



HAL
open science

Topological Characterization of Consensus under General Message Adversaries

Thomas Nowak, Ulrich Schmid, Kyrill Winkler

► **To cite this version:**

Thomas Nowak, Ulrich Schmid, Kyrill Winkler. Topological Characterization of Consensus under General Message Adversaries. ACM Symposium on Principles of Distributed Computing (PODC'19), Jul 2019, Toronto, Canada. 10.1145/3293611.3331624 . hal-02946833

HAL Id: hal-02946833

<https://hal.science/hal-02946833>

Submitted on 23 Sep 2020

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.

Topological Characterization of Consensus under General Message Adversaries

Dedicated to the 2018 Dijkstra Prize winners Bowen Alpern and Fred B. Schneider

Thomas Nowak*
Université Paris-Sud
Orsay, France
thomas.nowak@lri.fr

Ulrich Schmid
TU Wien
Vienna, Austria
s@ecs.tuwien.ac.at

Kyrill Winkler†
TU Wien
Vienna, Austria
kwinkler@ecs.tuwien.ac.at

ABSTRACT

In this paper, we provide a rigorous characterization of consensus solvability in synchronous directed dynamic networks controlled by an arbitrary message adversary using point-set topology: We extend the approach introduced by Alpern and Schneider in 1985 by introducing two novel topologies on the space of infinite executions: the process-view topology, induced by a distance function that relies on the local view of a given process in an execution, and the minimum topology, which is induced by a distance function that focuses on the local view of the process that is the last to distinguish two executions. We establish some simple but powerful topological results, which not only lead to a topological explanation of bivalence arguments, but also provide necessary and sufficient topological conditions on the admissible graph sequences of a message adversary for solving consensus. In particular, we characterize consensus solvability in terms of connectivity of the set of admissible graph sequences. For non-compact message adversaries, which are not limit-closed in the sense that there is a convergent sequence of graph sequences whose limit is not permitted, this requires the exclusion of all “fair” and “unfair” limit sequences that coincide with the forever bivalent runs constructed in bivalence proofs. For both compact and non-compact message adversaries, we also provide tailored characterizations of consensus solvability, i.e., tight conditions for impossibility and existence of algorithms, based on the broadcastability of the connected components of the set of admissible graph sequences.

CCS CONCEPTS

• **Theory of computation** → **Distributed algorithms**;

KEYWORDS

Topological characterization; point-set topology; consensus; dynamic networks; message adversaries

*Thomas Nowak has been supported by the Université Paris-Saclay project DEPEC MODE.

†Kyrill Winkler has been supported by the Austrian Science Fund (FWF) under project ADynNet (P28182) and RISE/SHiNE (S11405).

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 '19, July 29–August 2, 2019, Toronto, ON, Canada

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6217-7/19/07.

<https://doi.org/10.1145/3293611.3331624>

ACM Reference Format:

Thomas Nowak, Ulrich Schmid, and Kyrill Winkler. 2019. Topological Characterization of Consensus under General Message Adversaries: Dedicated to the 2018 Dijkstra Prize winners Bowen Alpern and Fred B. Schneider. In *2019 ACM Symposium on Principles of Distributed Computing (PODC '19)*, July 29–August 2, 2019, Toronto, ON, Canada. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3293611.3331624>

1 INTRODUCTION

We provide a complete characterization of consensus solvability in synchronous directed dynamic networks, controlled by a general message adversary, using a novel approach based on point-set topology as introduced by Alpern and Schneider [2]. Dynamic networks (see [14] for an overview) consist of a set of n fault-free processes that execute in communication-closed lock-step rounds. In every round, a directed communication graph (determined by the message adversary) defines which messages are delivered and which are lost. A *message adversary* (MA) [1] is thus just a set of infinite sequences of communication graphs.

Assuming that every process starts with some input value, the goal of a consensus algorithm is to irrevocably compute a common output value at every process eventually. A natural question to ask in this context is to describe the properties of the message adversaries that make deterministic consensus solvable. Until now, combinatorial methods were employed successfully to answer this question for certain classes of message adversaries [8, 9, 22, 24], but, to the best of our knowledge, there has been no success in finding a characterization for the general message adversary notion outlined above. Inspired by the 2018 Dijkstra Prize winning paper by Alpern and Schneider [2], our approach is thus based on point set topology.

Related work. In their celebrated paper [22], Santoro and Widmayer provided the first comprehensive characterization of consensus solvability in synchronous distributed systems prone to communication errors. Using bivalence arguments [10], they proved that consensus is impossible if up to $n - 1$ messages may be lost (by the same process) in each round. In [23], Schmid et al. showed that consensus can even be solved when a quadratic number of messages is lost per round, provided these losses do not isolate the processes. These approaches were generalized by the HO model by Charron-Bost and Schiper [7], which has been extended to also cover corrupted communication in [4].

All the above communication failure models (without corrupted messages) can also be expressed in terms of *oblivious* message adversaries [8], whose sequences are determined by all combinations




Figure 1: Comparison of the combinatorial topology approach and the point-set topology approach: The combinatorial topology approach (left) studies sequences of increasingly refined spaces in which the objects of interest are simplices (corresponding to configurations). The point-set topology approach (right) studies a single space in which the objects of interest are executions (i.e., infinite sequences of configurations).

of a set of possible communication graphs. For example, for $n = 2$ processes (the lossy link scenario), the result of [22] implies that if the message adversary may choose from the set of communication graphs $\{\leftarrow, \leftrightarrow, \rightarrow\}$ (where, e.g., \rightarrow represents successful communication from process 1 to process 2), consensus is impossible.

The above results were substantially refined by Coulouma, Godard and Peters in [8]. The authors identified a property of an equivalence relation defined on the set of communication graphs, which exactly captures consensus solvability in the oblivious setting. A universal consensus algorithm was also given, which, e.g., allows to solve consensus when the set of possible graphs for $n = 2$ consists of $\{\leftarrow, \rightarrow\}$.

The situation is considerably more complex for non-oblivious message adversaries, where the set of possible graphs may change over time. In sharp contrast to oblivious message adversaries, it need not be the case that the set of graph sequences is limit-closed, i.e., the model does not need to be compact, cf. [15]. For example, eventually stabilizing message adversaries like the *vertex-stable source component* (VSSC) message adversaries described in [6, 24] guarantee that some rounds with “good” communication graphs (a VSSC, which allows a consensus algorithm to terminate) occur eventually. However, limits of these sequences, in which the VSSC would never appear, are of course not in the message adversary.

For the special case of $n = 2$, Fevat and Godard [9] provided a complete characterization of consensus solvability for non-oblivious message adversaries as well: Using a bivalence argument, they showed that certain graph sequences (a “fair sequence” or a special pair of “unfair sequences”) must not be in the message adversary to render consensus solvable, and provided a universal algorithm for this case. However, a complete characterization of consensus solvability for arbitrary system sizes did not exist until now.

Regarding topological methods, one has to distinguish point-set topology, as introduced in [2], and combinatorial topology, as used for proving theorems characterizing solvable tasks in wait-free asynchronous shared memory systems with process crashes¹ [13] or under fair adversaries [15], for example. Combinatorial topology

studies the topology of the reachable states of admissible executions, captured by simplicial complexes as shown in Figure 1 (left), and has been developed into a widely applicable tool for the analysis of distributed systems [12].

By contrast, the primary objects of point-set topology are infinite executions, as shown in Figure 1 (right), where closed and dense sets precisely characterize safety and liveness properties, respectively. Apart from early work of one of the authors on characterizing consensus solvability in standard compact models [19], and a topological study of the strongly dependent decision problem [5], which both use the classic common prefix metric [2], the only distributed computing work we are aware of that utilizes point-set topology is [16]. In this paper, Lubitch and Moran introduced a construction for schedulers that lead to compact submodels of classic non-compact distributed computing models (like up to t crash failures). Whereas this greatly simplifies impossibility proofs, it does not lead to a precise characterization of solvability in non-compact models, however. Note that, in a similar spirit, [15] allows to reason, in the setting of combinatorial topology, about non-compact models by considering equivalent affine tasks that are compact.

Main contributions. We provide a complete² characterization of consensus solvability under compact and non-compact message adversaries, for an arbitrary number of processes. Rather than utilizing classic distributed computing techniques, we developed a novel modeling and analysis approach based on point-set topology. We had to add several new topological ideas to the setting of [2], as detailed below, which at the end provided us with a very powerful “toolbox” that, for example, allowed us to provide a topological explanation of bivalence [10] and bipotence [17] proofs. As our approach can be adapted to different distributed computing models, we believe that it could be of independent interest.

