H_q$ (ii) $H_r\beta_{\mathcal{G}}G$ and $K_r\beta_{\mathcal{G}}G$ for all $r \in [q]$, and (iii) $H_{r-1}\alpha_{K_r}H_r$ for all $r \in [q]$.

Condition (ii) implies $H_0, \dots, H_q \in \mathcal{G}'$ and $K_1, \dots, K_q \in \mathcal{G}'$ since they belong to the same $\beta_{\mathcal{G}}$ -class as G , i.e., $\mathcal{G}'$. But this means that the pair (G, H) is in the transitive closure of the union of the relations $\alpha_{K_1}, \dots, \alpha_{K_q}$. Further, all K_r are in $\mathcal{G}'$. Thus $G\alpha_{\mathcal{G}'}^*H$ holds. Hence $\alpha_{\mathcal{G}'}^* = \mathcal{G}' \times \mathcal{G}'$, i.e., the first part of the lemma.

To show the second part, define relation $\tilde{\beta} = \mathcal{G}' \times \mathcal{G}'$, which, as we just proved, is included in $\alpha_{\mathcal{G}'}^*$. But it also satisfies the closure property in $\mathcal{G}'$. Since $\tilde{\beta}$ is the coarsest equivalence relation on $\mathcal{G}'$, we thus have $\beta_{\mathcal{G}'} = \tilde{\beta} = \mathcal{G}' \times \mathcal{G}'$, i.e., the second part of the lemma. $\square$

Definition 7.5 ([22, Definition 4.5]). A set $\mathcal{G}$ of communication graphs is called *source-incompatible* if

$$\bigcap_{G \in \mathcal{G}} R(G) = \emptyset .$$

While Coulouma et al. [22] focus on binary consensus in their work and provide a characterization for binary consensus, their proof actually shows a stronger version of their theorem. Indeed the characterization also holds for exact consensus with inputs and outputs in $\mathbb{R}^d$ by the same proof [22]:

THEOREM 7.6 ([22, THEOREM 4.10]). *Let $\mathcal{M} = \mathcal{M}(\mathcal{G})$ be an oblivious message adversary. Exact consensus is solvable subject to $\mathcal{M}$ if and only if each $\beta_{\mathcal{G}}$ -class is not source-incompatible.*

We start with showing a generalization of Lemma 3.8 that allows us to induce non-empty intersection of valencies.

LEMMA 7.7. *Let C be a configuration of an asymptotic consensus algorithm $\mathcal{A}$ for $\mathcal{M} = \mathcal{M}(\mathcal{G})$. For all configurations C in an execution of $\mathcal{A}$ in $\mathcal{M}$, and for all $G, H, K \in \mathcal{G}$, if $G\alpha_K H$ then $Y_{\mathcal{M}, \mathcal{A}}^*(C.G) \cap Y_{\mathcal{M}, \mathcal{A}}^*(C.H) \neq \emptyset$.*

PROOF. By the definition of $G\alpha_K H$ it is $\text{In}_{R(K)}(G) = \text{In}_{R(K)}(H)$. Hence, together with Lemma 3.7, it follows that $C.G \sim_i C.H$ for all nodes i in $R(K)$. We consider an execution E in which C occurs at some $t_0 - 1$, G is the communication graph at t_0 and all following graphs are equal to K . Analogously, let E' be an execution identical to E except that the communication graph at round t_0 is H instead of G . By inductive application of Lemma 3.7, we show that for all $t \geq t_0$, we have $C_t \sim_i C'_t$. In particular, we obtain $y_E^i(t) = y_{E'}^i(t)$. Thus $y^*(E) = y^*(E')$, which shows that $Y_{\mathcal{M}, \mathcal{A}}^*(C.G)$ and $Y_{\mathcal{M}, \mathcal{A}}^*(C.H)$ intersect. $\square$

We next establish that for oblivious message adversaries subject to which exact consensus is not solvable, asymptotic consensus algorithms must have initial configurations that can be extended to executions with different limit outputs.

LEMMA 7.8. *Let $\mathcal{M}$ be an oblivious message adversary for which exact consensus is not solvable. Then for all asymptotic consensus algorithms $\mathcal{A}$, there exists an initial configuration C_0 such that $Y_{\mathcal{M}, \mathcal{A}}^*(C_0)$ is not a singleton.*

More precisely, for every $\Delta > 0$, there exists an initial configuration C_0 such that $\Delta(C_0) \leq \Delta$ and $\delta_{\mathcal{M}, \mathcal{A}}(C_0) \geq \Delta/n$.

PROOF. By embedding the initial values into the first axis, we assume $d = 1$ in the rest of the proof. Since the lemma claims an existence result, this is sufficient.

Let $\mathcal{M} = \mathcal{M}(\mathcal{G})$ and let $\mathcal{G}' \subseteq \mathcal{G}$ be any source-incompatible $\beta_{\mathcal{G}}$ -class, which exists by Theorem 7.6. Consider the $n + 1$ initial configurations $C_0^{(k)}$ where $0 \leq k \leq n$ with initial values

$$y_i^{(k)}(0) = \begin{cases} \Delta & \text{if } i \leq k \\ 0 & \text{if } i > k \end{cases} .$$

For all these initial configurations, we have $\Delta(C_0^{(k)}) \leq \Delta$. Define $a(k) = \inf Y_{\mathcal{M}, \mathcal{A}}^*(C_0^{(k)})$ and $b(k) = \sup Y_{\mathcal{M}, \mathcal{A}}^*(C_0^{(k)})$. By Validity, $Y_{\mathcal{M}, \mathcal{A}}^*(C_0^{(0)}) = \{0\}$ and $Y_{\mathcal{M}, \mathcal{A}}^*(C_0^{(n)}) = \{\Delta\}$, which means $a(0) = b(0) = 0$ and $a(n) = b(n) = \Delta$. There exists some k with $1 \leq k \leq n$ such that $b(k-1) \leq b(k) - \Delta/n$ since otherwise $0 = b(0) > b(n) - \Delta = 0$. Because $\mathcal{G}'$ is source-incompatible, for every agent k , there exists a communication graph $G^{(k)} \in \mathcal{G}'$ such that $k \notin S(G^{(k)})$. Since $C_0^{(k-1)} \sim_i C_0^{(k)}$ for all $i \in S(G^{(k)})$, choosing two executions with all communication graphs equal to $G^{(k)}$ shows

that $Y_{\mathcal{M}', \mathcal{A}}^*(C_0^{(k-1)}) \cap Y_{\mathcal{M}', \mathcal{A}}^*(C_0^{(k)}) \neq \emptyset$, which implies $a(k) \leq b(k-1)$. Combining both inequalities gives $a(k) \leq b(k) - \Delta/n$ and shows that $\delta_{\mathcal{M}', \mathcal{A}}(C_0^{(k)}) = b(k) - a(k) \geq \Delta/n$. We hence choose the initial configuration $C_0 = C_0^{(k)}$.

This shows $\delta_{\mathcal{M}, \mathcal{A}}(C_0) \geq \delta_{\mathcal{M}', \mathcal{A}}(C_0) \geq \Delta/n$ by Lemma 3.4 and concludes the proof. $\square$

This finally allows us to derive one of our main results of this section: a characterization of oblivious message adversaries for which exact consensus is solvable by the topological structure of valencies of asymptotic consensus algorithms.

THEOREM 7.9. *Let $\mathcal{M}$ be an oblivious message adversary. Exact consensus is solvable in $\mathcal{M}$ if and only if there exists an asymptotic consensus algorithm $\mathcal{A}$ for $\mathcal{M}$ such that $Y_{\mathcal{M}', \mathcal{A}}^*(C_0)$ is either a singleton or disconnected for all oblivious message adversaries $\mathcal{M}' \subseteq \mathcal{M}$ and all initial configurations C_0 of $\mathcal{A}$.*

PROOF. ($\Rightarrow$): Assume that exact consensus is solvable subject to $\mathcal{M}$, and let $\mathcal{A}'$ be an algorithm that solves exact consensus subject to $\mathcal{M}$. Let $\mathcal{A}$ be the algorithm derived from $\mathcal{A}'$ in which deciding is replaced by setting its output variable to the decision value of $\mathcal{A}'$ and not changing it anymore. Before the decision of algorithm $\mathcal{A}'$, algorithm $\mathcal{A}$ outputs its initial value. Then $\mathcal{A}$ is an asymptotic consensus algorithm subject to $\mathcal{M}$. Further, from Validity of exact consensus, for any initial configuration C_0 , the valency $Y_{\mathcal{M}, \mathcal{A}}^*(C_0)$ is a subset of the set of initial values in C_0 . As the set of initial values of C_0 is finite, so is $Y_{\mathcal{M}, \mathcal{A}}^*(C_0)$ and, by Lemma 3.4, also $Y_{\mathcal{M}', \mathcal{A}}^*(C_0)$ for all $\mathcal{M}' \subseteq \mathcal{M}$. Since any finite set is either a singleton or disconnected, the claim follows.

