



HAL
open science

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. ACM Symposium on Principles of Distributed Computing (PODC'18), Jul 2018, Egham, United Kingdom. 10.1145/3212734.3212762 . hal-01799039

HAL Id: hal-01799039

<https://hal.science/hal-01799039v1>

Submitted on 24 May 2018

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, LSV, ENS Paris-Saclay, Inria
mfuegger@lsv.fr

Thomas Nowak
Université Paris-Sud
thomas.nowak@lri.fr

Manfred Schwarz
ECS, TU Wien
mschwarz@ecs.tuwien.ac.at

ABSTRACT

We study the performance of asymptotic and approximate consensus algorithms under harsh environmental conditions. The asymptotic consensus problem requires a set of agents to repeatedly set their outputs such that the outputs converge to a common value within the convex hull of initial values. This problem, and the related approximate consensus problem, are fundamental building blocks in distributed systems where exact consensus among agents is not required or possible, e.g., man-made distributed control systems, and have applications in the analysis of natural distributed systems, such as flocking and opinion dynamics. We prove tight lower bounds on the contraction rates of asymptotic consensus algorithms in dynamic networks, 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 in [Charron-Bost et al., ICALP'16] for certain dynamic networks, including the weakest dynamic network model 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 our 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**;

KEYWORDS

asymptotic consensus; approximate consensus; lower bounds; dynamic networks

ACM Reference Format:

Matthias Függer, Thomas Nowak, and Manfred Schwarz. 2018. Tight Bounds for Asymptotic and Approximate Consensus. In *PODC '18: ACM Symposium on Principles of Distributed Computing, July 23–27, 2018, Egham, United Kingdom*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3212734.3212762>

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).

PODC '18, July 23–27, 2018, Egham, United Kingdom

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5795-1/18/07.

<https://doi.org/10.1145/3212734.3212762>

1 INTRODUCTION

In the *asymptotic consensus* problem a set of agents, each starting from an initial value in \mathbb{R}^d , update their values such that all agents' values converge to a common value 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. The latter is weaker than the *exact consensus* problem in which agents need to decide on the same value. 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 [4], clock synchronization [21], formation control [15], rendezvous in space [22], or load balancing [13], but also for analyzing natural systems like flocking [31], firefly synchronization [26], or opinion dynamics [20]. 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 asymptotic and approximate consensus algorithms under such harsh conditions. Specifically, we study algorithms in a *network model* \mathcal{N} with round-based computation and a dynamic communication topology whose *directed* communication graphs are chosen each round from a predefined set \mathcal{N} of communication graphs. While this model naturally captures highly unstable communication topologies, we later on show that it also allows to assess performance within classical, more stable, distributed fault models.

Solvability and Algorithms. In previous work [8], Charron-Bost et al. showed that asymptotic consensus is solvable precisely within *rooted network models* in which all communication graphs contain rooted spanning trees. These rooted spanning trees need not have any edges in common and can change from one round to the next.

An interesting special case of rooted network models are network models 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) In [8], Charron-Bost et al. showed that non-split graphs also play a central role in arbitrary rooted network models: they showed that any product of $n - 1$ rooted graphs with n nodes is non-split, allowing to transform asymptotic consensus algorithms for non-split network models into their *amortized* variants for rooted models.

Interestingly, solvability in any rooted network model is already provided by deceptively simple algorithms [8]: 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. One instance, presented by Charron-Bost et al. [9] 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.

Regarding time complexity, for dimension $d = 1$, the amortized midpoint algorithm was shown to have a contraction rate (of the range of reachable values; see Section 3 for a formal definition) of $\frac{1}{\sqrt[n]{1/2}}$ in arbitrary rooted network models with n agents, and the midpoint algorithm of $\frac{1}{2}$ in non-split network models [9]. The latter is *optimal* for “memoryless” averaging algorithms, which only depend on the values received in the current round [9].

A natural question is whether *non-averaging 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 rates. In the context of classical failure models, deriving lower bounds independent of such assumptions, is a long-standing open problem raised by Dolev et al. in [14]. 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 algorithm 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., [3].

Contribution. We prove asymptotically tight lower bounds on the contraction rate of any asymptotic consensus algorithm 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 averaging algorithms. In particular, the following lower bounds hold for a network model \mathcal{N} with n agents: If exact consensus is solvable in \mathcal{N} , an optimal contraction rate of 0 can be achieved. Otherwise:

In a system with $n = 2$ agents, the contraction rate is lower bounded by $1/3$ (Theorem 4.1). This is tight [9].

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{N} contains $\text{deaf}(G)$, then the contraction rate is lower bounded by $1/2$ (Theorem 5.1). This is tight in non-split network models because of the midpoint algorithm [9].

We then show that if \mathcal{N} contains certain rooted graphs Ψ , the contraction rate is lower bounded by $\frac{1}{\sqrt[n]{1/2}}$ (Theorem 6.1). This is asymptotically tight in rooted network models because of the amortized midpoint algorithm [9]. Specifically, this proves optimality of the amortized midpoint algorithm for the *weakest*, i.e., largest, network model in which asymptotic and approximate consensus is solvable: the set of all directed rooted communication graphs.

For arbitrary network models we show that in a system with $n \geq 3$ agents, any asymptotic consensus algorithm must have a contraction rate of at least $1/(D + 1)$, where D , the so-called α -diameter of \mathcal{N} , i.e., the smallest value which allows a connection of

any pair of communication graphs in \mathcal{N} via an indistinguishability chain of length at most D (Theorem 7.4).

We demonstrate how to apply the above mentioned bound to obtain new lower bounds on contraction rates 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 [10, 14, 24]: 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.

We show that no algorithm operating in asynchronous rounds can achieve a contraction rate 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. [14] and for systems of size $n > 2f$ with up to f crashes by Fekete [18] have asymptotically optimal contraction rates for round-based algorithms.

We then present an algorithm for $n > f$ that does *not operate in asynchronous rounds* and achieves a contraction rate of 0, demonstrating a large gap between round-based and non round-based algorithms for asymptotic consensus.

Table 1 summarizes lower and upper bounds. Central to the proofs is the concept of the *valency of a configuration* of an asymptotic consensus algorithm, defined as the set of limits reachable from this configuration. By studying the changes in valency along executions, we infer bounds on the contraction rate.