(i) We define two new topologies on the execution space, which allow to reason about sequences of local views of a certain process, rather than about configuration sequences. If this process is a fixed

¹Note that message adversaries do not fall into this category of models.

²Lacking space did not allow us to also include the proofs, which can be found in the full version [20] of our paper.

one (p) , the resulting $\{p\}$ -view topology is induced by a pseudo-metric $d_{\{p\}}(\alpha, \beta)$ based on the common prefix of p 's local views in the executions α, β . Alternatively, this process can be the last one to notice a difference between executions, which gives rise to the *minimum topology* induced by the pseudo-semi-metric $d_{\min}(\alpha, \beta) = \min_{p \in [n]} d_{\{p\}}(\alpha, \beta)$.

(ii) Since the set of possible views is not necessarily finite, the product topologies built on the execution space are not compact *a priori*. Fortunately, the space of sequences of process-time graphs [3], which are finite, comes to our rescue: Since the local transition function τ , which maps sequences of admissible process-time graphs to the corresponding admissible executions Γ , is continuous, the resulting product subspaces *are* compact.

(iii) We show that consensus can be modeled as a continuous decision function Δ , which maps an admissible execution to its unique decision value. In conjunction with the above results, this allows us to prove that consensus is solvable if and only if all the decision sets, i.e., the pre-images $\Gamma_v = \Delta^{-1}[\{v\}]$ resp. $PS_v = \tau^{-1}[\Delta^{-1}[\{v\}]]$ for every decision value v , are separated in our topologies. We also provide a universal consensus algorithm, which relies on this separation. Moreover, we show that separability is equivalent to broadcastability of the connected components (in the sense that there is a process that is heard by all processes in every sequence in a connected component).

(iv) Using some properties of the pseudo-metric $d_{\{p\}}$, we provide a topological definition of fair and unfair sequences [9]. They turn out to be the limits of two infinite sequences of executions lying in two different decision sets, which have distance 0, and happen to coincide with the forever bivalent/bipotent executions constructed in bivalence proofs. We show that a message adversary must not include such fair and unfair sequences for consensus to be solvable.

(v) We use our generic results to give a complete characterization of consensus solvability for both compact and non-compact message adversaries. For the former (like [8, 22, 23]), we introduce a simple ε -approximation of a decision set, which is equivalent to the set of the $(\log \frac{1}{\varepsilon})$ -prefixes of the sequences contained therein. We prove that consensus can be solved, using our universal algorithm, if and only if the ε -approximations are broadcastable for some ε . For non-compact message adversaries (like [6, 9, 24]), the ε -approximation does not work, so one needs to apply our universal algorithm based on the connected components directly.

Paper organization. In Section 2 and 3, we define the elements of the spaces that are endowed with our new topologies in Section 4. Section 5 introduces the consensus problem in topological terms and provides our abstract characterization result (Theorem 5.5, which also provides a universal algorithm). We refine our characterization according to the properties of the minimum topology (Theorem 5.11) and the process-view topology (Corollary 5.18). Section 6 is devoted to the application of our generic results to compact (Theorem 6.6) and non-compact (Theorem 6.7) message adversaries. Some conclusions in Section 7 round off our paper.

2 SYSTEM MODEL

We consider deterministic algorithms for a synchronous directed dynamic network consisting of n processes with unique identifiers taken from $[n] = \{1, \dots, n\}$.

A *communication graph* is any directed graph $G = ([n], E)$, $E \subseteq [n] \times [n]$ with node set $[n]$. A *message adversary* (MA) is a set of infinite sequences of communication graphs. Graph sequences in this set are called *admissible* under the message adversary.

An algorithm \mathcal{A} , an initial configuration C^0 of \mathcal{A} , and a graph sequence $G = (G^t)_{t \geq 1}$ uniquely determine an *execution* in a round-by-round fashion: During round t , every process p updates its previous local state C_p^{t-1} to its new state C_p^t in a deterministic way according to \mathcal{A} , based on the messages it received during round t . Rounds advance synchronously in a send–receive–compute order. Messages in round t are delivered according to communication graph G^t : Process q receives process p 's message in round t if and only if (p, q) is an edge of G^t .

An algorithm is defined by the local transition function and a set of initial states for every process. Since we are only interested in consensus algorithms, we stipulate that processes have an initial state for every initial value v of the finite input domain \mathcal{V}_I . Note carefully that we assume that a process running \mathcal{A} does not know n a priori, and cannot always infer n from the messages it receives in an execution either. By contrast, the message adversary need not be oblivious w.r.t. the algorithm, i.e., it may know \mathcal{A} and choose its graph sequences accordingly.

A *configuration* is a tuple (C_1, \dots, C_n) of process states. Whenever \mathcal{A} is clear from the context, we denote its set of possible configurations by C . C^ω denotes the set of all infinite sequences of configurations, $\Gamma \subseteq C^\omega$ is the set of admissible executions resulting from admissible graph sequences. Executions are represented by Greek letters α, β, \dots

3 PROCESS-TIME GRAPHS

We will consider *process-time graphs* [3], which are useful for reasoning about the *causal past* of a process in a given round, and which will play a crucial role in our topological framework.

For every graph sequence $G = (G^t)_{t \geq 1}$ and every assignment of initial values $x \in \mathcal{V}_I^n$ to the n processes, we inductively construct the following sequence of process-time graphs PT^t :

- The process-time graph PT^0 at time 0 contains the nodes $(p, 0, x_p)$ for all processes $p \in [n]$, with input value $x_p \in \mathcal{V}_I$, and no edges.
- The process-time graph PT^1 at time 1 contains the nodes $(p, 0, x_p)$ and $(p, 1)$ for all processes $p \in [n]$. It contains an edge from $(p, 0, x_p)$ to $(q, 1)$ if and only if $(p, q) \in G^1$.
- For $t \geq 2$, the process-time graph PT^t at time t contains the nodes of PT^{t-1} and the nodes (p, t) for all processes $p \in [n]$. It contains an edge from $(p, t-1)$ to (q, t) if and only if $(p, q) \in G^t$.

Figure 2 contains an example of a process-time graph at time 2.

Let \mathcal{PT}^t be the set of all possible process-time graphs at time $t \geq 0$, which is finite for any $t \geq 0$. Furthermore, let $\mathcal{PT}^\omega = \mathcal{PT}^0 \times \mathcal{PT}^1 \times \dots$ be the set of all infinite sequences of possible process-time graphs.³

Given an input domain \mathcal{V}_I , every message adversary uniquely corresponds to a subset PS of \mathcal{PT}^ω generated by the graph sequences admissible under the message adversary.

³Please note that we slightly abuse the notation \mathcal{PT}^ω here, which normally represents $\mathcal{PT} \times \mathcal{PT} \times \dots$




Figure 2: Example of a process-time graph PT^2 at time $t = 2$ with $n = 3$ processes and initial values $x = (1, 0, 1)$. Process 1's view $V_{\{1\}}(PT^2)$ is highlighted in bold green.

Note that process-time graphs are *independent* of the particular algorithm used, but do depend on the input values. For conciseness, we will use the term process-time graphs for the elements of \mathcal{PT}^ω as well. We will denote elements of \mathcal{PT}^ω by Roman letters a, b, \dots (or sometimes PT, PT', \dots).

4 TOPOLOGICAL STRUCTURE

In this section, we will endow the various sets introduced in Section 2 and 3 with suitable topologies. We first recall briefly the basic topological notions that are needed for our exposition. For a more thorough introduction, however, the reader is advised to refer to a textbook [18].

A topology for a set X is a family \mathcal{T} of subsets of X such that $\emptyset \in \mathcal{T}$, $X \in \mathcal{T}$, and \mathcal{T} contains all arbitrary unions as well as all finite intersections of its members. We call X endowed with \mathcal{T} , often written as (X, \mathcal{T}) , a (topological) space and the members of \mathcal{T} open sets. The complement of an open set is called closed and sets that are both open and closed, such as \emptyset and X itself, are called clopen. A metric on X is a function $d : X \times X \rightarrow \mathbb{R}_+$ such that for all $x, y, z \in X$ we have $d(x, y) = 0$ if and only if $x = y$, $d(x, y) = d(y, x)$, and $d(x, z) \leq d(x, y) + d(y, z)$. The topology induced by a metric d is the collection of sets U such that for each $u \in U$, there is an ε -ball $B_\varepsilon(u) = \{v \mid d(u, v) < \varepsilon\}$ with $\varepsilon > 0$ and $u \in B_\varepsilon(u) \subseteq U$.