($\Leftarrow$): We assume without loss of generality that $d = 1$. If not, we embed the initial values in any 1-dimensional affine subspace, e.g., choose them to lie on the first axis.

We proceed by means of contradiction. Assume that exact consensus is unsolvable subject to $\mathcal{M} = \mathcal{M}(\mathcal{G})$. We will show that for all asymptotic consensus algorithms $\mathcal{A}$ for $\mathcal{M}$, there exists an initial configuration C_0 and an oblivious message adversary $\mathcal{M}' = \mathcal{M}(\mathcal{G}') \subseteq \mathcal{M}$ such that $Y_{\mathcal{M}', \mathcal{A}}^*(C_0)$ is a nontrivial interval.

By Theorem 7.6, there is a source-incompatible $\beta_{\mathcal{G}}$ -class. Choose $\mathcal{G}'$ to be equal to such a class. We choose C_0 via Lemma 7.8 such that $Y_{\mathcal{M}', \mathcal{A}}^*(C_0)$ is not a singleton.

To show that $Y_{\mathcal{M}', \mathcal{A}}^*(C_0)$ is connected, we assume to the contrary that it is not and derive a contradiction. The set $Y_{\mathcal{M}', \mathcal{A}}^*(C_0)$ not being connected means the existence of some $z \notin Y_{\mathcal{M}', \mathcal{A}}^*(C_0)$ such that

$$\exists z_1, z_2 \in Y_{\mathcal{M}', \mathcal{A}}^*(C_0): z_1 < z < z_2. \quad (13)$$

We will inductively construct an execution $E = C_0, G_1, C_1, G_2, \dots$ such that

$$\exists z_1, z_2 \in Y_{\mathcal{M}', \mathcal{A}}^*(C_t): z_1 < z < z_2 \quad (14)$$

for all $t \geq 0$. Setting $m(t) = \inf Y_{\mathcal{M}', \mathcal{A}}^*(C_t)$ and $M(t) = \sup Y_{\mathcal{M}', \mathcal{A}}^*(C_t)$, we then have $m(t) \leq z \leq M(t)$ by (14) and $M(t) - m(t) = \delta_{\mathcal{M}', \mathcal{A}}(C_t) \rightarrow 0$ by Convergence and Agreement. Hence $\lim_{t \rightarrow \infty} m(t) = \lim_{t \rightarrow \infty} M(t) = z$, which means

$$\lim_{t \rightarrow \infty} Y_{\mathcal{M}', \mathcal{A}}^*(C_t) = \bigcap_{t \geq 0} Y_{\mathcal{M}', \mathcal{A}}^*(C_t) = \{z\},$$

where the first equality follows from Lemma 3.5. In particular $z \in Y_{\mathcal{M}', \mathcal{A}}^*(C_0)$, which gives the desired contradiction.

It thus suffices to construct an execution E satisfying (14). Assume that (14) holds for a given $t \geq 0$ and let $z_1^{(t)}, z_2^{(t)} \in Y_{\mathcal{M}', \mathcal{A}}^*(C_t)$ with $z_1^{(t)} < z < z_2^{(t)}$. By Lemma 3.5, it follows that there are communication graphs $G, H \in \mathcal{G}'$ with $z_1^{(t)} \in Y_{\mathcal{M}', \mathcal{A}}^*(C.G)$ and $z_2^{(t)} \in Y_{\mathcal{M}', \mathcal{A}}^*(C.H)$. By Lemma 7.4,

we have $G\alpha_{\mathcal{G}}^*H$. Thus there exists a chain $G = H_0, H_1, \dots, H_q = H \in \mathcal{G}'$ and communication graphs $K_1, \dots, K_q \in \mathcal{G}'$ such that $H_{r-1}\alpha_{K_r}H_r$ for all $r \in [q]$. From Lemma 7.7 we thus know that

$$Y_{\mathcal{M}', \mathcal{A}}^*(C.H_{r-1}) \cap Y_{\mathcal{M}', \mathcal{A}}^*(C.H_r) \neq \emptyset \quad (15)$$

for all $r \in [q]$. Set $f(r) = \inf Y_{\mathcal{M}', \mathcal{A}}^*(C.H_r)$ and $g(r) = \sup Y_{\mathcal{M}', \mathcal{A}}^*(C.H_r)$ for $r \in \{0, \dots, q\}$, and

$$\hat{r} = \min \{r \in \{0, \dots, q\} \mid g(r) > z\}.$$

Then $f(0) \leq z_1^{(t)} \leq g(0)$ and $f(q) \leq z_2^{(t)} \leq g(q)$. The quantity $\hat{r}$ is finite since $g(q) \geq z_2^{(t)} > z$. We show $f(\hat{r}) < z$ by distinguishing two cases:

(1) $\hat{r} = 0$: Then $f(\hat{r}) = f(0) \leq z_1^{(t)} < z$.

(2) $\hat{r} \geq 1$: Then, by (15) and the definition of $\hat{r}$, we have $f(\hat{r}) \leq g(\hat{r} - 1) < z$.

In both cases, we showed $f(\hat{r}) < z < g(\hat{r})$. Choosing $G_{t+1} = H_{\hat{r}}$ and $C_{t+1} = C_t.G_{t+1}$, we hence proved (14) for $t + 1$. This concludes the proof. $\square$

We next introduce the $\alpha_{\mathcal{G}}$ -diameter of an oblivious message adversary $\mathcal{M} = \mathcal{M}(\mathcal{G})$, which we will then (see Theorem 7.12 and Corollary 7.13) show to be directly linked to a nontrivial lower bound on the contraction ratio subject to $\mathcal{M}$ if exact consensus is not solvable subject to $\mathcal{M}$.

Definition 7.10. Let $\mathcal{G}$ be a set of communication graphs. We say that $\mathcal{G}$ is $\alpha_{\mathcal{G}}$ -connected if $\alpha_{\mathcal{G}}^*$ has a single equivalence class. Otherwise, we call it $\alpha_{\mathcal{G}}$ -disconnected.

The $\alpha_{\mathcal{G}}$ -diameter of an $\alpha_{\mathcal{G}}$ -connected $\mathcal{G}$ is the smallest $D \geq 1$ such that for all $G, H \in \mathcal{G}$ there exist communication graphs $H_0, \dots, H_q \in \mathcal{G}$ and $K_1, \dots, K_q \in \mathcal{G}$ with $q \leq D$ such that $G = H_0$, $H = H_q$, and $H_{r-1}\alpha_{K_r}H_r$ for all $r \in [q]$.

Example 7.11. Observe that for $\mathcal{G} = \{H_0, H_1, H_2\}$ from Theorem 4.1, the $\alpha_{\mathcal{G}}$ -diameter of $\mathcal{G}$ is $D = 2$: From Example 7.2 and Figure 5 one immediately verifies that $H_2\alpha_{H_1}H_0\alpha_{H_2}H_1$, but no graph G in $\mathcal{G}$ relates all three graphs via $\alpha_{\mathcal{G}}$.

By contrast, the set $\mathcal{G} = \{H_1, H_2\}$ is not $\alpha_{\mathcal{G}}$ -connected since neither $H_1\alpha_{H_1}H_2$ nor $H_1\alpha_{H_2}H_2$.

For the oblivious message adversary $\mathcal{M}(\text{deaf}(G))$, where G is an arbitrary communication graph G with $n \geq 3$ nodes, we have $D = 1$: Recall that $\text{deaf}(G) = \{F_1, \dots, F_n\}$ with $F_i = G \setminus \{(j, i) \mid j \in [n] \setminus \{i\}\}$. Let $i, j \in [n]$ and chose $k \in [n] \setminus \{i, j\}$. Then $R(F_k)$ is either $\emptyset$ or $\{k\}$. In both cases, $\text{In}_{R(F_k)}(F_i) = \text{In}_{R(F_k)}(F_j)$, and thus $F_i\alpha_{F_k}F_j$. $\square$

The following theorem and corollary thus generalize Theorems 4.1 and 5.1 to arbitrary oblivious message adversaries for which exact consensus is not solvable.