We extend the above results on contraction rates to derive new lower bounds on the decision time of any approximate consensus algorithm: Let $\Delta > 0$ be the largest distance between initial values. For $n = 2$ we obtain $\lceil \log_3 \frac{\Delta}{\epsilon} \rceil$ (Theorem 9.1). For $n \geq 3$ and models that include $\text{deaf}(G)$ for a communication graph G , we show $\lceil \log_2 \frac{\Delta}{\epsilon} \rceil$ (Theorem 9.2), and for $n \geq 4$ and models that include certain Ψ graphs, we obtain $(n - 2) \lceil \log_2 \frac{\Delta}{\epsilon} \rceil$ (Theorem 9.3). For arbitrary network models in which exact consensus is not solvable, we show $\log_{D+1} \frac{\Delta}{\epsilon n}$ (Theorem 9.4). Again, deciding versions of the asymptotic consensus algorithms from [9] have matching time complexities; showing optimality of these algorithms also for solving approximate consensus.

Related work. The problem of asymptotic consensus in dynamic networks has been extensively studied in distributed computing and control theory, see, e.g., [2, 5, 6, 11, 16, 27]. The question of guaranteed convergence rates 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, e.g., in [6].

Olshevsky and Tsitsiklis [30], 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 in [28]. Chazelle [11] proposed an averaging algorithm with polynomial convergence time, which works in any bidirectional connected network model.

To speed up convergence times, algorithms where agents set their output based on values that have been received in rounds

agents	network model			asynchronous + f crashes	
	general non-split	non-split with α -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 $\left[\frac{1}{D+1}^*, \frac{1}{2}\right]$	$\left[n^{-2}\sqrt{\frac{1}{2}}^*, n^{-1}\sqrt{\frac{1}{2}}\right]$	$\left[\frac{1}{\lceil n/f \rceil + 1}^*, \frac{1}{\lceil n/f \rceil - 1}\right]$	

Table 1: Summary of lower and upper bounds on contraction rates. New bounds proved in this work are marked with an $*$. The three left columns are worst-case contraction rates for the case the network model is (i) a general non-split, (ii) a non-split network model with α -diameter D , and (iii) a general rooted network model. For (ii) contraction rates are 0 iff exact consensus is solvable. The right two columns summarize the bounds for the classical model of an asynchronous system with crashes.

prior to the previous round have also been considered in literature: Olshevsky [29] proposed a linear convergence time algorithm that uses messages from two rounds, however, being restricted to fixed bidirectional communication graphs. In [32], a linear convergence time algorithm for a possibly non-bidirectional fixed topology was proposed. It requires storing all received values. In previous work [9], Charron-Bost et al. presented the midpoint algorithm, which has constant convergence time in non-split network models and the amortized midpoint algorithm with linear convergence time in rooted network models.

To the best of our knowledge, the only lower bound on convergence rate in dynamic networks has been shown in [7]: the authors proved that the convergence rate of a specific averaging algorithm in a non-split network model with n agents is at least $1 - \frac{1}{n}$.

In the context of classical distributed computing failure scenarios, Dolev et al. [14] studied the related 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 [19]. The latter result was improved to $n \geq 3f + 1$ in [1]. Both papers also address the question of optimal contraction rate in such systems. Since, however, in synchronous systems with $n \geq 3f + 1$ exact consensus is solvable, leading to a contraction rate of 0, the authors consider bounds for round-by-round contraction rates. In [14] they showed that the achieved round-by-round contraction rate 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. [25] proposed algorithms with convergence time of $d \cdot \lceil \log_2 \frac{\sqrt{d}\Delta}{\epsilon} \rceil$ under the optimal resiliency condition $n \geq f \cdot (d + 2) + 1$.

Fekete [17] also studied round-by-round contraction rates for several failure scenarios, again, all 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. In [18], Fekete presented an algorithm for asynchronous message-passing systems with minority of crashes, also proving a tight lower bound on the contraction rate of any algorithm operating in asynchronous rounds for such systems.

2 DYNAMIC SYSTEM MODEL

We consider a set $[n] = \{1, \dots, n\}$ of n agents (also classically called processes). We assume a distributed, round-based computational model in the spirit of the Heard-Of model [10]. Computation proceeds in *rounds*: In every round, each agent sends its state 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 .

Communications that occur in a round are modeled by a *directed graph* with a node for each agent. 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 two edges in G and H , respectively.

We fix a nonempty set of communication graphs \mathcal{N} which determines the *network model*. To fully model dynamic networks in which topology may change continually and unpredictably, the communication graph at each round is chosen arbitrarily among \mathcal{N} . Thus we form the infinite sequences of graphs in \mathcal{N} which we call *communication patterns in \mathcal{N}* . In each communication pattern, 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 .

Let us fix an algorithm \mathcal{A} ; a *configuration* is a collection of n agent states, one per agent. We assume that all agents pick their initial state from the same set of states. Obviously the picks can be different for different agents. 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 $G.C$ if no confusion can arise. Then the *execution E of \mathcal{A}* from the initial configuration C_0 and with the communication pattern $(G_t)_{t \geq 1}$ 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 , $C_t = G_t.C_{t-1}$. The set of executions with communication patterns in \mathcal{N} , denoted $\mathcal{E}_{\mathcal{A}}^{\mathcal{N}}$, with the distance $\text{dist}(E, E') = 1/2^\theta$, where θ is the first index at which E and E' differ, is a compact metric space (e.g., see [23]).

Finally, any configuration that occurs in some execution with a communication pattern in \mathcal{N} is said to be *reachable from C_0 by \mathcal{A} in \mathcal{N}* . In the sequel, the algorithm and the network model 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 . Then we let $y_E(t) = (y_E^1(t), \dots, y_E^n(t))$. We write

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

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

We say an algorithm *solves the asymptotic consensus problem* in a network model \mathcal{N} if the following holds for every execution E with a communication pattern in \mathcal{N} :

- *Convergence.* Each sequence $(y_E^i(t))_{t \geq 0}$ converges.
- *Agreement.* If $y_E^i(t)$ and $y_E^j(t)$ converge, then they have a common limit.
- *Validity.* If $y_E^i(t)$ converges, then its limit is in the convex hull of the initial values $y_E^1(0), \dots, y_E^n(0)$.

Observe that the *consensus function* defined by $y^* : E \in (\mathcal{E}, \text{dist}) \mapsto y_E^* \in (\mathbb{R}^d, \|\cdot\|)$, where y_E^* denotes the common limit of the n sequences $(y_E^i(t))_{t \geq 0}$, is a priori not continuous. And indeed, there exist asymptotic consensus algorithms whose consensus functions are not continuous.