A function from space X to space Y is continuous if the pre-image of every open set in Y is open in X . Given a space (X, \mathcal{T}) , $Y \subseteq X$ is called a subspace of X if Y is equipped with the subspace topology $\{Y \cap U \mid U \in \mathcal{T}\}$. Given $A \subseteq X$, the closure of A is the intersection of all closed sets containing A . For a space X , if $A \subseteq X$, we call x a limit point of A if it belongs to the closure of $A \setminus \{x\}$. It can be shown that the closure of A is the union of A with all limit points of A . A space X is called compact if every union of open sets that covers X contains a finite union of open sets that covers X . Note that our topologies are also sequentially compact, which ensures that every infinite sequence of elements has a convergent subsequence.

In previous work on point-set topology [19], the set of configurations C of some fixed algorithm \mathcal{A} was endowed with the discrete topology, induced by the discrete metric $d_{\max}(C, D) = 1$ if $C \neq D$

and 0 otherwise (for configurations $C, D \in C$)⁴ Moreover, C^ω was endowed with the corresponding product topology⁵, which happens to be induced by the *common prefix metric*

$$d_{\max}(\alpha, \beta) = 2^{-\inf\{t \geq 0 \mid \alpha^t \neq \beta^t\}}. \quad (1)$$

Our generalization will focus on the local views of the processes, which are obtained by suitable projection functions: For a tuple $x = (x_1, \dots, x_n)$ and any $\emptyset \neq P = \{p_1, \dots, p_k\} \subseteq [n]$, the P -projection function is defined as $\pi_P(x) = (x_{p_1}, \dots, x_{p_k})$ where $p_1 \leq \dots \leq p_k$. Similarly, for any infinite sequence $\theta = (\theta^0, \theta^1, \dots)$ and any $t \geq 0$, the t -projection function is defined as $\pi^t(\theta) = \theta^t$.

The views of a fixed non-empty subset of the processes in a configuration can be defined as follows:

Definition 4.1 (Views). For any configuration $C \in C$, the view of the processes in $\emptyset \neq P \subseteq [n]$ is $V_P(C) = \pi_P(C)$. The set of all possible P -views is defined as $\text{Views}_P = \{V_P(C) \mid C \in C\}$.

4.1 Process-View Topologies

We will now introduce a topology on the set C of configurations that relies on the corresponding set of views Views_P of a set of processes P .

Definition 4.2 (P-view topology for configurations). Let $P \subseteq [n]$ be a nonempty set of processes. The P -view topology \mathcal{T}_P^C on C is defined as the topology induced by the subbasis $\{V_P^{-1}[V] \mid V \subseteq \text{Views}_P\}$, i.e., is the arbitrary union of finite intersections of the elements $V_P^{-1}[V]$ for an arbitrary set $V \subseteq \text{Views}_P$.

The topology \mathcal{T}_P^C is induced by the pseudo-metric⁶

$$d_P(C, D) = \begin{cases} 0 & \text{if } V_P(C) = V_P(D) \\ 1 & \text{else} \end{cases}.$$

Note that the topology \mathcal{T}_P^C is much coarser than the discrete topology on C , as it does not distinguish the local states of processes outside of P .

The corresponding product topology on the set C^ω of sequences of configurations, also denoted P -view topology, can be defined as follows: The P -view topology $\mathcal{T}_P^{C^\omega}$ on C^ω is defined as the topology induced by the subbasis $\{(\pi^t)^{-1}[U] \mid t \geq 0, U \in \mathcal{T}_P^C\}$, i.e., is an arbitrary union of finite intersections of the elements $(\pi^t)^{-1}[U]$, where U is an arbitrary open set in \mathcal{T}_P^C .

We will show that the P -view topology is induced by the following pseudo-metric on C^ω :

$$d_P(\alpha, \beta) = 2^{-\inf\{t \geq 0 \mid V_P(\alpha^t) \neq V_P(\beta^t)\}}.$$

We call this the P -pseudo-metric on C^ω . Figure 3 shows an example for different sets P .

It follows immediately from the definition that, if $\alpha, \beta \in C^\omega$ satisfy $d_P(\alpha, \beta) < 2^{-t}$, then the processes in P have the same view of the first t configurations in α and β . Moreover, the P -pseudo-metric satisfies the following properties:

⁴The notation d_{\max} stems from the fact that it is equal to the maximum of the P -pseudo-metrics defined in the next subsection.

⁵The product topology on a product $\prod_{i \in I} X_i$ of topological spaces is defined as the coarsest topology such that all projections $\pi_i : \prod_{i \in I} X_i \rightarrow X_i$ are continuous.

⁶A pseudo-metric has the same properties as a metric, except that it lacks definiteness, i.e., we can have $d(x, y) = 0$ for $x \neq y$.




Figure 3: Comparison of the P -view, minimum, and common prefix topologies. The first three configurations of each of the two executions α and β with three processes and two different possible local states (dark blue and light yellow) are depicted. We have $d_{\max}(\alpha, \beta) = d_{\{3\}}(\alpha, \beta) = 1$, $d_{\{2\}}(\alpha, \beta) = 1/2$, and $d_{\min}(\alpha, \beta) = d_{\{1\}}(\alpha, \beta) = 1/4$.

THEOREM 4.3 PROPERTIES OF P -PSEUDO-METRIC. *The P -pseudo-metric $d_P(\alpha, \beta)$ on C^ω satisfies*

$$\begin{aligned} d_P(\alpha, \beta) &= d_P(\beta, \alpha) && \text{(symmetry),} \\ d_P(\alpha, \gamma) &\leq d_P(\alpha, \beta) + d_P(\beta, \gamma) && \text{(triangle inequality),} \\ d_P(\alpha, \beta) &\leq d_Q(\alpha, \beta) && \text{(monotonicity for } P \subseteq Q\text{),} \\ d_{[n]}(\alpha, \beta) &= d_{\max}(\alpha, \beta) && \text{(common prefix metric).} \end{aligned}$$

Despite of the lack of definiteness, most properties of metric spaces, including compactness, hold also in pseudo-metric spaces [11]. What is obviously lost is the uniqueness of the limit of a convergent sequence of executions, however: if $\alpha_k \rightarrow \hat{\alpha}$ and $d_P(\hat{\alpha}, \hat{\beta}) = 0$, then $\alpha_k \rightarrow \hat{\beta}$ as well.

THEOREM 4.4 P -VIEW TOPOLOGY INDUCED BY P -PSEUDO-METRIC. *The P -view topology $\mathcal{T}_P^{C^\omega}$ on C^ω is induced by the P -pseudo-metric.*

We could use exactly the same machinery as above for endowing the set \mathcal{G} of communication graphs and the set \mathcal{G}^ω of sequences of graph sequences, like C and C^ω , with an analogous P -view topology. It turns out, however, that this does not lead to a transition function that is continuous w.r.t. the P -pseudo-metric for all algorithms.

Fortunately, however, the P -view projection function $V_P(\cdot)$ is also meaningful for any process-time graph PT^t : It just represents the causal past of the processes in P at the end of round t , i.e., the sub-graph induced by all process-time nodes (q, t') that have a path to node (p, t) for some $p \in P$ in the process-time graph PT^t .

Moreover, we can define a P -view topology on \mathcal{PT}^ω , analogous to that on C^ω . The same is true for the proof of the correspondence Theorem 4.4, which hence also holds here.

In sharp contrast to the set of configurations C , however, the set of process-time graphs \mathcal{PT}^t is finite for any time t since $\mathcal{V}_I^n \times \mathcal{G}^t$ is finite. Tychonoff's theorem⁷ hence implies compactness of the P -view topology on \mathcal{PT}^ω , which is not necessarily the case for C^ω .

⁷Tychonoff's theorem states that any product of compact spaces is compact (with respect to the product topology).

Since the algorithms take decisions based on local states only, we can define the *transition function* $\tau : \mathcal{PT}^\omega \rightarrow C^\omega$ that provides $C^t = \tau(PT^t)$ for every $t \geq 0$, i.e., maps process-time graphs to the corresponding configurations. The following Lemma 4.5 shows that τ is continuous w.r.t. the P -pseudo-metric.