THEOREM 7.12. *Let $\mathcal{M} = \mathcal{M}(\mathcal{G})$ be an oblivious message adversary for which exact consensus is not solvable. If $\mathcal{G}$ is $\alpha_{\mathcal{G}}$ -connected, then the contraction ratio of any asymptotic consensus algorithm subject to $\mathcal{M}$ is greater or equal to $1/(D + 1)$ where D is the $\alpha_{\mathcal{G}}$ -diameter of $\mathcal{G}$.*

PROOF. We show the stronger statement that for every initial configuration C_0 there is an execution $E = C_0, G_1, C_1, G_2, \dots$ starting at C_0 such that

$$\delta_{\mathcal{M}, \mathcal{A}}(C_t) \geq \frac{1}{(D + 1)^t} \delta_{\mathcal{M}, \mathcal{A}}(C_0) \quad (16)$$

for all $t \geq 0$. This, applied to an initial configuration with $\delta_{\mathcal{M}, \mathcal{A}}(C_0) > 0$, which exists by Lemma 7.8, then shows the theorem.

The proof is by inductive construction of an execution $E = C_0, G_1, C_1, G_2, \dots$ whose configurations C_t satisfy (16).

For $t = 0$ the inequality trivially holds.

Now let t be any nonnegative integer and assume that Equation (16) holds for t . By Lemma 3.6, there exist $G, H \in \mathcal{G}$ such that $\text{diam}(Y_{\mathcal{M}, \mathcal{A}}^*(C_t)) = \text{diam}(Y_{\mathcal{M}, \mathcal{A}}^*(C_t.G) \cup Y_{\mathcal{M}, \mathcal{A}}^*(C_t.H))$. Because the $\alpha_{\mathcal{G}}$ -diameter of $\mathcal{G}$ is equal to $D < \infty$, there exist communication graphs $H_0, \dots, H_q \in \mathcal{G}$ and $K_1, \dots, K_q \in \mathcal{G}$ with $q \leq D$ such that $G = H_0$, $H = H_q$, and $H_{r-1} \alpha_{K_r} H_r$ for all $r \in [q]$.

Define $Y = Y_{\mathcal{M}, \mathcal{A}}^*(C_t)$ and $Y_r = Y_{\mathcal{M}, \mathcal{A}}^*(C_t.H_r)$. We have $\text{diam}(Y) = \text{diam}(Y_0 \cup Y_q)$ by choice of $G = H_0$ and $H = H_q$. We show that there exists some $r \in \{0, \dots, q\}$ such that $\text{diam}(Y_r) \geq \text{diam}(Y)/(q+1)$ and then set $G_{t+1} = H_r$ and $C_{t+1} = C_t.H_r$. Then, by the induction hypothesis, we have

$$\delta_{\mathcal{M}, \mathcal{A}}(C_{t+1}) \geq \frac{\delta_{\mathcal{M}, \mathcal{A}}(C_t)}{q+1} \geq \frac{\delta_{\mathcal{M}, \mathcal{A}}(C_t)}{D+1} \geq \frac{1}{(D+1)^{t+1}} \delta_{\mathcal{M}, \mathcal{A}}(C_0) , \quad (17)$$

i.e., Equation (16) holds for $t+1$.

Assume by contradiction that for all $r \in \{0, \dots, q\}$ $\text{diam}(Y_r) < \text{diam}(Y)/(q+1)$. By Lemma 7.7, we have $Y_{r-1} \cap Y_r \neq \emptyset$ for all $r \in [q]$. Inductively, we prove

$$\text{diam}\left(\bigcup_{s=0}^r Y_s\right) < \frac{r+1}{q+1} \cdot \text{diam}(Y) \quad (18)$$

for all $r \in \{0, \dots, q\}$. In particular for $r = q$, which leads to $\text{diam}(Y) \leq \text{diam}(Y_0 \cup Y_q) < \text{diam}(Y)$, which is a contradiction and concludes the proof. $\square$

Direct application of Theorem 7.12 to an oblivious message adversary $\mathcal{M} = \mathcal{M}(\mathcal{G})$ in which exact consensus is not solvable may yield a trivial bound of 0 if $\mathcal{G}$ is not $\alpha_{\mathcal{G}}$ -connected. We can, however, use Lemma 3.4 to derive a strictly positive bound for any $\mathcal{M}$ subject to which exact consensus is not solvable:

COROLLARY 7.13. *Let $\mathcal{M} = \mathcal{M}(\mathcal{G})$ be an oblivious message adversary for which exact consensus is not solvable. The contraction ratio of any asymptotic consensus algorithm subject to $\mathcal{M}$ is greater or equal to $1/(D+1)$ where D is the infimum of $\alpha_{\mathcal{G}'}$ -diameters of any $\alpha_{\mathcal{G}'}$ -connected $\mathcal{G}' \subseteq \mathcal{G}$ such that exact consensus is not solvable subject to $\mathcal{M}(\mathcal{G}')$.*

PROOF. Set $\mathcal{G}' \subseteq \mathcal{G}$ equal to the set with the smallest $\alpha_{\mathcal{G}'}$ -diameter such that exact consensus is not solvable subject to $\mathcal{M}(\mathcal{G}')$. Applying Theorem 7.12 to $\mathcal{M}'$, and Lemma 3.4 (iv) to $\mathcal{M}'$ and $\mathcal{M}$ yields the corollary. $\square$

Corollary 7.13 gives a trivial lower bound of 0 on the contraction ratio if exact consensus is solvable subject to $\mathcal{M}$ (and thus all its sub-message adversaries). It turns out that the converse also holds: If a contraction ratio of 0 can be achieved, then exact consensus is solvable. This is shown in the next theorem. We thus indeed have a nontrivial lower bound in Corollary 7.13 whenever possible.

THEOREM 7.14. *Let $\mathcal{M}$ be an oblivious message adversary. Exact consensus is solvable subject to $\mathcal{M}$ if and only if there is an asymptotic consensus algorithm that achieves a contraction ratio of 0 subject to $\mathcal{M}$.*

PROOF. ($\Rightarrow$): Starting from an exact consensus algorithm, we can construct an asymptotic consensus algorithm by having each agent output its initial value before the exact consensus algorithm would have decided, and its decision value after. The valency of a configuration after all agents would have decided in the exact consensus algorithm is a singleton. The diameter of the valency is thus eventually 0 in every execution of the asymptotic consensus algorithm. But then the algorithm's contraction ratio is 0.

($\Leftarrow$): Assume by contradiction that there is an asymptotic consensus algorithm that has contraction ratio 0 subject to $\mathcal{M} = \mathcal{M}(\mathcal{G})$ and that exact consensus is not solvable subject to $\mathcal{M}$. Then,

by Corollary 7.13, every $\mathcal{G}' \subseteq \mathcal{G}$ such that exact consensus is not solvable subject to $\mathcal{M}(\mathcal{G}')$ is $\alpha_{\mathcal{G}}$ -disconnected. We show that this is impossible.

By Theorem 7.6 there is a $\beta_{\mathcal{G}}$ -class C that is source-incompatible. By the Closure Property of the $\beta_{\mathcal{G}}$ -relation, the class C is also a β_C -class. Now, by invoking Theorem 7.6 for the oblivious message adversary $\mathcal{M}(C)$, since C is source-incompatible, exact consensus is not solvable subject to $\mathcal{M}(C)$. But the single β_C -class is α_C -connected, and thus also $\alpha_{\mathcal{G}}$ -connected. The class C is thus $\alpha_{\mathcal{G}}$ -connected and exact consensus is not solvable subject to $\mathcal{M}(C)$, a contradiction to Corollary 7.13. $\square$

8 ALMOST-TIGHT BOUNDS FOR ASYNCHRONOUS SYSTEMS WITH CRASHES: THE PRICE OF ROUNDS

In this section we show that Corollary 7.13 provides a tool to clearly separate time complexities of algorithms that operate in rounds to general algorithms in the classical static fault model of asynchronous message-passing systems with crashes. Our result applies to algorithms without any restriction: we do not make assumptions on the nature of the functions used by the agents, and agents are not required to be memoryless.