3 VALENCY AND CONTRACTION RATE

We now extend the notion of valency for a consensus algorithm to asymptotic consensus algorithms. We fix an asymptotic consensus algorithm \mathcal{A} that solves d -dimensional asymptotic consensus in a certain network model \mathcal{N} with $n \geq 2$ agents. Let C be a configuration reachable by \mathcal{A} in \mathcal{N} . Then we define the *valency* of C by $Y_{\mathcal{N}, \mathcal{A}}^*(C) = \{y_E^* \in \mathbb{R}^d \mid C \text{ occurs in } E \in \mathcal{E}_{\mathcal{A}}^{\mathcal{N}}\}$.

If the algorithm \mathcal{A} is clear from the context, we skip it from the subscript. 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 $\delta_{\mathcal{N}}(C) = \text{diam}(Y_{\mathcal{N}}^*(C))$ the diameter of the set of reachable limits starting from configuration C .

We have $\delta_{\mathcal{N}}(C_t) \rightarrow 0$ in any execution $E = G_0, C_1, G_1, C_2, \dots$ by Convergence and Agreement. To study the speed of convergence, we introduce the *contraction rate* of algorithm \mathcal{A} in network model \mathcal{N} as

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

where $E = C_0, G_1, C_1, G_2, \dots$. In particular, any algorithm that guarantees $\delta_{\mathcal{N}}(C_t) \leq \gamma^t \delta_{\mathcal{N}}(C_0)$ for all $t \geq 0$ has a contraction rate of at most γ .

We obtain the following for subsets of network models:

LEMMA 3.1. *Let $\mathcal{N}, \mathcal{N}'$ be two network models with $\mathcal{N}' \subseteq \mathcal{N}$. If \mathcal{A} is an algorithm that solves asymptotic consensus in \mathcal{N} , then (i) it*

also solves asymptotic consensus in \mathcal{N}' , (ii) for every configuration C reachable by \mathcal{A} in \mathcal{N}' , we have $Y_{\mathcal{N}'}^(C) \subseteq Y_{\mathcal{N}}^*(C)$, (iii) $\delta_{\mathcal{N}'}(C) \leq \delta_{\mathcal{N}}(C)$, and (iv) the contraction rate in \mathcal{N}' is less or equal to the contraction rate in \mathcal{N} .*

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

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

which concludes the proof. \square

We establish two branching properties of valency of configurations in execution trees.

LEMMA 3.2. *Let C be a configuration reachable by algorithm \mathcal{A} in network model \mathcal{N} . Then*

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

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

Now let $G \in \mathcal{N}$ and $y^* \in Y_{\mathcal{N}}^*(G.C)$. Then there is an execution $E = C_0, G_1, C_1, G_2, \dots$ in $\mathcal{E}_{\mathcal{A}}^{\mathcal{N}}$ and a $t \geq 0$ such that $y^* = y_E^*$ and $G.C = C_t$. Since C is a reachable configuration, there exists an execution $E' = C'_0, G'_1, C'_1, G'_2, \dots$ in $\mathcal{E}_{\mathcal{A}}^{\mathcal{N}}$ 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 $\mathcal{E}_{\mathcal{A}}^{\mathcal{N}}$ with $y_{E''}^* = y_E^* = y^*$. Hence $y^* \in Y_{\mathcal{N}}^*(C)$ because C occurs in E'' . This shows inclusion of the right-hand side in the left-hand side and concludes the proof. \square

LEMMA 3.3. *Let C be a configuration reachable by algorithm \mathcal{A} in network model \mathcal{N} . Then there exist $G, H \in \mathcal{N}$ such that*

$$\text{diam}(Y_{\mathcal{N}}^*(C)) = \text{diam}(Y_{\mathcal{N}}^*(G.C) \cup Y_{\mathcal{N}}^*(H.C)).$$

PROOF. Set $Y = Y_{\mathcal{N}}^*(C)$, and $Y_G = Y_{\mathcal{N}}^*(G.C)$ for $G \in \mathcal{N}$. By Lemma 3.2 it is $Y = \bigcup_{G \in \mathcal{N}} Y_G$, which means that every sequence of pairs of points in Y whose distances converge to $\text{diam}(Y)$ includes an infinite subsequence in some product $Y_G \times Y_H$ because there are only finitely many. Thus $\text{diam}(Y) \leq \text{diam}(Y_G \cup Y_H)$. The other inequality follows from $Y_G \cup Y_H \subseteq Y$. \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 the above definition, we obtain:

LEMMA 3.4. *Let C and C' be two reachable configurations, and let G and G' be communication graphs in \mathcal{N} . 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 $G.C \sim_i G'.C'$.*

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

LEMMA 3.5. *If agent i has the same in-neighbors in two communication graphs G and G' in \mathcal{N} , and if there exists a communication graph in \mathcal{N} in which i is deaf and $C \sim_j C'$ for the in-neighbors j of i , then $Y_{\mathcal{N}}^*(G.C) \cap Y_{\mathcal{N}}^*(G'.C') \neq \emptyset$.*

PROOF. From Lemma 3.4, we have $G.C \sim_i G'.C'$.

Let D_i be a communication graph in \mathcal{N} 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 3.4, 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 $Y_{\mathcal{N}}^*(G.C)$ and $Y_{\mathcal{N}}^*(G'.C')$ intersect. \square

From Lemma 3.5 we determine the valency of any initial configuration when the network model contains certain communication graphs. If every agent is deaf in some communication graph of the network model \mathcal{N} , 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.

LEMMA 3.6. *If, for every agent i , there is a communication graph in \mathcal{N} in which i is deaf, then each initial configuration C_0 satisfies $\delta_{\mathcal{N}}(C_0) = \Delta(y(0))$. In particular, there is an initial configuration for which $\delta_{\mathcal{N}}(C_0) > 0$.*

PROOF. Since $Y_{\mathcal{N}}^*(C_0)$ 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 $\Delta(y(0))$, we have the inequality $\delta_{\mathcal{N}}(C_0) \leq \Delta(y(0))$.

To show the converse inequality, let i and j be two agents such that $\|y^i(0) - y^j(0)\| = \Delta(y(0))$. 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 pattern as in E .