LEMMA 4.5 CONTINUITY OF τ W.R.T. d_P . *Let $P \subseteq [n]$. The transition function $\tau : \mathcal{PT}^\omega \rightarrow C^\omega$ is continuous when both \mathcal{PT}^ω and C^ω are endowed with d_P .*

Since the image of a compact space under a continuous function is compact, it hence follows that the set $\tau[\mathcal{PT}^\omega] \subseteq C^\omega$ of admissible executions is a compact subspace of $(C^\omega, \mathcal{T}_P^{C^\omega})$.

The common structure of \mathcal{PT}^ω and its image under the transition function τ , implied by the continuity of τ , allows us to reason in either of these spaces. We will usually reason in \mathcal{PT}^ω or in its subspace PS . Note also that any sequence $a \in \mathcal{PT}^\omega$ or $\gamma \in \tau[\mathcal{PT}^\omega]$ can also be identified by specifying a vector $x \in \mathcal{V}_I^n$ of input values and a graph sequence $G \in \mathcal{G}^\omega$ that leads to a and hence γ .

4.2 Minimum Topology

For our characterization of consensus solvability, we endow the set C^ω of configuration sequences and the set of process-time graphs \mathcal{PT}^ω with a slightly different but related topology, induced by the distance function

$$d_{\min}(\alpha, \beta) = \min_{p \in [n]} d_{\{p\}}(\alpha, \beta), \quad (2)$$

which reflects the common prefix length of the process p that is the *last* one to distinguish its q -view in α and β .

Note carefully, however, that d_{\min} only satisfies symmetry and nonnegativity but not the triangle inequality, i.e., is only a *pseudo-semi-metric*: There may be sequences with $d_{\{p\}}(\alpha, \beta) = 0$ and $d_{\{q\}}(\beta, \gamma) = 0$ but $d_{\{q\}}(\alpha, \gamma) > 0$ for all $q \in [n]$. Hence, the topology $\mathcal{T}_{\min}^{C^\omega}$ on C^ω induced by d_{\min} lacks many of the properties of (pseudo-)metric spaces, but will turn out to be already sufficient for the characterization of the possibility/impossibility of consensus (see Theorem 5.4). Alternatively, however, one can also consider all individual pseudo-metric spaces induced by $d_{\{p\}}$, $p \in [n]$, separately and compute the minimum afterwards.

Pseudo-semi-metrics induce a topology just like (pseudo-)metrics. In fact, a much more general result holds:

LEMMA 4.6 PSEUDO-SEMI-METRICS INDUCE TOPOLOGIES. *Let X be a nonempty set and $d : X \times X \rightarrow \mathbb{R}$ be a function. Define $\mathcal{T} \subseteq 2^X$ by setting $U \in \mathcal{T}$ if and only if for all $x \in U$ there exists some $\varepsilon > 0$ such that*

$$B_\varepsilon(x) = \{y \in X \mid d(x, y) < \varepsilon\} \subseteq U.$$

Then \mathcal{T} is a topology on X .

Denote by \mathcal{T}_{\min}^C the topology on the set C of configurations induced by

$$d_{\min}(C, D) = \min_{p \in [n]} d_{\{p\}}(C, D).$$

Then the function d_{\min} defined on C^ω induces the product topology $\mathcal{T}_{\min}^{C^\omega}$, where every copy of C is endowed with \mathcal{T}_{\min}^C :

LEMMA 4.7 PSEUDO-SEMI-METRIC FOR PRODUCT TOPOLOGIES. *Let X be a nonempty set and let $d : X \times X \rightarrow \{0, 1\}$ be a function. Then*

the product topology of X^ω , where every copy of X is endowed by the topology \mathcal{T}^ω induced by d , is induced by

$$d^\omega : X^\omega \times X^\omega \rightarrow \mathbb{R} \quad , \quad d^\omega(\alpha, \beta) = 2^{-\inf\{t \geq 0 \mid d(\alpha^t, \beta^t) > 0\}} \quad .$$

Specializing to $X = C$ and $d = d_{\min}$, we get:

LEMMA 4.8 MINIMUM TOPOLOGY IN TERMS OF PROCESS-VIEW TOPOLOGIES. *We have*

$$d^\omega(\alpha, \beta) = \min_{p \in [n]} d_{\{p\}}(\alpha, \beta) = d_{\min}(\alpha, \beta)$$

for all $\alpha, \beta \in C^\omega$.

Like for $\mathcal{T}_P^{\mathcal{P}\mathcal{T}^\omega}$, we get from Tychonoff's theorem and the finiteness, and hence compactness, of every $\mathcal{P}\mathcal{T}^t$ that $(\mathcal{P}\mathcal{T}^\omega, \mathcal{T}_{\min}^{\mathcal{P}\mathcal{T}^\omega})$ is compact.

We finally show that the function $\tau : \mathcal{P}\mathcal{T}^\omega \rightarrow C^\omega$ is continuous also for the distance d_{\min} .

LEMMA 4.9 CONTINUITY OF τ W.R.T. d_{\min} . *The local transition function $\tau : \mathcal{P}\mathcal{T}^\omega \rightarrow C^\omega$ is continuous when $\mathcal{P}\mathcal{T}^\omega$ and C^ω are endowed with d_{\min} .*

Like for the P -view topologies, the continuity of τ also implies the compactness of $\tau[\mathcal{P}\mathcal{T}^\omega]$ in the minimum-topology.

5 CONSENSUS

In this section, we will develop a topological condition for consensus under message adversaries. Unlike in [19], where the admissible graph sequences needed to be compact in the space induced by the common prefix metric, this is usually not the case here: The set of admissible graph sequences need not be limit-closed in general, see Section 6.2.

We will consider the consensus problem, which is defined as follows:

Definition 5.1 (Consensus). Every process $p \in [n]$ has an input value $x_p \in \mathcal{V}_I$, taken from a finite input domain \mathcal{V}_I , which is set in the initial state, and an output value $y_p \in \mathcal{V}_O \cup \{\perp\}$, with a finite output domain $\mathcal{V}_O \supseteq \mathcal{V}_I$, initially $y_p = \perp$. In every admissible execution, a correct consensus algorithms \mathcal{A} must ensure the following properties:

- (T) Eventually, every p must decide, i.e., change to $y_p \neq \perp$, exactly once (termination).
- (A) If p and q have decided, then $y_p = y_q$ (agreement).
- (V) If $x_p = v$ for all $p \in [n]$, then v is the only possible decision value (validity).

Note that our framework can be easily adapted to different validity conditions, like strong validity, where every decision value must satisfy $y_p = x_q$ for some $q \in [n]$.

If we endow the set \mathcal{V}_O with the discrete topology, it turns out that consensus can be described by a continuous map from the set of admissible executions, endowed with any P -view topology or the minimum topology, to \mathcal{V}_O .

LEMMA 5.2 CONTINUITY OF CONSENSUS. *Let $\Gamma \subseteq C^\omega$ be the set of admissible executions of some consensus algorithm \mathcal{A} . Define the map $\Delta : \Gamma \rightarrow \mathcal{V}_I$ such that $\Delta(y)$ is the common decision value of algorithm \mathcal{A} in y . Then, Δ is continuous, both with respect to any P -view topology and the minimum topology.*

In the following definition, we introduce the sets of process-time graphs and executions that lead to a given decision value:

Definition 5.3 (Decision sets). For every output value $v \in \mathcal{V}_O$, let $PS(v) = \tau^{-1}[\Delta^{-1}[\{v\}]] \subseteq PS$ and $\Gamma(v) = \Delta^{-1}[\{v\}]$ be the set of admissible process-time graph sequences and admissible executions that lead to a common decision value v , respectively.

We need a few more basic topological terms: A set in a topological space is *clopen*, if it is both closed and open. A topological space is *disconnected*, if it contains a nontrivial clopen set, which means that it can be partitioned into two disjoint open sets. It is *connected*, if it is not disconnected.

With these preparations, we can already provide a topological consensus impossibility result:

THEOREM 5.4 CONSENSUS IMPOSSIBILITY. *If an algorithm solves consensus, then all of its decision sets $\Gamma(v) = \Delta^{-1}[\{v\}]$, $v \in \mathcal{V}_O$, and $PS(v) = \tau^{-1}[\Delta^{-1}[\{v\}]]$ are clopen in the subspace topology of Γ and PS , respectively, both w.r.t. any d_P and d_{\min} .*

In particular, consensus is impossible if the set Γ of admissible executions or the set PS of admissible process-time graphs is connected.