We start with recalling and adapting notation for classical asynchronous message-passing systems. We consider a distributed system where agents perform receive-compute-send steps. An agent may crash, i.e., stop making steps. Crashes can be unclean: the final broadcast message may be received by a proper subset of correct, i.e., non-crashed, agents, only. Since an agent that crashes stops to make steps, we require Convergence, Validity, and Agreement of asymptotic consensus to hold only for the set of correct agents. Analogously, the consensus function y^* , and thus the valencies, are restricted to correct agents only. Further, we apply the standard convention of measuring time in asynchronous systems by normalizing to the longest end-to-end message delay from a broadcast to the respective reception in an execution.

8.1 Round-based Algorithms

An algorithm is said to operate in rounds if each agent waits for $n - f$ messages corresponding to the current round, updates its state based on the received messages and its previous state, and broadcasts the next round's messages. If up to f agents can crash, and thus are muted, i.e., stop sending messages, then waiting for more than $n - f$ messages can lead to a deadlock. Algorithms that operate in rounds are widely used in asynchronous systems (see, e.g., [19, 25, 44]).

We next show that Corollary 7.13 can be applied to obtain new asymptotically tight bounds for round-based algorithms. Specifically, we prove a lower bound for asynchronous systems of size $n \geq 3$ with up to $f < n/2$ crashes whose agents operate in rounds.

Let us construct the following set of communication graphs: denote by $\mathcal{G}_n$ the set of communication graphs with n nodes and let

$$\mathcal{G}_A = \{G \in \mathcal{G}_n \mid \forall i \in [n]: |\text{In}_i(G)| \geq n - f\}$$

for some $f < n/2$.

LEMMA 8.1. *The $\alpha_{\mathcal{G}_A}$ -diameter of $\mathcal{G}_A$ is at most $\lceil n/f \rceil$.*

PROOF. Let $G, H \in \mathcal{G}_A$ and $q = \lceil n/f \rceil$. We choose the communication graphs H_r , where $0 \leq r \leq q$, such that

$$\text{In}_i(H_r) = \begin{cases} \text{In}_i(G) & \text{if } 1 \leq i \leq rf \\ \text{In}_i(H) & \text{if } rf + 1 \leq i \leq n \end{cases}.$$

Further, for $r \in [q]$, we choose the communication graphs K_r such that

$$\text{In}_i(K_r) = [n] \setminus \{j \mid (r-1)f + 1 \leq j \leq rf\}.$$

Fig. 6. Visualization of the argument in the proof of Lemma 8.1 stating that $H_{r-1}\alpha_{K_r}H_r$ for the first two $r \in [q]$. The figure shows the nodes of the graphs H_0 , H_1 , and H_2 . Nodes are labeled by whether their in-neighbors are as in H or as in G . The nodes in $R(K_1)$ are shown in blue, the nodes in $R(K_2)$ in red. One immediately observes that $H_0\alpha_{K_1}H_2\alpha_{K_3}H_3$ by comparing in-neighbors for the nodes in $R(K_1)$ in graphs H_0 and H_1 , and in-neighbors for the nodes in $R(K_2)$ in graphs H_1 and H_2 .

Note that all such H_r and K_r are uniquely defined via the in-neighbors of each node $i \in [n]$.

Since, for $r = 0$, it is $\text{In}_i(H_r) = \text{In}_i(H)$ for all nodes $i \in [n]$, and for $r = q$, it is $\text{In}_i(H_r) = \text{In}_i(G)$ for all nodes $i \in [n]$, we have $H_0 = H$ and $H_q = G$.

Further, for all $r \in [q]$, we can write $R(K_r) = \text{In}_i(K_r) = [n] \setminus \{j \mid (r-1)f + 1 \leq j \leq rf\}$. Together with the fact that for all $r \in [q]$ and all nodes $i \in R(K_r)$, it is $\text{In}_i(H_{r-1}) = \text{In}_i(H_r)$, we obtain $H_{r-1}\alpha_{K_r}H_r$ for all $r \in [q]$; see Figure 6. Noting $H_r \in \mathcal{G}_A$ and $K_r \in \mathcal{G}_A$ for all $r \in [q]$, this concludes the proof. $\square$

From Lemma 8.1 and Corollary 7.13 we immediately obtain the lower bound:

THEOREM 8.2. *The contraction ratio for any asymptotic consensus algorithm for $n \geq 3$ agents and at most $f < n/2$ crashes that operates in rounds is greater or equal to $\frac{1}{\lceil n/f \rceil + 1}$.*

Note that the contraction ratio in Theorem 8.2 is with respect to rounds. However, we can construct an execution where a single round requires $1 + \varepsilon$ time for arbitrarily small $\varepsilon > 0$: we assign all messages that are delivered according to the communication graph of the respective round, delay 1, and all others delay $1 + \varepsilon$. Theorem 8.2 thus also holds for a contraction ratio with respect to time.

8.2 General Algorithms

Within this section, consider the classical static fault model of asynchronous message-passing systems with crashes; see, e.g., [31, 44] for a detailed definition of this model. We show that there

is an algorithm that does not operate in rounds that ensures that all agents' outputs are equal by time $f + 1$. This gives a contraction ratio of 0.

The following algorithm *MinRelay* is inspired by the exact consensus algorithm for synchronous systems with crash faults (see, e.g., [44]), and is based on a non-terminating reliable broadcast protocol: Initially, at time 0, each agent i sets S^i to the set containing only its initial value, and broadcasts S^i . Whenever an agent i receives a set $S \neq S^i$, it sets $S^i \leftarrow S^i \cup S$, $y^i \leftarrow \min(S^i)$, and broadcasts S^i .

THEOREM 8.3. *The MinRelay algorithm solves asymptotic consensus in asynchronous message-passing systems with up to $f < n$ crashes. Specifically, all correct agents' sets S^i , and thus y^i , are equal by time $f + 1$, and the algorithm's contraction ratio is 0.*

PROOF. We first show equality of sets S^i by time $f + 1$. Assume by means of contradiction that there exist two correct agents i, j with $S^i \neq S^j$ after time $f + 1$. Then there exists an x in S^i that is not in S^j . We distinguish between two cases:

Case i: x was added to S^i at latest by time f . By the algorithm and the maximum message delay of 1, x is added to S^j by time $f + 1$; a contradiction.

Case ii: Otherwise, x was added to S^i after time f . Consider the causal chain of messages that lead to adding x at agent i . By the algorithm, its origin must be a message broadcast at time 0. Together with the maximum message delay of 1, the chain must contain at least $f + 1$ broadcasts during the time interval $[0, f]$. At most f of these broadcasts may be ones where an agent crashed and stopped making steps. Thus there is at least one broadcast among them that happened at an agent that did not crash during the broadcast. By the maximum message delay 1, node j received this message by time $f + 1$, adding x to S^j ; a contradiction to the assumption. The claim follows.

Convergence, Agreement, and Validity follow from equality of all correct agents' S^i after time $f + 1$, the fact all elements in S^i are initial values, and the properties of the function $\min$. $\square$

9 APPROXIMATE CONSENSUS

Alternatively to asymptotic consensus, one may also consider the *approximate consensus* problem, in which convergence is replaced by a decision in a finite number of rounds and where agreement should be achieved with an arbitrarily small error tolerance (see, e.g., [44]). Like for exact consensus, formally, the local state of i is augmented with variable d^i initialized to $\perp$. Again, agent i is allowed to set d^i to some value $v \neq \perp$ only once, in which case we say that i *decides* v . In addition to the initial values $y^i(0)$, agents initially receive the error tolerance ε and an upper bound Δ on the maximum distance of initial values. An algorithm *solves approximate consensus* subject to $\mathcal{M}$ if for all $\varepsilon > 0$ and all Δ , each execution E with a communication sequence in $\mathcal{M}$ with initial diameter at most Δ satisfies:

- *Termination.* Each agent eventually decides.
- *ε -Agreement.* If agents i and j decide v and v' , then we have $\|v - v'\| \leq \varepsilon$.
- *Validity.* If agent i decides v , then v is in the convex hull of initial values $y_E^1(0), \dots, y_E^n(0)$.

Asymptotic consensus and approximate consensus are clearly closely related. However, the ε -Agreement condition does not preclude that the decisions of a given agent in a sequence of executions with same initial values and communication sequences, but with different error tolerance parameters $\varepsilon \rightarrow 0$, diverges. It may thus lead to unstable decisions with respect to this parameter, as shown by the following example.

Example 9.1. Consider the oblivious message adversary $\mathcal{M} = \mathcal{M}(\{H_0\})$, with H_0 from Figure 3, and the following Algorithm $\mathcal{A}$ that is easily seen to solve approximate consensus subject to $\mathcal{M}$: If

the initially provided $\varepsilon > 0$ fulfills $\lceil 1/\varepsilon \rceil \equiv 0 \pmod{2}$, then both agents output $y^1(0)$ at the end of round $\lceil 1/\varepsilon \rceil$; otherwise, both agents output $y^2(0)$ at the end of round $\lceil 1/\varepsilon \rceil$.