By a repeated application of Lemma 3.4, 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 $y^*(E^{(i)}) = y^i(0)$. It then follows that $y^i(0) \in Y_{\mathcal{N}}^*(C_0)$. By a similar argument, we see $y^j(0) \in Y_{\mathcal{N}}^*(C_0)$. Hence

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

which concludes the proof. \square

4 TIGHT BOUND FOR TWO AGENTS

In this section, we prove a lower bound of $1/3$ on the contraction rate of algorithms that solve asymptotic consensus in the network model of all rooted (and here also non-split) communication graphs with two agents. Combined with Algorithm 1, which achieves this




Figure 1: The rooted communication graphs H_0 , H_1 , and H_2 for $n = 2$

lower bound [9], we have indeed identified a tight bound on the contraction rate for $n = 2$. Moreover, the algorithm also shows that the lower bound is achieved by a simple convex combination algorithm.

Algorithm 1 Algorithm with contraction rate $1/3$ for $n = 2$

Initialization:

1: $y^i \in \mathbb{R}$

In round $t \geq 1$ do:

2: send y^i to other agent

3: **if** y^j was received from other agent **then**

4: $y^i \leftarrow y^i/3 + 2y^j/3$

5: **end if**

A straightforward analysis of Algorithm 1 shows that its contraction rate is equal to $1/3$.

Note that for $n = 2$, there are 3 possible rooted communication graphs that may occur, all of which are non-split; see Figure 1: (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.

THEOREM 4.1. *The contraction rate of any asymptotic consensus algorithm for $n = 2$ agents in a network model that includes the three 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

$$\delta_{\mathcal{N}}(C_t) \geq \frac{1}{3^t} \cdot \delta_{\mathcal{N}}(C_0) \quad (1)$$

for $t \geq 0$. This, applied to an initial configuration with $\delta_{\mathcal{N}}(C_0) > 0$, which exists by Lemma 3.6, then shows the theorem.

Note that it suffices to show (1) for the specific network model $\mathcal{N}' = \{H_0, H_1, H_2\}$ shown in Figure 1 because $\delta_{\mathcal{N}}(C_t) \geq \delta_{\mathcal{N}'}(C_t)$ by Lemma 3.1 and $\delta_{\mathcal{N}'}(C_0) = \delta_{\mathcal{N}}(C_0)$ by Lemma 3.6 whenever $\mathcal{N} \supseteq \mathcal{N}'$. We hence suppose $\mathcal{N} = \mathcal{N}'$ 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 (1). Equation (1) is trivial for $t = 0$.

Now assume $t \geq 0$ and that Equation (1) 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{N}' . Set $C_{t+1}^k = H_k.C_t$. Further let $Y = Y_{\mathcal{N}'}^*(C_t)$, and $Y_k = Y_{\mathcal{N}'}^*(C_{t+1}^k)$.

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

$$\delta_{N'}(C_{t+1}) \geq \delta_{N'}(C_t)/3 \geq \delta_{N'}(C_0)/3^{t+1} ,$$

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

Assume by contradiction that $\text{diam}(Y_{\hat{k}}) < \text{diam}(Y)/3$ for all $k \in \{0, 1, 2\}$. From Lemma 3.2 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 3.5 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 MODEL: CONTRACTION IN PRESENCE OF DEAF GRAPHS

In this section, we prove a lower bound of $1/2$ on the contraction rate of asymptotic consensus algorithms for $n \geq 3$ agents, in a network model that includes graphs derived from a communication graph G , where agents are made deaf in the derived graphs. As a special case this includes the network model of all non-split communication graphs. Charron-Bost et al. [9] presented the midpoint algorithm (given in Algorithm 2) for dimension one with contraction rate $1/2$ for non-split communication graphs. Together this shows tightness of our lower bound in dimension one.

Algorithm 2 Midpoint algorithm

Initialization:

1: $y^i \in \mathbb{R}$

In round $t \geq 1$ do:

2: send y^i to all agents

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

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

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

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) : let $\text{deaf}(G) = \{F_1, \dots, F_n\}$ with $F_i = G \setminus \{(j, i) : j \in [n] \setminus \{i\}\}$.

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 rate of any asymptotic consensus algorithm for $n \geq 3$ agents in a network model that includes $\text{deaf}(G)$ is greater or equal to $1/2$.*




Figure 2: Rooted communication graph Ψ_i for $n = 6$

Note that the network model $\text{deaf}(K_n)$, where K_n is the complete digraph on n nodes, is a subset of the network model that contains all non-split communication graphs. Hence the lower bound holds and, since Algorithm 2 is applicable, a tight bound 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 $i, j, l \in [n]$.

6 TIGHT BOUND FOR ROOTED MODEL: CONTRACTION IN PRESENCE OF Ψ GRAPHS

We next prove a lower bound of $n^{-2/3}/2$ on the contraction rate of asymptotic consensus algorithms for $n \geq 4$ agents.

For $i \in \{1, 2, 3\}$, let Ψ_i (see Figure 2) 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 sequence of graphs Ψ_i of length $n - 2$.

First observe that any communication pattern arising from the concatenation of σ_i sequences necessarily is a communication pattern of the network model of Ψ_i graphs, which are rooted. The analysis of the set of these communication patterns necessitates a generalization of our system model: generalizing from sets of allowed graphs to arbitrary sets of allowed communication patterns.

THEOREM 6.1. *The contraction rate of any asymptotic consensus algorithm in a network model including the Ψ graphs is greater or equal to $n^{-2/3}/2$.*

From [9] we have that the amortized midpoint algorithm guarantees a contraction of $n^{-2/3}/2$ for rooted network models. Theorem 6.1 shows that this is asymptotically optimal.

6.1 From Network Models to Sequences

To prove Theorem 6.1, we generalize the system model from Section 2 and some of the basic lemmas we proved for the specific case of network models. While we previously allowed the adversary to choose any sequence of communication graphs from the network model, we next consider more general properties on graph sequences, including safety and liveness properties.

A *property* is a set of communication patterns. A *snapshot* is a pair $S = (C, \pi)$ where C is a configuration, i.e., a collection of states of the agents, and π is a finite sequence of communication graphs. Given a snapshot $S = (C, \pi)$ and a communication graph G , define $G.S = (G.C, \pi \cdot G)$ where $\pi \cdot G$ is the addition of G to the end of π . We extend this definition to a finite sequence σ of communication

graphs. We write $S \sim_i S'$ if agent i has the same local state in both S and S' .