5.1 Characterization in the Minimum Topology

We call a process-time graph z_v , for $v \in \mathcal{V}_O$, v -valent, if it starts from an initial configuration where all processes $p \in [n]$ have the same input value $x_p(z_v) = v$. Let PS_{z_v} denote the connected component of PS that contains the v -valent $z_v \in PS$. In the minimum topology $\mathcal{T}_{\min}^{\mathcal{P}\mathcal{T}^\omega}$, we get the following characterization of consensus solvability:

THEOREM 5.5 CONSENSUS CHARACTERIZATION. *Consensus is solvable with a message adversary generating the set of admissible process-time graph sequences PS if and only if there exists a partition of PS into sets $PS(v)$, $v \in \mathcal{V}_O$ such that the following holds:*

- (1) Every $PS(v)$ is open in PS with respect to the minimum topology.
- (2) Every admissible v -valent $z_v \in PS$ satisfies $z_v \in PS(v)$.

PROOF. (\Rightarrow): Define $PS(v) = (\Delta \circ \tau)^{-1}[\{v\}] \subseteq PS$ using the functions τ and Δ defined by an algorithm that solves consensus. This is clearly a partition of PS by the termination and validity property of consensus. The validity condition of the algorithm also implies property (2). It thus remains to show openness of the $PS(v)$, which follows from the continuity of $\Delta \circ \tau : PS \rightarrow \mathcal{V}_O$, as every singleton $\{v\}$ is open (and closed) in the discrete topology.

(\Leftarrow): We construct an algorithm as follows. Each process's state contains a variable V whose value is equal to the projection of process-time graphs onto its own view. State updates happen in such a way that process p 's variable at the end of round t in the execution with process-time graph a is equal to $V = \pi_{\{p\}}(a^t)$. Process p decides value v in round t if the ball of radius $\varepsilon = 2^{-t}$ around the set of sequences of process-time graphs $\pi_{\{p\}}^{-1}[\{V\}]$ compatible with its locally recorded view V is contained in $PS(v)$, i.e., if

$$\{b \in PS \mid \pi_{\{p\}}(b^t) = V\} \subseteq PS(v) \quad .$$

We first show termination of the resulting algorithm. Let $a \in PS$ and let $v \in \mathcal{V}_O$ such that $a \in PS(v)$. Since $PS(v)$ is open w.r.t. d_{\min} ,

there exists some $\varepsilon > 0$ such that

$$\{b \in PS \mid d_{\min}(b, a) < \varepsilon\} \subseteq PS(v) .$$

Let $t \geq 0$ such that $2^{-t} \leq \varepsilon$. We hence have

$$\{b \in PS \mid d_{\min}(b, a) < 2^{-t}\} \subseteq \{b \in PS \mid d_{\min}(b, a) < \varepsilon\} \subseteq PS(v) .$$

By Lemma 4.8, we have

$$\begin{aligned} \{b \in PS \mid d_{\{p\}}(b, a) < 2^{-t}\} &\subseteq \{b \in PS \mid d_{\min}(b, a) < 2^{-t}\} \\ &\subseteq PS(v) \end{aligned}$$

for every process $p \in [n]$. By the definition of the distance $d_{\{p\}}$, we have

$$\{b \in PS \mid d_{\{p\}}(b, a) < 2^{-t}\} = \{b \in PS \mid \pi_{\{p\}}(b^t) = \pi_{\{p\}}(a^t)\} .$$

Since, by construction, process p 's variable V at the end of round t is equal to $\pi_{\{p\}}(a^t)$, this shows that process p decides in or before round t .

To show agreement, assume by contradiction that there exists some $a \in PS$ such that process p_1 decides v_1 in round t_1 and process p_2 decides v_2 in round t_2 in execution $\tau(a)$ with $q_1 \neq q_2$. By the definition of the algorithm, we have

$$\{b \in PS \mid \pi_{\{p_1\}}(b^{t_1}) = \pi_{\{p_1\}}(a^{t_1})\} \subseteq PS(v_1)$$

and

$$\{b \in PS \mid \pi_{\{p_2\}}(b^{t_2}) = \pi_{\{p_2\}}(a^{t_2})\} \subseteq PS(v_2) .$$

In particular, $a \in PS(v_1)$ and $a \in PS(v_2)$, a contradiction to the fact that the $PS(v)$ form a partition of PS .

Validity is an immediate consequence of property (2). \square

This characterization gives rise to the following meta-procedure for determining whether consensus is solvable and constructing an algorithm if it is, which will be instantiated for some examples in Section 6. It requires knowledge of the connected components of the space PS with respect to the minimum topology:

- (1) Initially, start with an empty set $PS(v)$ for every value $v \in \mathcal{V}_O$.
- (2) Add to $PS(v)$ every connected component PS_{z_v} of every v -valent $z_v \in PS$.
- (3) Add every remaining connected component of PS to an arbitrarily chosen set $PS(v)$ (i.e., decide on default value v).
- (4) If the sets $PS(v)$ are pairwise disjoint, then consensus is solvable. In this case, the sets $PS(v)$ determine a consensus algorithm via the universal construction in the proof of Theorem 5.5. If the $PS(v)$ are not pairwise disjoint, then consensus is not solvable.

In particular, this meta-procedure gives rise to the following succinct characterization of consensus solvability.

COROLLARY 5.6. *Consensus is solvable with a message adversary MA generating the set of admissible process-time graph sequences PS if and only if none of its connected components with respect to the minimum topology contains z_v and z_w , $v, w \neq v \in \mathcal{V}_O$ that are v -valent and w -valent, respectively.*

We will now develop another characterization of consensus solvability, with rests on broadcastability of the PS_{z_v} .

Definition 5.7 (Diameter of a set). For $A \subseteq \mathcal{PT}^\omega$ and P' denoting either $P \subseteq [n]$ or \min , define A 's diameter as $d_{P'}(A) = \sup\{d_{P'}(a, b) \mid a, b \in A\}$.

Definition 5.8 (Broadcastability). We call a subset $A \subseteq PS$ of admissible process-time graphs *broadcastable* by the broadcaster $p \in [n]$, if for every $a \in A$ there is some round $T(a) < \infty$ where every process $q \in [n]$ knows p 's input value $x_p(a)$ in a , i.e., $(p, 0, x_p(a))$ is in $V_{\{q\}}(a^{T(a)})$.

We will now prove the essential fact that *connected* broadcastable sets have a diameter strictly smaller than 1:

THEOREM 5.9 DIAMETER OF BROADCASTABLE CONNECTED SETS. *If a connected set $A \subseteq PS$ of admissible process-time graph sequences is broadcastable by some process p , then $d_{\min}(A) \leq d_{\{p\}}(A) \leq 1/2$, i.e., p 's input value $x_p(a)$ is the same for all $a \in A$.*

Corollary 5.10 follows immediately from Theorem 5.9:

COROLLARY 5.10 DIAMETER OF BROADCASTABLE PS_{z_v} . *If PS_{z_v} for a v -valent $z_v \in PS$ is broadcastable for p , then $d_{\min}(PS_{z_v}) \leq d_{\{p\}}(PS_{z_v}) \leq 1/2$ since p 's input value $x_p(a) = v$ is the same for all $a \in PS_{z_v}$.*

We can now prove the following necessary and sufficient condition for solving consensus based on broadcastability:

THEOREM 5.11 CONSENSUS CHARACTERIZATION VIA BROADCASTABILITY. *A message adversary allows to solve consensus if and only if it guarantees that the connected components of the set PS of admissible processes-time graphs are broadcastable for some process.*

5.2 Characterization in the P -View Topologies

It is possible to shed some additional light on the consensus characterization given in Theorem 5.5, by exploiting the fact that d_P (unlike d_{\min}) is a pseudo-metric: Since most of the convenient properties of metric spaces, including sequential compactness, also hold in pseudo-metric spaces, we can further explore the border of the decision sets $PS(v)$. It will turn out in Corollary 5.18 that consensus is impossible if and only if certain limit points in the P -view topology $\mathcal{T}_P^{\mathcal{PT}^\omega}$ are admissible.