We can observe that, if the sequence $\varepsilon_s = 1/s$ is provided as initial parameter to $\mathcal{A}$ and given that the initial values $y^1(0)$ and $y^2(0)$ differ, the (identical) decision values of both agents oscillate between $y^1(0)$ and $y^2(0)$ and thus do not converge as $\varepsilon \rightarrow 0$. $\square$

We next extend our lower bounds on the contraction ratio of asymptotic consensus to lower bounds on the decision time of approximate consensus. In particular, we show optimality of the decision times of the algorithms presented by Charron-Bost et al. [18]: For $n = 2$, running Algorithm 1 and deciding y^i after $\lceil \log_3 \frac{\Delta}{\varepsilon} \rceil$ rounds is optimal (Theorem 9.2). For $n \geq 3$ and the oblivious message adversary of all non-split graphs, running the midpoint algorithm and deciding after $\lceil \log_2 \frac{\Delta}{\varepsilon} \rceil$ rounds is optimal (Theorem 9.3). For $n \geq 4$ and the strongest oblivious message adversary of all rooted graphs, running the amortized midpoint algorithm and deciding after $(n-1)\lceil \log_2 \frac{\Delta}{\varepsilon} \rceil$ rounds is optimal within a multiplicative term of at most $\frac{n-1}{n-2}$ (Theorem 9.4).

We start with the case of two agents in Theorem 9.2. The proof is by reducing asymptotic consensus to approximate consensus, arriving at a contradiction with Theorem 4.1 for too fast approximate consensus algorithms.

THEOREM 9.2. *Let $\Delta > 0$ and $\varepsilon > 0$. Subject to an oblivious message adversary of $n = 2$ agents whose communication graphs are H_0, H_1 , and H_2 from Figure 3, all approximate consensus algorithms have an execution with initial diameter $\Delta(C_0) \leq \Delta$ and decision time greater or equal to $\log_3 \frac{\Delta}{\varepsilon}$.*

PROOF. Assume to the contrary that algorithm $\mathcal{A}$ solves approximate consensus subject to the oblivious message adversary $\mathcal{M} = \mathcal{M}(\{H_0, H_1, H_2\})$ that decides in $T < \log_3 \frac{\Delta}{\varepsilon}$ rounds for all initial configurations C_0 with $\Delta(C_0) \leq \Delta$ and some $\varepsilon > 0$.

Choose any C_0 with $\Delta(C_0) = \Delta$. Define algorithm $\tilde{\mathcal{A}}$ by running algorithm $\mathcal{A}$, updating y to the agents' decision values in round T , and then running Algorithm 1 with the initial values $y^i(T) = d^i$ from round $T+1$ on. Because Algorithm 1 is an asymptotic consensus algorithm and the decision values $y(T)$ of $\mathcal{A}$ satisfy the Validity condition of approximate consensus, algorithm $\tilde{\mathcal{A}}$ is an asymptotic consensus algorithm.

Let C_0 be an initial configuration of $\tilde{\mathcal{A}}$ with initial values $y(0)$. By the proof of Theorem 4.1, namely (5), there is an execution $E = C_0, G_1, C_1, G_2, \dots$ starting from C_0 such that

$$\delta_{\mathcal{M}, \tilde{\mathcal{A}}}(C_T) \geq \frac{1}{3^T} \cdot \delta_{\mathcal{M}, \tilde{\mathcal{A}}}(C_0) .$$

We have $\delta_{\mathcal{M}, \tilde{\mathcal{A}}}(C_0) = \Delta(C_0) = \Delta$ by Lemma 3.9 and $\delta_{\mathcal{M}, \tilde{\mathcal{A}}}(C_T) \leq \Delta(y(T)) \leq \varepsilon$ by Validity of Algorithm 1 and ε -Agreement of algorithm $\mathcal{A}$. But this means $T \geq \log_3 \frac{\Delta}{\varepsilon}$, a contradiction. $\square$

With a similar proof, we also get the lower bound for approximate consensus with $n \geq 3$ agents:

THEOREM 9.3. *Let $\Delta > 0$ and $\varepsilon > 0$. Subject to an oblivious message adversary of $n \geq 3$ agents that includes the communication graphs $\text{deaf}(G)$, all approximate consensus algorithms have an execution with initial diameter $\Delta(C_0) \leq \Delta$ and decision time greater or equal to $\log_2 \frac{\Delta}{\varepsilon}$.*

Analogously, for oblivious message adversaries with rooted Ψ graphs, using (11), we obtain:

THEOREM 9.4. *Let $\Delta > 0$ and $\varepsilon > 0$. Subject to an oblivious message adversary of $n \geq 4$ agents that includes the Ψ communication graphs, all approximate consensus algorithms have an execution with initial diameter $\Delta(C_0) \leq \Delta$ and decision time greater or equal to $(n-2)\log_2 \frac{\Delta}{\varepsilon}$.*

In case the oblivious message adversary does not include any of the above graphs, we obtain the following general bound on the termination time:

THEOREM 9.5. *Let $\Delta > 0$ and $\varepsilon > 0$. For an oblivious message adversary $\mathcal{M} = \mathcal{M}(\mathcal{G})$ subject to which exact consensus is not solvable, all approximate consensus algorithms have an execution with initial diameter $\Delta(C_0) \leq \Delta$ and decision time greater or equal to $\log_{D+1} \frac{\Delta}{\varepsilon n}$, where D is the $\alpha_{\mathcal{G}}$ -diameter of $\mathcal{G}$.*

PROOF. Assume to the contrary that algorithm $\mathcal{A}$ solves approximate consensus subject to some oblivious message adversary $\mathcal{M}$ subject to which exact consensus is not solvable and that decides in $T < \log_{D+1} \frac{\Delta}{\varepsilon n}$ rounds for all initial configurations C_0 with $\Delta(C_0) \leq \Delta$ and some $\varepsilon > 0$.

Define algorithm $\tilde{\mathcal{A}}$ by repeatedly running algorithm $\mathcal{A}$, updating y to the agents' decision values in round kT , and then restarting $\mathcal{A}$ in round $kT + 1$ with the decision values from the previous phase. Then $\tilde{\mathcal{A}}$ is an asymptotic consensus algorithm.

Let C_0 be an initial configuration of $\tilde{\mathcal{A}}$ with $\Delta(C_0) \leq \Delta$ and $\delta_{\mathcal{M}, \tilde{\mathcal{A}}}(C_0) \geq \Delta/n$. By the proof of Theorem 7.12, namely (16), there is an execution $E = C_0, G_1, C_1, G_2, \dots$ starting from C_0 such that

$$\delta_{\mathcal{M}, \tilde{\mathcal{A}}}(C_T) \geq \frac{1}{(D+1)^T} \cdot \delta_{\mathcal{M}, \tilde{\mathcal{A}}}(C_0). \quad (19)$$

It is $\delta_{\mathcal{M}, \tilde{\mathcal{A}}}(C_0) \leq \Delta(C_0) \leq \Delta/n$ and $\delta_{\mathcal{M}, \tilde{\mathcal{A}}}(C_T) \leq \Delta(y(T)) \leq \varepsilon$ by ε -Agreement of algorithm $\mathcal{A}$. But this means $T \geq \log_{D+1} \frac{\Delta}{\varepsilon n}$, a contradiction. $\square$

From Theorem 9.5 and the fact that $\mathcal{M}' \subseteq \mathcal{M}$ implies $\mathcal{E}_{\mathcal{M}', \mathcal{A}} \subseteq \mathcal{E}_{\mathcal{M}, \mathcal{A}}$, we get:

COROLLARY 9.6. *Let $\Delta > 0$ and $\varepsilon > 0$. For an oblivious message adversary $\mathcal{M} = \mathcal{M}(\mathcal{G})$ subject to which exact consensus is not solvable, all approximate consensus algorithms have an execution with initial diameter $\Delta(C_0) \leq \Delta$ and decision time greater or equal to $\log_{D+1} \frac{\Delta}{\varepsilon n}$, where D is the smallest $\alpha_{\mathcal{G}}$ -diameter of $\alpha_{\mathcal{G}}$ -connected $\mathcal{G}' \subseteq \mathcal{G}$ such that exact consensus is not solvable subject to $\mathcal{M}(\mathcal{G}')$.*

10 CONCLUSIONS

We introduced the notion of valency for asymptotic consensus algorithms, generalizing the concept of valency from exact consensus algorithms. Based on the study of valency diameters along executions we proved lower bounds on the contraction ratios of asymptotic consensus algorithm in arbitrary oblivious message adversaries: In particular, together with previously published convex-combination algorithms [18], we showed tight bounds for the oblivious message adversary containing all non-split graphs, and the strongest oblivious message adversary for which asymptotic consensus is solvable, the oblivious message adversary of all rooted graphs. Furthermore we obtained a general lower bound of $1/(D+1)$ for any oblivious message adversary for which exact consensus is not solvable; here D denotes the newly introduced $\alpha_{\mathcal{G}}$ -diameter of the set $\mathcal{G}$ of communication graphs. Interestingly, this result also immediately provides new tight lower bounds on classical static failure models, as exemplified in the case of asynchronous message-passing systems with crashes and shows a fundamental discrepancy in performance between round-based and general algorithms. We finally demonstrated how to obtain corresponding results for approximate consensus algorithms.