A *trace* of an algorithm \mathcal{A} in a property \mathcal{P} is an infinite sequence $T = (S_0, S_1, \dots)$ of snapshots such that there exists a communication pattern $P \in \mathcal{P}$ with $S_t = G_t \cdot S_{t-1}$ for all $t \geq 1$. We denote by $\mathcal{T}_{\mathcal{A}}^{\mathcal{P}}$ the set of all traces of \mathcal{A} in \mathcal{P} . If \mathcal{A} solves asymptotic consensus in \mathcal{P} , then we write y_T^* for the common limit of the agents' values in trace $T \in \mathcal{T}_{\mathcal{A}}^{\mathcal{P}}$.

We define the valency of snapshots and the contraction rate of an algorithm in \mathcal{P} analogously to the case of network models as

$$Y_{\mathcal{P}}^*(S) = \left\{ y_T^* \in \mathbb{R}^d \mid S \text{ occurs in } T \in \mathcal{T}_{\mathcal{A}}^{\mathcal{P}} \right\}$$

and the contraction rate as

$$\sup_{T \in \mathcal{T}_{\mathcal{A}}^{\mathcal{P}}} \limsup_{t \rightarrow \infty} \sqrt[t]{\delta_{\mathcal{P}}(S_t)}$$

where $\delta_{\mathcal{P}}(S_t) = \text{diam}(Y_{\mathcal{P}}^*(S_t))$.

LEMMA 6.2. *Let $\mathcal{P}, \mathcal{P}'$ be two properties with $\mathcal{P}' \subseteq \mathcal{P}$. If \mathcal{A} is an algorithm that solves asymptotic consensus in \mathcal{P} , then (i) it also solves asymptotic consensus in \mathcal{P}' , (ii) for every snapshot S reachable by \mathcal{A} in \mathcal{P}' , we have $Y_{\mathcal{P}'}^*(S) \subseteq Y_{\mathcal{P}}^*(S)$, (iii) $\delta_{\mathcal{P}'}(S) \leq \delta_{\mathcal{P}}(S)$, and (iv) the contraction rate in \mathcal{P}' is less or equal to the contraction rate in \mathcal{P} .*

For a snapshot $S = (C, \pi)$ reachable by algorithm \mathcal{A} in property \mathcal{P} , we define $\Sigma(S)$ to be the set of communication graphs G such that $\pi \cdot G$ is a prefix of a communication pattern in \mathcal{P} .

LEMMA 6.3. *Let S be a snapshot reachable by algorithm \mathcal{A} in property \mathcal{P} . Then*

$$Y_{\mathcal{P}}^*(S) = \bigcup_{G \in \Sigma(S)} Y_{\mathcal{P}}^*(G.S) .$$

LEMMA 6.4. *Let S be a configuration reachable by algorithm \mathcal{A} in property \mathcal{P} . Then there exist $G, H \in \Sigma(S)$ such that*

$$\text{diam}(Y_{\mathcal{P}}^*(S)) = \text{diam}(Y_{\mathcal{P}}^*(G.S) \cup Y_{\mathcal{P}}^*(H.S)) .$$

LEMMA 6.5. *Let $S = (C, \pi)$ and $S' = (C', \pi')$ be two snapshots with $S \sim_i S'$. If there exist sequences of communication graphs α and α' such that $\pi \cdot \alpha \in \mathcal{P}$, $\pi' \cdot \alpha' \in \mathcal{P}$ and i is deaf in all communication graphs in α and α' , then $Y_{\mathcal{P}}^*(S) \cap Y_{\mathcal{P}}^*(S') \neq \emptyset$.*

LEMMA 6.6. *Let $\Delta \geq 0$. If there exist agents $i \neq j$ and communication patterns $P_i, P_j \in \mathcal{P}$ such that agent i is deaf in P_i and agent j is deaf in P_j , then there is an initial snapshot S_0 with $\delta_{\mathcal{P}}(S_0) = \Delta$. In particular, there is an initial snapshot for which $\delta_{\mathcal{P}}(S_0) > 0$.*

6.2 Proof of Theorem 6.1

LEMMA 6.7. *For $i, j, \ell \in \{1, 2, 3\}$ with $\ell \neq i, j$: $\sigma_i \cdot C_t \sim_{\ell} \sigma_j \cdot C_t$.*

PROOF. We inductively show the following stronger statement. Let σ_i^k be the sequence of graphs 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 $\sigma_i^k \cdot C_t \sim_{\ell, m} \sigma_j^k \cdot C_t$.

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 3.4. 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 3.4, the inductive step follows. \square

Let property \mathcal{P}_{seq} contain any communication pattern arising from the concatenation of σ_i sequences defined at the start of the section and property \mathcal{P} contain all communication patterns generated by rooted graphs. We show the stronger statement that for every initial snapshot S_0 there is a trace $T = S_0, S_1, \dots$ starting at S_0 such that

$$\delta_{\mathcal{P}}(S_t) \geq \frac{1}{2^{\lceil \frac{t}{n-2} \rceil}} \delta_{\mathcal{P}}(S_0) \quad (2)$$

for all $t \geq 0$. It suffices to show (2) for \mathcal{P}_{seq} because $\delta_{\mathcal{P}}(S_t) \geq \delta_{\mathcal{P}_{\text{seq}}}(S_t)$ by Lemma 6.2 and $\delta_{\mathcal{P}_{\text{seq}}}(S_0) = \delta_{\mathcal{P}}(S_0)$ by Lemma 6.6 whenever $\mathcal{P} \supseteq \mathcal{P}_{\text{seq}}$. We hence suppose $\mathcal{P} = \mathcal{P}_{\text{seq}}$ in the rest of the proof. The proof is by inductive construction of a trace $T = S_0, S_1, \dots$ whose snapshots S_t satisfy (2). This, applied to an initial snapshot with $\delta_{\mathcal{P}}(S_0) > 0$, which exists by Lemma 6.6, 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 (2) holds for $t = (n-2)k$. First observe that, by construction of any $P \in \mathcal{P}$, there are three possible successor patterns until round $t+n-2$: $\sigma_1, \sigma_2, \sigma_3$. We thus have, $Y_{\mathcal{P}}^*(S_t) = Y_{\mathcal{P}}^*(S_{t+1}^1) \cup Y_{\mathcal{P}}^*(S_{t+1}^2) \cup Y_{\mathcal{P}}^*(S_{t+1}^3) = \dots = Y_{\mathcal{P}}^*(S_{t+n-2}^1) \cup Y_{\mathcal{P}}^*(S_{t+n-2}^2) \cup Y_{\mathcal{P}}^*(S_{t+n-2}^3)$, where $S_{t+n-2}^u = \sigma_u \cdot S_t$ for agent $u \in \{1, 2, 3\}$.