For a given consensus algorithm, we again consider the set of all admissible process-time graph sequences PS resp. the corresponding set of admissible executions Γ . We endow PS with the subspace topology generated by $\mathcal{PT}^\omega \cap PS$ resp. Γ with the subspace topology⁸ generated by $C^\omega \cap \Gamma$, both in the P -view topology. Recall that PS and Γ are not closed in general, hence not compact, even though \mathcal{PT}^ω resp. $\tau(\mathcal{PT}^\omega)$ are compact.

Definition 5.12 (Distance of sets). For $A, B \subseteq \mathcal{PT}^\omega$, let $d_P(A, B) = \inf\{d_P(a, b) \mid a \in A, b \in B\}$.

In [19, Theorem 4.3], the following theorem has been proved:

THEOREM 5.13 COMPACT SET DISTANCE CONDITION. *Let $A \subseteq \mathcal{PT}^\omega$ be closed and $B \subseteq \mathcal{PT}^\omega$ be compact with respect to the $[n]$ -view topology. If $A \cap B = \emptyset$, then $d_{[n]}(A, B) > 0$.*

⁸Whenever we state a topological property w.r.t. the subspace topology, we will refer to Γ (resp. PS), otherwise to C^ω (resp. \mathcal{PT}^ω).

We prove the following result, which also holds when A, B are not closed/compact. Note that it also implies Theorem 5.13 as a simple corollary for any P -view topology: Corollary 5.15 shows that it also holds in the minimum topology.

THEOREM 5.14 GENERAL SET DISTANCE CONDITION. *Let $P \subseteq [n]$ with $P \neq \emptyset$. Let A, B be arbitrary subsets of \mathcal{PT}^ω . Then, $d_P(A, B) = 0$ if and only if there are infinite sequences $(a_k) \in A^\omega$ and $(b_k) \in B^\omega$ of process-time graphs as well as $\hat{a}, \hat{b} \in \mathcal{PT}^\omega$ with $a_k \rightarrow \hat{a}$ and $b_k \rightarrow \hat{b}$ with respect to the P -view topology and $d_P(\hat{a}, \hat{b}) = 0$.*

The representation $d_{\min} = \min_{p \in [n]} d_{\{p\}}$ from (2) allows us to extend this result from the P -topologies to the minimum topology:

COROLLARY 5.15. *Let $P \subseteq [n]$ with $P \neq \emptyset$. Let A, B be arbitrary subsets of \mathcal{PT}^ω . Then, $d_{\min}(A, B) = 0$ if and only if there are infinite sequences $(a_k) \in A^\omega$ and $(b_k) \in B^\omega$ of process-time graph sequences as well as $\hat{a}, \hat{b} \in \mathcal{PT}^\omega$ with $a_k \rightarrow \hat{a}$ and $b_k \rightarrow \hat{b}$ with respect to the minimum topology and $d_{\min}(\hat{a}, \hat{b}) = 0$.*

The above Theorem 5.14 allows us to distinguish 3 main cases that cause $d_P(A, B) = 0$: (i) If $\hat{a} \in A \cap B \neq \emptyset$, one can choose the sequences defined by $a_k = b_k = \hat{a} = \hat{b}$, $k \geq 1$. (ii) If $A \cap B = \emptyset$ and $\hat{a} = \hat{b}$, there is a “fair sequence” as the common limit. (iii) If $A \cap B = \emptyset$ and $\hat{a} \neq \hat{b}$, there is a pair of “unfair sequences” acting as limits, which have distance 0 (and are hence also common w.r.t. the pseudo-metric d_P). We note, however, that due to the non-uniqueness of the limits in our pseudo-metric, (iii) are actually two instances of (ii). We kept the distinction for compatibility with the existing results [9, 21] for $n = 2$.

Definition 5.16 (Fair and unfair process-time graph sequences). Consider two process-time graph sequences $r, r' \in \mathcal{PT}^\omega$ of some consensus algorithm with partitions $PS(v)$, $v \in \mathcal{V}_O$, in the P' -topology with $P' = \{p\}$, $p \in [n]$, or $P' = \min$:

- r is called *fair*, if for some $v, w \neq v \in \mathcal{V}_O$ there are convergent sequences $(a_k) \in PS(v)$ and $(b_k) \in PS(w)$ with $a_k \rightarrow r$ and $b_k \rightarrow r$ with respect to $\mathcal{T}_{P'}^{\mathcal{PT}^\omega}$.
- r, r' are called a pair of *unfair* sequences, if for some $v, w \neq v \in \mathcal{V}_O$ there are convergent sequences $(a_k) \in PS(v)$ with $a_k \rightarrow r$ and $(b_k) \in PS(w)$ with $b_k \rightarrow r'$ and $d_{P'}(r, r') = 0$ with respect to $\mathcal{T}_{P'}^{\mathcal{PT}^\omega}$.

An illustration is shown in Figure 5.

The above findings go nicely with the alternative characterization of consensus solvability given in Corollary 5.18, which results from applying the following Lemma 5.17 from [18] to Theorem 5.5.

LEMMA 5.17 SEPARATION LEMMA [18, LEMMA 23.12]. *If Y is a subspace of X , a separation of Y is a pair of disjoint nonempty sets A and B whose union is Y , neither of which contains a limit point of the other. The space Y is connected if and only if there exists no separation of Y .*

COROLLARY 5.18 SEPARATION-BASED CHARACTERIZATION. *Consensus is solvable with a message adversary MA generating the set of admissible process-time graph sequences PS if and only if there exists a partition of PS into sets $PS(v)$, $v \in \mathcal{V}_O$ such that the following holds:*

- (1) *No $PS(v)$ contains a limit point of any other $PS(w)$ w.r.t. the minimum topology in \mathcal{PT}^ω .*

- (2) *Every v -valent admissible sequence z_v satisfies $z_v \in PS(v)$.*

We hence immediately obtain:

COROLLARY 5.19 FAIR/UNFAIR CONSENSUS IMPOSSIBILITY. *The set of admissible process-time graphs PS of a consensus algorithm \mathcal{A} with partitions $PS(v)$, $v \in \mathcal{V}_O$, does not contain any fair process-time graph sequence r or any pair r, r' of unfair process-time graph sequences.*

6 APPLICATIONS

In this section, we will apply our various topological characterizations of consensus solvability to particular classes of examples. We start in Section 6.1 with the broad class of message adversaries that have been shown, or can be shown, to make consensus impossible by means of bivalence proofs. We then proceed with a complete characterization of consensus solvability the class of compact message adversaries in Section 6.2. A characterization for the class of non-compact message adversaries will be provided in Section 6.3.

6.1 Bivalence-based Impossibilities

Our topological results shed some new light on the now standard technique of bivalence-based impossibility proofs introduced in the celebrated FLP paper [10], which have been generalized in [17] and used in many different contexts: Our results reveal that the forever bivalent executions constructed inductively in bivalence proofs such as [22, 23] and [6, 24] are just the common limit of two infinite sequence of executions $\alpha_0, \alpha_1, \dots$ all contained in, say, $\Gamma(0)$ and β_0, β_1, \dots all contained in $\Gamma(1)$ that have a common limit $\alpha_k \rightarrow \hat{\alpha}$ and $\beta_k \rightarrow \hat{\beta}$ in some $\{p\}$ -view topology with $d_{\{p\}}(\hat{\alpha}, \hat{\beta}) = 0$.

More specifically, what is common to these proofs is that one shows that, for any consensus algorithm, there is an admissible forever bivalent run. This is usually done inductively, by showing that there is a bivalent initial configuration and that, given a bivalent configuration C^{t-1} at the end of round $t-1$, there is a 1-round extension leading to a bivalent configuration C^t at the end of round t . By definition, bivalence of C^t means that there are two admissible executions α_t with decision value 0 and β_t with decision value 1 starting out from C^t , i.e., having a common prefix that leads to C^t . Consequently, their distance, in any $\{p\}$ -view topology, satisfies $d_{\{p\}}(\alpha_t, \beta_t) < 2^{-t}$. Note that this is also true for the more general concept of a bipotent configuration C^t , as introduced in [17].

By construction, the $(t-1)$ -prefix of α_t and α_{t-1} are the same, for all t , which implies that they converge to a limit $\hat{\alpha}$ (and analogously for $\hat{\beta}$), see Figure 5 for an illustration. Therefore, these executions match Definition 5.16, and Corollary 5.19 implies that the stipulated consensus algorithm cannot be correct. Concrete examples are the lossy link impossibility [22], i.e., the impossibility of consensus under an oblivious message adversary for $n = 2$ that may choose any graph out of the set $\{\leftarrow, \leftrightarrow, \rightarrow\}$, and the impossibility of solving consensus with vertex-stable source components with insufficient stability interval [6, 24]. In the case of the oblivious lossy link message adversary using the reduced set $\{\leftarrow, \rightarrow\}$ considered in [8], consensus is solvable and there is no forever bivalent run. Indeed, there exists a consensus algorithm, such that all configurations reached after the first round are already univalent.