ACKNOWLEDGMENTS

We would like to thank Bernadette Charron-Bost and the anonymous reviewers, whose comments greatly improved the paper. The research was partially funded by the Austrian Science Fund (FWF) projects SIC (P26436) and ADynNet (P28182), the Centre National de la Recherche Scientifique (CNRS) project PEPS DEMO, and the DigiCosme working group HicDiesMeus.

REFERENCES

- [1] Ittai Abraham, Yonatan Amit, and Danny Dolev. Optimal resilience asynchronous approximate agreement. In Teruo Higashino, editor, *Proceedings of the 8th International Conference On Principles Of Distributed Systems (OPODIS 2004)*, pages 229–239, Heidelberg, 2004. Springer.
- [2] Manuel Alcántara, Armando Castañeda, David Flores-Peñaloza, and Sergio Rajsbaum. The topology of look-compute-move robot wait-free algorithms with hard termination. *Distributed Computing*, 32:235–255, 2019.
- [3] Bowen Alpern and Fred B. Schneider. Defining liveness. *Information Processing Letters*, 21(4):181–185, 1985.
- [4] David Angeli and Pierre-Alexandre Bliman. Stability of leaderless discrete-time multi-agent systems. *Mathematics of Control, Signals, and Systems*, 18(4):293–322, 2006.
- [5] Karl Johan Åström and Richard M. Murray. *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press, Princeton, 2008.
- [6] J. A. Benediktsson and P. H. Swain. Consensus theoretic classification methods. *IEEE Transactions on Systems, Man, and Cybernetics*, 22(4):688–704, 1992.
- [7] Ofer Biran, Shlomo Moran, and Shmuel Zaks. A combinatorial characterization of the distributed 1-solvable tasks. *Journal of Algorithms*, 11(3):420–440, 1990.
- [8] Pierre-Alexandre Bliman, Angelia Nedic, and Asuman E. Ozdaglar. Rate of convergence for consensus with delays. In *Proceedings of the 47th IEEE Conference on Decision and Control, and the European Control Conference (CDC-ECC 2008)*, pages 2226–2231. IEEE, New York, 2008.
- [9] Nicolas Braud-Santoni, Swan Dubois, Mohamed-Hamza Kaaouachi, and Franck Petit. The next 700 impossibility results in time-varying graphs. *International Journal of Networking and Computing*, 6(1):27–41, 2016.
- [10] Ming Cao, A. Stephen Morse, and Brian D. O. Anderson. Reaching a consensus in a dynamically changing environment: convergence rates, measurement delays, and asynchronous events. *SIAM Journal on Control and Optimization*, 47(2):601–623, 2008.
- [11] Ming Cao, Daniel A. Spielman, and A. Stephen Morse. A lower bound on convergence of a distributed network consensus algorithm. In Hannes Frey, Xu Li, and Stefan Rührup, editors, *Proceedings of the 44th IEEE Conference on Decision and Control, and the European Control Conference (CDC-ECC 2005)*, pages 2356–2361, New York, 2005. IEEE.
- [12] Armando Castañeda, Pierre Fainaugaud, Ami Paz, Sergio Rajsbaum, Matthieu Roy, and Corentin Travers. A topological perspective on distributed network algorithms. *Theoretical Computer Science*, 2020. In press.
- [13] Armando Castañeda, Sergio Rajsbaum, and Matthieu Roy. Convergence and covering on graphs for wait-free robots. *Journal of the Brazilian Computer Society*, 24(1):1, 2018.
- [14] Arnaud Casteigts, Paola Flocchini, Walter Quattrociocchi, and Nicola Santoro. Time-varying graphs and dynamic networks. *International Journal of Parallel, Emergent and Distributed Systems*, 27(5):387–408, 2012.
- [15] Bernadette Charron-Bost, Matthias Függer, and Thomas Nowak. Approximate consensus in highly dynamic networks. *CoRR*, abs/1408.0620, 2014.
- [16] Bernadette Charron-Bost, Matthias Függer, and Thomas Nowak. Amortized averaging algorithms for approximate consensus, 2015. <http://arxiv.org/abs/1512.04222>.
- [17] Bernadette Charron-Bost, Matthias Függer, and Thomas Nowak. Approximate consensus in highly dynamic networks: the role of averaging algorithms. In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming (ICALP 2015)*, pages 528–539, Heidelberg, 2015. Springer.
- [18] Bernadette Charron-Bost, Matthias Függer, and Thomas Nowak. Fast, robust, quantizable approximate consensus. In Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi, editors, *Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, pages 137:1–137:14, Dagstuhl, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [19] Bernadette Charron-Bost and André Schiper. The Heard-Of model: computing in distributed systems with benign faults. *Distributed Computing*, 22(1):49–71, 2009.
- [20] Bernard Chazelle. The total s -energy of a multiagent system. *SIAM Journal on Control and Optimization*, 49(4):1680–1706, 2011.
- [21] Reuven Cohen and David Peleg. Convergence properties of the gravitational algorithm in asynchronous robot systems. *SIAM Journal on Computing*, 34(6):1516–1528, 2005.
- [22] Étienne Coulouma, Emmanuel Godard, and Joseph Peters. A characterization of oblivious message adversaries for which consensus is solvable. *Theoretical Computer Science*, 584:80–90, 2015.
- [23] George Cybenko. Dynamic load balancing for distributed memory multiprocessors. *Journal of Parallel and Distributed Computing*, 7(2):279–301, 1989.
- [24] Carole Delporte-Gallet, Hugues Fauconnier, and Sergio Rajsbaum. Communication complexity of wait-free computability in dynamic networks. In Andrea Richa and Christian Scheideler, editors, *Proceedings of the 27th International Colloquium on Structural Information and Communication Complexity (SIROCCO)*, volume 12156 of *Lecture Notes in*

Computer Science, pages 291–309. Springer, Heidelberg, 2020.