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

$$\text{diam}(Y_{\hat{u}}) \geq \text{diam}(Y)/2 . \quad (3)$$

We then define $S_{t+n-2} = S_{t+n-2}^{\hat{u}}$. By (3) and the induction hypothesis, we then have

$$\delta_{\mathcal{P}'}(S_{t+n-2}) = \dots = \delta_{\mathcal{P}'}(S_{t+1}) \geq \frac{\delta_{\mathcal{P}'}(S_t)}{2} \geq \frac{1}{2^{\lceil \frac{t}{n-2} \rceil + 1}} \delta_{\mathcal{P}'}(S_0) ,$$

i.e., Equation (2) 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, by Lemma 6.7, $(C_t, \pi \cdot \sigma_i) \sim_{\ell} (C_t, \pi \cdot \sigma_j)$ together with the fact that $\pi \cdot \sigma_i \cdot \sigma_{\ell}^{\omega} \in \mathcal{P}_{\text{seq}}$ and $\pi \cdot \sigma_j \cdot \sigma_{\ell}^{\omega} \in \mathcal{P}_{\text{seq}}$ we can apply Lemma 6.5 which shows that, for any pair $i, j \in \{1, 2, 3\}$ we have

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

By Lemma 6.4, 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 [12], Coulouma et al. characterized the network models in which exact consensus is solvable. In [8], Charron-Bost et al. showed that asymptotic consensus is solvable in a significantly broader class: it is solvable if and only if a network model is rooted. In this section

we aim to shed light on the deeper relation between these two problems by studying valencies and convergence rates. Our main results are a characterization of the topological structure of valencies with respect to solvability of exact consensus (Theorem 7.2) and nontrivial lower bounds on the contraction rates whenever exact consensus is not solvable (Theorem 7.4 and Corollary 7.5).

We start with recalling some definitions from Coulouma et al. [12]. 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 (Definition 4.7 in [12]). Let \mathcal{N} be a network model. Given $G, H, K \in \mathcal{N}$, we define $G \alpha_{\mathcal{N}, K} H$ if $\text{In}_{R(K)}(G) = \text{In}_{R(K)}(H)$. The relation $\alpha_{\mathcal{N}}^*$ is the transitive closure of the union of the relations $\alpha_{\mathcal{N}, K}$ where K varies in \mathcal{N} .

The following theorem is a characterization of network models in which exact consensus is solvable by the topological structure of valencies of asymptotic consensus algorithms. Of course, having to decide in a discrete set of value for exact consensus, valencies should be disconnected. However, the theorem shows that this condition is actually already sufficient to solve exact consensus.

THEOREM 7.2. *Let \mathcal{N} be a network model. Exact consensus is solvable in \mathcal{N} if and only if there exists an asymptotic consensus algorithm \mathcal{A} for \mathcal{N} such that $Y_{\mathcal{N}, \mathcal{A}}^*(C_0)$ is either a singleton or disconnected for all network models $\mathcal{N}' \subseteq \mathcal{N}$ and all initial configurations C_0 of \mathcal{A} .*

We next introduce the α -diameter of a network model \mathcal{N} , which we will then (see Theorem 7.4 and Corollary 7.5) show to be directly linked to a nontrivial lower bound on the contraction rate in \mathcal{N} if exact consensus is not solvable in \mathcal{N} . Note, that in case exact consensus is solvable in \mathcal{N} , the optimal contraction rate always is 0, obtained by a reduction argument to exact consensus.

Definition 7.3. Let \mathcal{N} be a network model. The α -diameter of \mathcal{N} is the smallest $D \geq 1$ such that for all $G, H \in \mathcal{N}$ there exist communication graphs $H_0, \dots, H_q \in \mathcal{N}$ and $K_1, \dots, K_q \in \mathcal{N}$ with $q \leq D$ such that $G = H_0$, $H = H_q$, and $H_{r-1} \alpha_{\mathcal{N}, K_r} H_r$ for all $r \in [q]$. In case it does not exist we set $D = \infty$.

Observe, that for the network model $\{H_0, H_1, H_2\}$ from Theorem 4.1, it is $D = 2$. Further, for network model $\text{deaf}(G)$, where G is an arbitrary communication graph, we have $D = 1$. The following theorem and corollary thus generalize Theorems 4.1 and 5.1 to arbitrary network models in which exact consensus is not solvable.

THEOREM 7.4. *Let \mathcal{N} be a network model in which exact consensus is not solvable. The contraction rate of any asymptotic consensus algorithm in \mathcal{N} is greater or equal to $1/(D+1)$ where D is the α -diameter of \mathcal{N} .*

Direct application of Theorem 7.4 to a network model \mathcal{N} in which exact consensus is not solvable may yield a trivial bound of 0 in case its α -diameter is ∞ . Indeed, we can, however, use Lemma 3.1 to derive a strictly positive bound for any \mathcal{N} in which exact consensus is not solvable:

COROLLARY 7.5. *Let \mathcal{N} be a network model in which exact consensus is not solvable. The contraction rate of any asymptotic consensus*

algorithm in \mathcal{N} is greater or equal to $1/(D+1)$ where D is the smallest α -diameter of $\mathcal{N}' \subseteq \mathcal{N}$ in which exact consensus is not solvable.

PROOF. Set $\mathcal{N}' \subseteq \mathcal{N}$ equal to the network model with the smallest α -diameter in which exact consensus is not solvable. Applying Theorem 7.4 to \mathcal{N}' , and Lemma 3.1 (iv) to \mathcal{N}' and \mathcal{N} yields the corollary. \square

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

In this section we show that Corollary 7.5 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 the classical asynchronous message passing systems. We consider a distributed system where agents perform receive-compute-broadcast steps. An agent may crash, i.e., stop making steps. Crashes can be unclear: 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 receive 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. Indeed algorithms that operate in rounds are widely used in asynchronous systems; see, e.g., [10, 14, 24].