6.2 Compact Message Adversaries

In this section, we consider message adversaries (like oblivious ones [8, 22]) that are limit-closed, in the sense that every convergent sequence of process-time graphs a_0, a_1, \dots with $a_i \in PS$ for every i has a limit $\hat{a} \in PS$. An illustration is shown in Fig. 4, where the blue dots represent the a_i 's and \times the limit point \hat{a} at the boundary.




Figure 4: Examples of two connected components of the decision sets $PS(0) = PS_{z_0} \cup PS_{z'_0}$ and $PS(1) = PS_{z_1} \cup PS_{z'_1}$ for a compact message adversary. They are closed in \mathcal{PT}^ω , hence contain all their limit points (marked by \times) and have a distance > 0 by Theorem 5.13.

In this case, the set of admissible process-time graph sequences PS is closed and hence a compact subspace both in any P -view topology and in the minimum topology. Moreover, we obtain:

COROLLARY 6.1 DECISION SETS FOR COMPACT MAs ARE COMPACT. *For every correct consensus algorithm for a compact message adversary and every $v \in \mathcal{V}_O$, $PS(v)$ is closed in PS and compact, and $d_{\{p\}}(PS(v), PS(w)) > 0$ for any $v, w \neq v \in \mathcal{V}_O$ and $p \in [n]$, and hence also $d_{\min}(PS(v), PS(w)) > 0$.*

Moreover, there are only finitely many different connected components PS_x , $x \in PS$, which are all compact, and for every x, y with $PS_x \neq PS_y$, it holds that $d_{\{p\}}(PS_x, PS_y) > 0$ and hence also $d_{\min}(PS_x, PS_y) > 0$.

We now make the abstract characterization of Theorem 5.5 and our meta-procedure more operational, by introducing the ε -approximation of the connected component PS_z that contains a process-time graph $z \in PS$, typically for some $\varepsilon = 2^{-t}$, $t \geq 0$. It is constructed iteratively, using finitely many iterations (since the number of different possible t -prefixes satisfies $|\mathcal{PT}^t| < \infty$) of the following algorithm:

Definition 6.2 (ε -approximations). Let $z \in PS$ be an admissible process-time graph. In the minimum topology, we iteratively define PS_z^ε , for $\varepsilon > 0$, as follows: $PS_z^\varepsilon[0] = \{z\}$; for $\ell > 0$, $PS_z^\varepsilon[\ell] = \bigcup_{a \in PS_z^\varepsilon[\ell-1]} (B_\varepsilon(a) \cap PS)$; and $PS_z^\varepsilon = PS_z^\varepsilon[m]$ where $m < \infty$ is such that $PS_z^\varepsilon[m] = PS_z^\varepsilon[m+1]$. For $v \in \mathcal{V}_O$, the ε -approximation $PS^\varepsilon(v)$ is defined as $PS^\varepsilon(v) = \bigcup_{z_v \in PS} PS_{z_v}^\varepsilon$, where every z_v denotes a v -valent process-time graph.

LEMMA 6.3 PROPERTIES OF ε -APPROXIMATION. *For every $\varepsilon > 0$, every $v, w \in \mathcal{V}_O$, every $z \in PS$, v -valent z_v , and every w -valent z_w , the ε -approximations have the following properties:*

- (i) *For a compact message adversary, there are only finitely many different PS_z^ε , $z \in PS$.*
- (ii) *For every $0 < \varepsilon' \leq \varepsilon$, it holds that $PS_{z_v}^{\varepsilon'} \subseteq PS_{z_v}^\varepsilon$.*
- (iii) *$PS_{z_v}^\varepsilon \cap PS_{z_w}^\varepsilon \neq \emptyset$ implies $PS_{z_v}^\varepsilon = PS_{z_w}^\varepsilon$.*

(iv) $PS_z \subseteq PS_z^\varepsilon$.

We now show that PS_x^ε and PS_y^ε for sequences x and y with $PS_x \neq PS_y$ have a distance > 0 , provided ε is sufficiently small:

LEMMA 6.4 SEPARATION OF ε -APPROXIMATIONS FOR COMPACT MAs. *For a compact message adversary that allows to solve consensus, let $x \in PS$ and $y \in PS$ be such that $PS_x \neq PS_y$. Then, there is some $\varepsilon > 0$ such that, for any $0 < \varepsilon' \leq \varepsilon$, it holds that $d_{\min}(PS_x^{\varepsilon'}, PS_y^{\varepsilon'}) > 0$.*

We immediately get the following corollary, which allows us to reformulate Theorem 5.11 as given in Theorem 6.6.

COROLLARY 6.5 MATCHING ε -APPROXIMATION. *For a compact message adversary, if $\varepsilon > 0$ is chosen in accordance with Lemma 6.4, then $PS_z^\varepsilon = PS_z$ for every $z \in PS$.*

THEOREM 6.6 CONSENSUS CHARACTERIZATION FOR COMPACT MAs. *A compact message adversary allows to solve consensus if and only if there is some $\varepsilon > 0$ such that every v -valent $PS_{z_v}^\varepsilon$, $v \in \mathcal{V}_O$, is broadcastable for some process.*

Note carefully that if consensus is solvable, then, for every $0 < \varepsilon' \leq \varepsilon$, the universal algorithm from Theorem 5.5 with $PS(v) = PS^{\varepsilon'}(v) \cup PS \setminus \bigcup_{w \neq v \in \mathcal{V}_O} PS^{\varepsilon'}(w)$ for some arbitrary value $v \in \mathcal{V}_O$, and $PS(w) = PS^{\varepsilon'}(w)$ for the remaining $w \in \mathcal{V}_O$, can be used.

We conclude this section with noting that checking the broadcastability of $PS_{z_v}^\varepsilon$ can be done by checking broadcastability only in *finite* prefixes. More specifically, like the decision function Δ of consensus, the function $T(a)$ that gives the round by which every process in $a \in PS$ has $(p, 0, x_p(a))$ of the broadcaster p in its view is locally constant for a sufficiently small neighborhood, namely, $B_{2^{-T(a)}}(a)$, and hence continuous. Since $PS_{z_v} = PS_{z_v}^\varepsilon$ is compact, $T(a)$ is in fact uniformly continuous and hence attains its maximum \hat{T} in $PS_{z_v}^\varepsilon$. It hence suffices to check broadcastability in the t -prefixes of $PS_{z_v}^\varepsilon$ for $t = \max\{\lfloor \log_2(1/\varepsilon) \rfloor, \hat{T}\}$ in Theorem 6.6.

6.3 Non-compact Message Adversaries

In this section, we finally consider message adversaries that are not limit-closed, like the ones of [9, 21, 24]. Unfortunately, we cannot use the ε -approximations according to Definition 6.2 for non-compact message adversaries. Even if ε is made arbitrarily small, Lemma 6.4 does not hold. An illustration is shown in Fig. 5. Hence, adding a ball $B_\varepsilon(a)$ in the iterative construction of PS_z^ε , where $d_{\min}(a, r) < \varepsilon$ for some forbidden limit sequence r , inevitably lets the construction grow into some $PS_{z'}^\varepsilon$, where z' has a different valence than z . Whereas this could be avoided by adapting ε when coming close to r , the resulting approximation does not provide any advantage over directly using the connected components.

Actually, this is in accordance with the existing solutions for non-compact message adversaries we are aware of, albeit they typically use their complement, namely, excluded sequences. For example, the binary consensus algorithms for $n = 2$ given in [9, 21] assume that the algorithm knows a fair or a pair of unfair sequences a priori, which effectively partition the sequence space into two connected components.⁹ The $(D + 1)$ -VSRC message adversary of [24] even excludes all sequences without a root component that is vertex-stable for at least $D + 1$ rounds, where D is the dynamic diameter

⁹Note that there are uncountably many choices for separating $PS(0)$ and $PS(1)$ here.