- [25] Danny Dolev, Nancy A. Lynch, Shlomit S. Pinter, Eugene W. Stark, and William E. Weihl. Reaching approximate agreement in the presence of faults. *Journal of the ACM*, 33(2):499–516, 1986.
- [26] Magnus Egerstedt and Xiaoming Hu. Formation constrained multi-agent control. *IEEE Transactions on Robotics and Automation*, 17(6):947–951, 2001.
- [27] Seyed Rasoul Etesami and Tamer Başar. Convergence time for unbiased quantized consensus. *IEEE Transactions on Automatic Control*, 61(2):443–455, 2016.
- [28] Alan D. Fekete. Asymptotically optimal algorithms for approximate agreement. *Distributed Computing*, 4(1):9–29, 1990.
- [29] Alan D. Fekete. Asynchronous approximate agreement. *Information and Computation*, 115(1):95–124, 1994.
- [30] Michael J Fischer, Nancy A Lynch, and Michael Merritt. Easy impossibility proofs for distributed consensus problems. *Distributed Computing*, 1(1):26–39, 1986.
- [31] Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM*, 32(2):374–382, 1985.
- [32] Matthias Függer and Thomas Nowak. Fast multidimensional asymptotic and approximate consensus. In *32nd International Symposium on Distributed Computing (DISC 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [33] Matthias Függer, Thomas Nowak, and Manfred Schwarz. Tight bounds for asymptotic and approximate consensus. In *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing - PODC '18*, pages 325–334, New York, NY, USA, 2018. ACM Press.
- [34] Eli Gafni. Round-by-round fault detectors (extended abstract): unifying synchrony and asynchrony. In *Proceedings of the Seventeenth Annual ACM Symposium on Principles of Distributed Computing*, pages 143–152, New York, NY, USA, 1998. ACM Press.
- [35] Eli Gafni, Petr Kuznetsov, and Ciprian Manolescu. A generalized asynchronous computability theorem. In *Proceedings of the 33rd ACM Symposium on Principles of Distributed Computing (PODC 2014)*, pages 222–231. ACM, New York, 2014.
- [36] Rainer Hegselmann and Ulrich Krause. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3):1–33, 2002.
- [37] Maurice Herlihy, Dmitry Kozlov, and Sergio Rajsbaum. *Distributed Computing Through Combinatorial Topology*. Morgan Kaufmann, Waltham, 2014.
- [38] Maurice Herlihy, Sergio Rajsbaum, Michel Raynal, and Julien Stainer. From wait-free to arbitrary concurrent solo executions in colorless distributed computing. *Theoretical Computer Science*, 683:1–21, 2017.
- [39] Gunnar Hoest and Nir Shavit. Toward a topological characterization of asynchronous complexity. *SIAM Journal on Computing*, 36(2):457–497, 2006.
- [40] Idit Keidar and Alex Shraer. Timeliness, failure detectors, and consensus performance. In *Proceedings of the twenty-fifth annual ACM SIGACT-SIGOPS symposium on Principles of Distributed Computing (PODC'06)*, pages 169–178, New York, NY, USA, 2006. ACM Press.
- [41] Fabian Kuhn and Rotem Oshman. Dynamic networks: Models and algorithms. *SIGACT News*, 42(1):82–96, 2011.
- [42] Qun Li and Daniela Rus. Global clock synchronization in sensor networks. *IEEE Transactions on Computers*, 55(2):214–226, 2006.
- [43] J. Lin, A.S. Morse, and B.D.O. Anderson. The multi-agent rendezvous problem. In Vijay Kumar, Naomi Leonard, and A. Stephen Morse, editors, *Cooperative Control: A Post-Workshop Volume, 2003 Block Island Workshop on Cooperative Control*, pages 257–289. Springer, Heidelberg, 2005.
- [44] Nancy A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, San Francisco, 1996.
- [45] Hammurabi Mendes, Maurice Herlihy, Nitin Vaidya, and Vijay K Garg. Multidimensional agreement in Byzantine systems. *Distributed Computing*, 28(6):423–441, 2015.
- [46] Renato E. Mirollo and Steven H. Strogatz. Synchronization of pulse-coupled biological oscillators. *SIAM Journal on Applied Mathematics*, 50(6):1645–1662, 1990.
- [47] Luc Moreau. Stability of multiagent systems with time-dependent communication links. *IEEE Transactions on Automatic Control*, 50(2):169–182, 2005.
- [48] Yoram Moses and Sergio Rajsbaum. A layered analysis of consensus. *SIAM J. Comput.*, 31(4):989–1021, 2002.
- [49] Angelia Nedic, Alexander Olshevsky, Asuman E. Ozdaglar, and John N. Tsitsiklis. On distributed averaging algorithms and quantization effects. *IEEE Transactions on Automatic Control*, 54(11):2506–2517, 2009.
- [50] Thomas Nowak. Topology in distributed computing. Master thesis, TU Wien, March 2010. <http://www.ub.tuwien.ac.at/dipl/2010/AC07807067.pdf>.
- [51] Thomas Nowak, Ulrich Schmid, and Kyrill Winkler. Topological characterization of consensus under general message adversaries. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing - PODC '19*, PODC '19, pages 218–227, New York, NY, USA, 2019. ACM Press.
- [52] Alex Olshevsky. Linear time average consensus on fixed graphs. *IFAC-PapersOnLine*, 48(22):94–99, 2015.

- [53] Alex Olshevsky and John N. Tsitsiklis. Convergence speed in distributed consensus and averaging. *SIAM Review*, 53(4):747–772, 2011.
- [54] Nicola Santoro and Peter Widmayer. Time is not a healer. In B. Monien and R. Cori, editors, *6th Symposium on Theoretical Aspects of Computer Science*, volume 349 of *LNCS*, pages 304–313. Springer, Heidelberg, 1989.
- [55] Ulrich Schmid. How to model link failures: A perception-based fault model. In *Proceedings of the International Conference on Dependable Systems and Networks (DSN’01)*, pages 57–66, Göteborg, Sweden, 2001.
- [56] Tamás Vicsek, András Czirók, Eshel Ben-Jacob, Inon Cohen, and Ofer Shochet. Novel type of phase transition in a system of self-driven particles. *Physical Review Letters*, 75(6):1226–1229, 1995.
- [57] Kyrill Winkler and Ulrich Schmid. An overview of recent results for consensus in directed dynamic networks. *Bulletin of the EACTS*, 128, 2019.
- [58] Y. Yuan, G.-B. Stan, L. Shi, M. Barahona, and J. Goncalves. Decentralized minimum-time consensus. *Automatica*, 49(5):1227–1235, 2013.

A AUXILIARY RESULTS

LEMMA A.1. *Let $m \geq 1$ and $A_1, \dots, A_m \subseteq \mathbb{R}^d$. Then $\text{diam} \left(\bigcup_{\ell=1}^m A_\ell \right) = \max_{1 \leq i, j \leq m} \text{diam} (A_i \cup A_j)$.*

PROOF. Let $D = \text{diam} \left(\bigcup_{\ell=1}^m A_\ell \right)$ and $D_{i,j} = \text{diam} (A_i \cup A_j)$ for $1 \leq i, j \leq m$. The set inclusion $A_i \cup A_j \subseteq \bigcup_{\ell=1}^m A_\ell$ implies the inequality $D_{i,j} \leq D$ and thus $\max_{i,j} D_{i,j} \leq D$.

To prove the converse direction, take a sequence of pairs (x_k, y_k) of points in $\bigcup_{\ell=1}^m A_\ell$ whose distance converges to the set’s diameter D . That is, $D = \lim_{k \rightarrow \infty} \|x_k - y_k\|$.

Since the set $\bigcup_{\ell=1}^m A_\ell$ is a union, for every k , there exists some index i_k such that $x_k \in A_{i_k}$. Because there are only finitely many possible values for the indices i_k , there is an infinite subsequence of the sequence of indices i_k that is constant, say, equal to i . The corresponding points x_k of this subsequence all lie in the set A_i .

Similarly, starting from this subsequence, because there are only finitely many possible values for the indices j_k , there is an infinite subsequence for which j_k is also constant, say equal to j . Again, the corresponding points y_k of this subsequence all lie in the set A_j .

The limit of distances $\|x_k - y_k\|$ in this second subsequence is also equal to D , as the original sequence. But, also, it is upper-bounded by $D_{i,j}$ since all x_k and y_k are in $A_i \cup A_j$. Hence $D \leq D_{i,j} \leq \max_{i,j} D_{i,j}$, which shows the converse inequality and concludes the proof. $\square$

Similarly to the metric on executions, we define a metric on the set of communication sequences by $\text{dist}(\sigma, \sigma') = 1/2^\theta$ where θ is the first index at which σ and σ' differ.

LEMMA A.2. *Let $\mathcal{G}$ be the set of communication graphs of an oblivious message adversary. Then the set $\mathcal{G}^\omega$ of communication sequences is compact.*

PROOF. It is straightforward to show that the metric dist induces the product topology on $\mathcal{G}^\omega$. The set $\mathcal{G}$ is compact since it is finite. Hence, by Tychonoff’s theorem, the product space $\mathcal{G}^\omega$ is compact as well. $\square$

LEMMA A.3. *Let $\mathcal{A}$ be an algorithm and let $\mathcal{M}$ be a message adversary. For every initial configuration C_0 , the scheduler function $S_{C_0} : \mathcal{M} \rightarrow \mathcal{E}_{\mathcal{M}, \mathcal{A}}$,*

$$S_{C_0}(G_1, G_2, \dots) = C_0, G_1, C_1, G_2, \dots \quad (20)$$

which maps communication sequences to executions is continuous.

PROOF. By definition of the metric, $\text{dist}(\sigma, \sigma') < 1/2^\theta$ means that the first θ communication graphs of σ and σ' coincide. But then, since the algorithm is deterministic, the first θ configurations of the executions $S_{C_0}(\sigma)$ and $S_{C_0}(\sigma')$ also coincide. Hence $\text{dist}(S_{C_0}(\sigma), S_{C_0}(\sigma')) < 1/2^\theta$.