We next show that Corollary 7.5 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 network model: Denote by \mathcal{G}_n the set of communication graphs with n nodes and let

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

for some $f < n/2$.

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

PROOF. Let $G, H \in \mathcal{N}_A$. Setting $q = \lceil n/f \rceil$, we choose the communication graphs H_r and K_r defined by

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

and

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

Clearly, it is $H_0 = G$ and $H_q = H$. Since we can write $R(K_r) = [n] \setminus \{i \mid (r-1)f + 1 \leq i \leq rf\}$ and $\text{In}_i(H_{r-1}) = \text{In}_i(H_r)$ for all $i \in K_r$, we also have $H_{r-1} \alpha_{\mathcal{N}_A, K_r} H_r$. Noting $H_r \in \mathcal{N}_A$ and $K_r \in \mathcal{N}_A$, this concludes the proof. \square

From Lemma 8.1 and Corollary 7.5 we immediately obtain the lower bound:

THEOREM 8.2. *The contraction rate 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 rate in Theorem 8.2 is with respect to rounds. However, we can easily 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 rate with respect to time.

8.2 General Algorithms

We next 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 rate of 0.

The following algorithm *MinRelay* is inspired by the exact consensus algorithm for synchronous systems with crash faults (see, e.g., [24]), 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 rate is 0.*

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., [24]). Formally, the local state of 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 . 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* in \mathcal{N} if for all $\varepsilon > 0$ and all Δ , each execution E with a communication pattern in \mathcal{N} 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 $y_E^1(0), \dots, y_E^n(0)$.

The above two problems are clearly closely related. However, the ε -Agreement condition does not preclude the decisions of a given agent, as a function of the error tolerance parameter ε , to

diverge, i.e., a priori may lead to unstable decisions with respect to this parameter.

We next extend our lower bounds on the contraction rate 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. [9]: For $n = 2$, running Algorithm 1 and deciding y^i after $\lceil \log_3 \frac{\Delta}{\varepsilon} \rceil$ rounds is optimal (Theorem 9.1). For $n \geq 3$ and the network model of all non-split graphs, running the midpoint algorithm and deciding after $\lceil \log_2 \frac{\Delta}{\varepsilon} \rceil$ rounds is optimal (Theorem 9.2). For $n \geq 4$ and the weakest network model 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 9.3).

We start with the case of two agents in Theorem 9.1. 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 9.1. *Let $\Delta > 0$ and $\varepsilon > 0$. In a network model of $n = 2$ agents that includes the three communication graphs H_0, H_1 , and H_2 , all approximate consensus algorithms have an execution with initial diameter $\Delta(y(0)) \leq \Delta$ and decision time greater or equal to $\log_3 \frac{\Delta}{\varepsilon}$.*

PROOF. Assume to the contrary that algorithm \mathcal{A} solves approximate consensus in some network model $\mathcal{N} \supseteq \{H_0, H_1, H_2\}$ that decides in $T < \log_3 \frac{\Delta}{\varepsilon}$ rounds for all vectors of initial values $y(0)$ with $\Delta(y(0)) \leq \Delta$ and some $\varepsilon > 0$.

Choose any $y(0)$ with $\Delta(y(0)) = \Delta$. 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 (1), there is an execution $E = C_0, G_1, C_1, G_2, \dots$ starting from C_0 such that

$$\delta_{\mathcal{N}}(C_T) \geq \frac{1}{3^T} \cdot \delta_{\mathcal{N}}(C_0) .$$

We have $\delta_{\mathcal{N}}(C_0) = \Delta(y(0)) = \Delta$ by Lemma 3.6 and $\delta_{\mathcal{N}}(C_T) \leq \Delta(y(T)) \leq \varepsilon$ 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 9.2. *Let $\Delta > 0$ and $\varepsilon > 0$. In a network model of $n \geq 3$ agents that includes the communication graphs $\text{deaf}(G)$, all approximate consensus algorithms have an execution with initial diameter $\Delta(y(0)) \leq \Delta$ and decision time greater or equal to $\log_2 \frac{\Delta}{\varepsilon}$.*

Analogously, for network models with rooted Ψ graphs, using (2), we obtain:

THEOREM 9.3. *Let $\Delta > 0$ and $\varepsilon > 0$. In a network model of $n \geq 4$ agents that includes the Ψ communication graphs, all approximate consensus algorithms have an execution with initial diameter $\Delta(y(0)) \leq \Delta$ and decision time greater or equal to $(n-2) \log_2 \frac{\Delta}{\varepsilon}$.*

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

THEOREM 9.4. *Let $\Delta > 0$ and $\varepsilon > 0$. In a network model in which exact consensus is not solvable, all approximate consensus algorithms have an execution with initial diameter $\Delta(y(0)) \leq \Delta$ and decision time greater or equal to $\log_{D+1} \frac{\Delta}{\varepsilon n}$, where D is the α -diameter of the network model.*

From Theorem 9.4 and the fact that $N' \subseteq N$ implies $\mathcal{E}' \subseteq \mathcal{E}$ for the corresponding sets of executions of algorithm \mathcal{A} , we get:

COROLLARY 9.5. *Let $\Delta > 0$ and $\varepsilon > 0$. In a network model in which exact consensus is not solvable, all approximate consensus algorithms have an execution with initial diameter $\Delta(y(0)) \leq \Delta$ and decision time greater or equal to $\log_{D+1} \frac{\Delta}{\varepsilon n}$, where D is the smallest α -diameter of a network model $N' \subseteq N$ in which exact consensus is not solvable.*

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 rates of asymptotic consensus algorithm in arbitrary network models: In particular, together with previously published averaging algorithms in [9], we showed tight bounds for the network model containing all non-split graphs, and the weakest network model in which asymptotic consensus is solvable, the network model of all rooted graphs. Furthermore we obtained a general lower bound of $1/(D + 1)$ for any network model in which exact consensus is not solvable; here D denotes the newly introduced α -diameter of the network model. 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 for the many fruitful discussions and her valuable input, which greatly helped improve the paper. The research was partially funded by the Austrian Science Fund (FWF) projects SIC (P26436) and ADynNet (P28182), and by the CNRS project PEPS DEMO.