Figure 5: Examples of two connected components of the decision sets $PS(0) = PS_{z_0} \cup PS_{z'_0}$ and $PS(1) = PS_{z_1} \cup PS_{z'_1}$ for a non-compact message adversary. They are not closed in \mathcal{PT}^ω and may have distance 0; common limit points (like for PS_{z_0} and PS_{z_1} , marked by \times) must hence be excluded by Corollary 5.18.

that ensures broadcastability by all root members, which renders the remaining connected components broadcastable.

We restate the following necessary and sufficient condition for solving consensus with a non-compact message adversary from Theorem 5.11:

THEOREM 6.7 CONSENSUS CHARACTERIZATION FOR NON-COMPACT MAS. *A non-compact message adversary allows to solve consensus if and only if every v -valent PS_{z_v} , $v \in \mathcal{V}_O$, is broadcastable for some process.*

We will finally prove that the set of to be excluded limit sequences for any decision set $PS(v)$ is compact for a message adversary that allows to solve consensus:

LEMMA 6.8 COMPACTNESS OF EXCLUDED SEQUENCES. *Let $PS(v)$, $v \in \mathcal{V}_O$, be any decision set of a correct consensus algorithm for an arbitrary message adversary, $\overline{PS}(v)$ be its closure in \mathcal{PT}^ω and $\text{Int}(PS)(v)$ its interior. Then, $\overline{PS}(v) = \overline{PS}(v) - \text{Int}(PS)(v)$, which is the set of to be excluded limit points, is compact.*

7 CONCLUSIONS

We provided a complete topological characterization of the solvability of consensus in synchronous directed dynamic networks of arbitrary size controlled by a general message adversary: Consensus can only be solved when the space of admissible process-time graph sequences can be partitioned into at least two non-empty sets that are both closed and open in specific topologies. This requires exclusion of certain fair and unfair limit sequences, which limit broadcastability and happen to coincide with the forever bivalent executions constructed in bivalence and bipotence proofs.

Part of our future work will be devoted to a generalization of our topological framework to other distributed computing models and, most importantly, to other decision problems. Another very interesting area of future research is to study the homology of non-compact message adversaries, i.e., the topological structure of the space of admissible executions using combinatorial topology.

REFERENCES

- [1] Yehuda Afek and Eli Gafni. 2013. Asynchrony from Synchrony. In *Distributed Computing and Networking*, Davide Frey, Michel Raynal, Saswati Sarkar, RudrapatnaK. Shyamasundar, and Prasun Sinha (Eds.). Lecture Notes in Computer Science, Vol. 7730. Springer Berlin Heidelberg, 225–239. https://doi.org/10.1007/978-3-642-35668-1_16
- [2] Bowen Alpern and Fred B. Schneider. 1985. Defining Liveness. *Inform. Process. Lett.* 21, 4 (1985), 181–185.
- [3] Ido Ben-Zvi and Yoram Moses. 2014. Beyond Lamport’s Happened-before: On Time Bounds and the Ordering of Events in Distributed Systems. *J. ACM* 61, 2, Article 13 (April 2014), 26 pages. <https://doi.org/10.1145/2542181>
- [4] Martin Biely, Bernadette Charron-Bost, Antoine Gaillard, Martin Hutle, André Schiper, and Josef Widder. 2007. Tolerating Corrupted Communication. In *Proceedings of the 26th ACM Symposium on Principles of Distributed Computing (PODC’07)*. ACM, Portland, OR, USA, 244–253. <http://www.vmars.tuwien.ac.at/php/pserver/extern/docdetail.php?DID=2240&viewmode=paper&year=2007>
- [5] Martin Biely and Peter Robinson. 2019. On the hardness of the strongly dependent decision problem. In *Proceedings of the 20th International Conference on Distributed Computing and Networking, ICDCN 2019, Bangalore, India, January 04-07, 2019*. 120–123. <https://doi.org/10.1145/3288599.3288614>
- [6] Martin Biely, Peter Robinson, Ulrich Schmid, Manfred Schwarz, and Kyrill Winkler. 2018. Gracefully degrading consensus and k-set agreement in directed dynamic networks. *Theoretical Computer Science* 726 (2018), 41–77. <https://doi.org/10.1016/j.tcs.2018.02.019>
- [7] Bernadette Charron-Bost and André Schiper. 2009. The Heard-Of model: computing in distributed systems with benign faults. *Distributed Computing* 22, 1 (April 2009), 49–71. <https://doi.org/10.1007/s00446-009-0084-6>
- [8] Étienne Coulouma, Emmanuel Godard, and Joseph G. Peters. 2015. A characterization of oblivious message adversaries for which Consensus is solvable. *Theor. Comput. Sci.* 584 (2015), 80–90. <https://doi.org/10.1016/j.tcs.2015.01.024>
- [9] Tristan Fevat and Emmanuel Godard. 2011. Minimal Obstructions for the Coordinated Attack Problem and Beyond. In *25th IEEE International Symposium on Parallel and Distributed Processing, IPDPS 2011, Anchorage, Alaska, USA, 16-20 May, 2011 - Conference Proceedings*. 1001–1011. <https://doi.org/10.1109/IPDPS.2011.96>
- [10] Michael J. Fischer, Nancy A. Lynch, and M. S. Paterson. 1985. Impossibility of Distributed Consensus with one Faulty Process. *J. ACM* 32, 2 (April 1985), 374–382.
- [11] Ronald C. Freiwald. 2014. *An Introduction to Set Theory and Topology*. Washington University in St. Louis. <https://doi.org/10.7936/K7D798QH>
- [12] Maurice Herlihy, Dmitry N. Kozlov, and Sergio Rajsbaum. 2013. *Distributed Computing Through Combinatorial Topology*. Morgan Kaufmann. <https://store.elsevier.com/product.jsp?isbn=9780124045781>
- [13] Maurice Herlihy and Nir Shavit. 1999. The Topological Structure of Asynchronous Computability. *J. ACM* 46, 6 (Nov. 1999), 858–923. <https://doi.org/10.1145/331524.331529>
- [14] F. Kuhn and R. Oshman. 2011. Dynamic networks: Models and algorithms. *SIGACT News* 42(1) (2011), 82–96.
- [15] Petr Kuznetsov, Thibault Rieutord, and Yuan He. 2018. An Asynchronous Computability Theorem for Fair Adversaries. In *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing, PODC 2018, Egham, United Kingdom, July 23-27, 2018*. 387–396. <https://dl.acm.org/citation.cfm?id=3212765>
- [16] Ronit Rubinfeld and Shlomo Moran. 1995. Closed Schedulers: A Novel Technique for Analyzing Asynchronous Protocols. *Distrib. Comput.* 8, 4 (June 1995), 203–210. <https://doi.org/10.1007/BF02242738>
- [17] Yoram Moses and Sergio Rajsbaum. 2002. A Layered Analysis of Consensus. *SIAM J. Comput.* 31, 4 (2002), 989–1021.
- [18] James Munkres. 2000. *Topology (2nd Edition)*. Pearson.
- [19] Thomas Nowak. 2010. *Topology in Distributed Computing*. Master Thesis. Embedded Computing Systems Group, Technische Universität Wien.
- [20] Thomas Nowak, Ulrich Schmid, and Kyrill Winkler. 2019. Topological Characterization of Consensus under General Message Adversaries. *arXiv:1905.09590* (May 2019). <https://arxiv.org/abs/1905.09590> (appeared in Proc. PODC’19).
- [21] Daniel Pflieger. 2018. *Knowledge and Communication Complexity*. Master’s thesis. Technische Universität Wien, Institut für Computer Engineering, Treitlstr. 3/2/182-2, 1040 Vienna, Austria. http://publik.tuwien.ac.at/files/publik_273447.pdf
- [22] Nicola Santoro and Peter Widmayer. 1989. Time is not a healer. In *Proc. 6th Annual Symposium on Theor. Aspects of Computer Science (STACS’89) (LNCS 349)*. Springer-Verlag, Paderborn, Germany, 304–313.
- [23] Ulrich Schmid, Bettina Weiss, and Idit Keidar. 2009. Impossibility Results and Lower Bounds for Consensus under Link Failures. *SIAM J. Comput.* 38, 5 (2009), 1912–1951. <https://doi.org/10.1137/S009753970443999X>
- [24] Kyrill Winkler, Manfred Schwarz, and Ulrich Schmid. 2019. Consensus in Directed Dynamic Networks with Short-Lived Stability. *Distributed Computing* (2019). <https://doi.org/10.1007/s00446-019-00348-0>