This shows that $\text{dist}(S_{C_0}(\sigma), S_{C_0}(\sigma')) \leq \text{dist}(\sigma, \sigma')$. In particular, S_{C_0} is continuous. $\square$

LEMMA A.4. *Let $\mathcal{A}$ be an algorithm with a finite number of initial configurations. The set of executions of algorithm $\mathcal{A}$ subject to an oblivious message adversary is compact.*

PROOF. Let $\mathcal{G}$ be the set of communication graphs of the oblivious message adversary $\mathcal{M} = \mathcal{M}(\mathcal{G})$. For every initial configuration, the scheduler function $S_{C_0} : \mathcal{M} \rightarrow \mathcal{E}_{\mathcal{M}, \mathcal{A}}$ is continuous (Lemma A.3). Since the set $\mathcal{M} = \mathcal{G}^\omega$ is compact (Lemma A.2), and the image of a compact set under a continuous function is compact, the subset of $\mathcal{E}_{\mathcal{M}, \mathcal{A}}$ whose executions share the same initial configuration C_0 is compact. Now, since there are only finitely many initial configurations, and any finite union of compact sets is compact, the set $\mathcal{E}_{\mathcal{M}, \mathcal{A}}$ of all executions is also compact. $\square$

LEMMA A.5. *Let $\mathcal{M}$ be a message adversary and let $\mathcal{A}$ be an algorithm that solves exact consensus subject to $\mathcal{M}$. Then its consensus function $y : \mathcal{E}_{\mathcal{M}, \mathcal{A}} \rightarrow \mathcal{V}$ that maps executions to their common decision value is continuous.*

PROOF. The function y is even locally constant: Let $E \in \mathcal{E}_{\mathcal{M}, \mathcal{A}}$ be an execution. Denote by T a round number in which all agents have already decided in execution E , and set $\varepsilon = 2^{-T}$. Then, in every execution $E' \in \mathcal{E}_{\mathcal{M}, \mathcal{A}}$ with $\text{dist}(E, E') < \varepsilon$, all agents have decided at round T , and on the same value as in E . In other words, $y(E) = y(E')$. This shows that the consensus function y is locally constant, hence continuous. $\square$

B ADDITIONAL PROOFS

LEMMA 6.2. *Let C be a configuration reachable by algorithm $\mathcal{A}$ subject to message adversary $\mathcal{M}$ via the finite communication sequence π . Then*

$$Y_{\mathcal{M}, \mathcal{A}}^*(C) = \bigcup_{G \in \Sigma(\pi)} Y_{\mathcal{M}, \mathcal{A}}^*(C.G) .$$

PROOF. Assume that C is reachable by algorithm $\mathcal{A}$ subject to message adversary $\mathcal{M}$ via the finite communication sequence $\pi = G_1, G_2, \dots, G_t$.

Let $y^* \in Y_{\mathcal{M}, \mathcal{A}}^*(C)$. By the assumption on the configuration C and the definition of the valency $Y_{\mathcal{M}, \mathcal{A}}^*(C)$, there exists an execution $E = C_0, G_1, C_1, G_2, \dots$ in $\mathcal{E}_{\mathcal{M}, \mathcal{A}}$ such that $y^* = y^*(E)$ and $C = C_t$. Set $G = G_{t+1} \in \Sigma(\pi)$. Hence we have $C_{t+1} = C.G$. But this shows that $y^* \in Y_{\mathcal{M}, \mathcal{A}}^*(C.G)$ since $C.G$ occurs in execution E whose limit is y^* . This shows inclusion of the left-hand side in the right-hand side.

Now let $G \in \Sigma(\pi)$ and $y^* \in Y_{\mathcal{M}, \mathcal{A}}^*(C.G)$. Then there exists an execution $E = C_0, G_1, C_1, G_2, \dots$ in $\mathcal{E}_{\mathcal{M}, \mathcal{A}}$ such that $C_t = C$, $G_{t+1} = G$, and $C_{t+1} = C.G$. But then trivially $y^* \in Y_{\mathcal{M}, \mathcal{A}}^*(C)$ because C occurs in E . This shows inclusion of the right-hand side in the left-hand side and concludes the proof. $\square$

LEMMA 6.3. *Let C be a configuration reachable by algorithm $\mathcal{A}$ subject to message adversary $\mathcal{M}$ via the finite communication sequence π . Then there exist $G, H \in \Sigma(\pi)$ such that*

$$\delta_{\mathcal{M}, \mathcal{A}}(C) = \text{diam} \left(Y_{\mathcal{M}, \mathcal{A}}^*(C.G) \cup Y_{\mathcal{M}, \mathcal{A}}^*(C.H) \right) .$$

PROOF. By Lemma 6.2 it is $Y_{\mathcal{M}, \mathcal{A}}^*(C) = \bigcup_{G \in \Sigma(\pi)} Y_{\mathcal{M}, \mathcal{A}}^*(C.G)$. The claimed equality now follows from Lemma A.1. $\square$

LEMMA 6.4. *Let C and C' be two reachable configurations subject to message adversary $\mathcal{M}$ via the finite communication sequences π and π' , respectively. If $C \sim_i C'$ and there exist communication sequence α and α' such that $\pi \cdot \alpha \in \mathcal{M}$, $\pi' \cdot \alpha' \in \mathcal{M}$, and i is deaf in all communication graphs in α and α' , then $Y_{\mathcal{M}, \mathcal{A}}^*(C) \cap Y_{\mathcal{M}, \mathcal{A}}^*(C') \neq \emptyset$.*

PROOF. Let $E = C_0, G_1, C_1, G_2, \dots$ be an execution in $\mathcal{E}_{\mathcal{M}, \mathcal{A}}$ with communication sequence $\pi \cdot \alpha$ that includes configuration C , and let $E' = C'_0, G'_1, C'_1, G'_2, \dots$ be an execution in $\mathcal{E}_{\mathcal{M}, \mathcal{A}}$ with communication sequence $\pi' \cdot \alpha'$ and that includes configuration C' . Then there exists some time t such that $C_t = C$ and some time s such that $C'_s = C'$.

We show by induction that $C_{t+k} \sim_i C'_{s+k}$ for all $k \geq 0$. The base case $k = 0$ is true since $C \sim_i C'$ by assumption. For the induction step, assume that $C_{t+k} \sim_i C'_{s+k}$. Since agent i does not receive any messages in communication graphs G_{t+k+1} and G_{s+k+1} , we also have $C_{t+k+1} \sim_i C'_{s+k+1}$.

But then we necessarily have $y^*(E) = y^*(E')$ and thus $y^*(E) \in Y_{\mathcal{M}, \mathcal{A}}^*(C) \cap Y_{\mathcal{M}, \mathcal{A}}^*(C') \neq \emptyset$. $\square$

LEMMA 6.5. *Let $\Delta \geq 0$. If there exist agents $i \neq j$ and communication sequences $\alpha^{(i)}, \alpha^{(j)} \in \mathcal{M}$ such that agent i is deaf in $\alpha^{(i)}$ and agent j is deaf in $\alpha^{(j)}$, then there is an initial configuration C_0 with $\delta_{\mathcal{M}, \mathcal{A}}(C_0) = \Delta$. In particular, there is an initial configuration for which $\delta_{\mathcal{M}, \mathcal{A}}(C_0) > 0$.*

PROOF. Let x and y be any two points at distance $\|x - y\| = \Delta$. Let C_0 be an initial configuration in which all agents except i have initial value x and agent j has initial value y . Also, denote by $C_0^{(x)}$ an initial configuration in which all agents have initial value x , and by $C_0^{(y)}$ one in which all agents have initial value y .

Since agent i cannot distinguish the two executions with communication sequence $\alpha^{(i)}$ and respective initial configurations C_0 and $C_0^{(x)}$, we necessarily have $x \in Y_{\mathcal{M}, \mathcal{A}}^*(C_0)$ by the Validity condition. Similarly, we have $y \in Y_{\mathcal{M}, \mathcal{A}}^*(C_0)$. Also by the Validity condition, the set $Y_{\mathcal{M}, \mathcal{A}}^*(C)$ is contained in the convex hull of the points x and y . Hence, $\delta_{\mathcal{M}, \mathcal{A}}(C_0) = \text{diam}\left(Y_{\mathcal{M}, \mathcal{A}}^*(C_0)\right) = \|x - y\| = \Delta$. $\square$