REFERENCES

- [1] Ittai Abraham, Yonatan Amit, and Danny Dolev. 2004. Optimal resilience asynchronous approximate agreement. In *Proceedings of the 8th International Conference On Principles Of Distributed Systems (OPODIS 2004)*, Teruo Higashino (Ed.), Springer, Heidelberg, 229–239.
- [2] David Angeli and Pierre-Alexandre Bliman. 2006. Stability of leaderless discrete-time multi-agent systems. *Mathematics of Control, Signals, and Systems* 18, 4 (2006), 293–322.
- [3] Karl Johan Aström and Richard M. Murray. 2008. *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press, Princeton.
- [4] J. A. Benediktsson and P. H. Swain. 1992. Consensus theoretic classification methods. *IEEE Transactions on Systems, Man, and Cybernetics* 22, 4 (1992), 688–704.
- [5] Pierre-Alexandre Bliman, Angelia Nedic, and Asuman E. Ozdaglar. 2008. 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)*. IEEE, New York, 2226–2231.
- [6] Ming Cao, A. Stephen Morse, and Brian D. O. Anderson. 2008. Reaching a consensus in a dynamically changing environment: convergence rates, measurement delays, and asynchronous events. *SIAM Journal on Control and Optimization* 47, 2 (2008), 601–623.
- [7] Ming Cao, Daniel A. Spielman, and A. Stephen Morse. 2005. A lower bound on convergence of a distributed network consensus algorithm. In *Proceedings of the 44th IEEE Conference on Decision and Control, and the European Control Conference (CDC-ECC 2005)*, Hannes Frey, Xu Li, and Stefan Rührup (Eds.). IEEE, New York, 2356–2361.
- [8] Bernadette Charron-Bost, Matthias Függer, and Thomas Nowak. 2015. Approximate consensus in highly dynamic networks: the role of averaging algorithms. In *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming (ICALP 2015)*, Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann (Eds.). Springer, Heidelberg, 528–539.
- [9] Bernadette Charron-Bost, Matthias Függer, and Thomas Nowak. 2016. Fast, robust, quantizable approximate consensus. In *Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi (Eds.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, 137:1–137:14.
- [10] Bernadette Charron-Bost and André Schiper. 2009. The Heard-Of model: computing in distributed systems with benign faults. *Distributed Computing* 22, 1 (2009), 49–71.
- [11] Bernard Chazelle. 2011. The total s -energy of a multiagent system. *SIAM Journal on Control and Optimization* 49, 4 (2011), 1680–1706.
- [12] Étienne Coulouma, Emmanuel Godard, and Joseph Peters. 2015. A characterization of oblivious message adversaries for which consensus is solvable. *Theoretical Computer Science* 584 (2015), 80–90.
- [13] George Cybenko. 1989. Dynamic load balancing for distributed memory multiprocessors. *J. Parallel and Distrib. Comput.* 7, 2 (1989), 279–301.
- [14] Danny Dolev, Nancy A. Lynch, Shlomit S. Pinter, Eugene W. Stark, and William E. Weihl. 1986. Reaching approximate agreement in the presence of faults. *J. ACM* 33, 2 (1986), 499–516.
- [15] Magnus Egerstedt and Xiaoming Hu. 2001. Formation constrained multi-agent control. *IEEE Transactions on Robotics and Automation* 17, 6 (2001), 947–951.
- [16] Seyed Rasoul Etesami and Tamer Başar. 2016. Convergence time for unbiased quantized consensus. *IEEE Trans. Automat. Control* 61, 2 (2016), 443–455.
- [17] Alan D. Fekete. 1990. Asymptotically optimal algorithms for approximate agreement. *Distributed Computing* 4, 1 (1990), 9–29.
- [18] Alan D. Fekete. 1994. Asynchronous approximate agreement. *Information and Computation* 115, 1 (1994), 95–124.
- [19] Michael J Fischer, Nancy A Lynch, and Michael Merritt. 1986. Easy impossibility proofs for distributed consensus problems. *Distributed Computing* 1, 1 (1986), 26–39.
- [20] Rainer Hegselmann and Ulrich Krause. 2002. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation* 5, 3 (2002), 1–33.
- [21] Qun Li and Daniela Rus. 2006. Global clock synchronization in sensor networks. *IEEE Trans. Comput.* 55, 2 (2006), 214–226.
- [22] J. Lin, A.S. Morse, and B.D.O. Anderson. 2005. The multi-agent rendezvous problem. In *Cooperative Control: A Post-Workshop Volume, 2003 Block Island Workshop on Cooperative Control*, Vijay Kumar, Naomi Leonard, and A. Stephen Morse (Eds.). Springer, Heidelberg, 257–289.
- [23] Ronit Rubinfeld and Shlomo Moran. 1995. Closed schedulers: a novel technique for analyzing asynchronous protocols. *Distributed Computing* 8, 4 (1995), 203–210.
- [24] Nancy A. Lynch. 1996. *Distributed Algorithms*. Morgan Kaufmann, San Francisco.
- [25] Hammurabi Mendes, Maurice Herlihy, Nitin Vaidya, and Vijay K Garg. 2015. Multidimensional agreement in Byzantine systems. *Distributed Computing* 28, 6 (2015), 423–441.
- [26] Renato E. Mirolo and Steven H. Strogatz. 1990. Synchronization of pulse-coupled biological oscillators. *SIAM J. Appl. Math.* 50, 6 (1990), 1645–1662.
- [27] Luc Moreau. 2005. Stability of multiagent systems with time-dependent communication links. *IEEE Trans. Automat. Control* 50, 2 (2005), 169–182.
- [28] Angelia Nedic, Alexander Olshevsky, Asuman E. Ozdaglar, and John N. Tsitsiklis. 2009. On distributed averaging algorithms and quantization effects. *IEEE Trans. Automat. Control* 54, 11 (2009), 2506–2517.
- [29] Alex Olshevsky. 2015. Linear time average consensus on fixed graphs. *IFAC-PapersOnLine* 48, 22 (2015), 94–99.
- [30] Alex Olshevsky and John N. Tsitsiklis. 2011. Convergence speed in distributed consensus and averaging. *SIAM Rev.* 53, 4 (2011), 747–772.
- [31] Tamás Vicsek, András Czirók, Eshel Ben-Jacob, Inon Cohen, and Ofer Shochet. 1995. Novel type of phase transition in a system of self-driven particles. *Physical Review Letters* 75, 6 (1995), 1226–1229.
- [32] Y. Yuan, G.-B. Stan, L. Shi, M. Barahona, and J. Goncalves. 2013. Decentralized minimum-time consensus. *Automatica* 49, 5 (2013), 1227–1235.