

A Topological View of Partitioning Arguments: Reducing k -Set Agreement to Consensus*

Hugo Rincon Galeana¹[0000-0002-8152-1275], Kyrill Winkler²[0000-0002-7310-1748],
Ulrich Schmid²[0000-0001-9831-8583], and Sergio Rajsbaum¹[0000-0002-0009-5287]

¹ UNAM, Ciudad Universitaria, D.F. 04510, Mexico

² TU Wien, ECS Group (E191-02), Treitlstrasse 1-3, 1040 Vienna, Austria

Abstract. The objective of this paper is to understand the effect of partitioning in distributed computing models. In spite of being quite similar agreement problems, (deterministic) consensus (1-set agreement) and k -set agreement (for $k > 1$) require surprisingly different techniques for proving impossibilities. There is a widely applicable generic theorem, however, which allows to reduce the impossibility of k -set agreement to consensus in message-passing models that allow some partitioning. In this paper, we provide the topological representation of this theorem, which reveals how partitioning is reflected in the protocol complex: It turns out that this leads to a “color splitting” of the algorithm’s decision map, which separates the sub-complexes representing the partitioned processes. We also harvest a general advantage of topological results, which allowed us to carry over our findings to shared memory systems. We first demonstrate the utility of our reduction theorem by proving that d -set agreement cannot be solved in the d -solo asynchronous read-write model even when a single process may crash, not just in the wait-free case. Moreover, our new insights into the structure of protocol complexes gave us the idea for a simple proof of the fact that no partitioning argument can provide a valid impossibility proof for wait-free set agreement in the iterated immediate snapshot model: For any set of partition-compatible runs (which do not contain runs where all processes always have a complete view), we provide a way to construct a simple algorithm that solves set agreement.

Keywords: Algebraic topology · consensus · set agreement · partitioning arguments · shared memory.

1 Introduction

Partitions, i.e., sets of processes that cannot always communicate with each other, are a fundamental combinatorial notion in distributed computability [13,16]. Studying the effect of partitioning on the (im)possibility of agreement problems has been a focal point at least since the realization that asynchronous consensus is solvable with a minority of initially crashed processes [15] and the observation

* This work has been supported by the PAPIIT-UNAM grant IN109917 and the Austrian Science Fund (FWF) projects RiSE/SHiNE (S11405) and ADynNet (P28182).

that a shared register cannot be implemented on top of a message passing system with a majority of faulty processes [5]. Since then, partitioning arguments have been applied successfully to many distributed computing problems [14]. Essentially, these arguments exploit the fact that one cannot guarantee agreement among processes of a distributed system that never, neither directly nor indirectly, communicate with each other.

The objective of this paper is to understand the computational power of models where partitioning can occur. For this purpose, we focus on the k -set agreement problem and its special cases consensus and set agreement. In these problems, every process owns a local input value taken from a finite domain \mathcal{V} (to rule out trivial solutions, we assume that $|\mathcal{V}| \geq k$), and must irrevocably assign a local output value (also called decision value), that must be the input value of some process and satisfy certain properties. For consensus, no two processes may decide on, i.e., assign, different values. For set agreement among n processes, the number of different decision values must be at most $n - 1$ system-wide. The general case is k -set agreement, which requires that the number of different decision values is at most k . We will focus solely on deterministic algorithms.

Due to the landmark FLP impossibility result [15], which employs (now classic) bivalence proofs, it is well-known that consensus is impossible to solve in asynchronous systems if a single process may crash. The corresponding result for general k -set agreement is the impossibility of solving this problem in asynchronous systems where $f \geq k$ processes may crash. Surprisingly, establishing the latter result requires quite involved techniques based on algebraic topology resp. a variant of Sperner’s lemma [11,20,22], which have not been matched by combinatorial proofs so far in their full generality (see the related work for some exceptions for special cases).

Despite this apparent “proof-incompatibility” of consensus and k -set agreement for $k > 1$, Biely, Robinson, and Schmid showed in [9] that, in message-passing models that allow partitioning, impossibility proofs for k -set agreement can be reduced to impossibility proofs for consensus: They provided a theorem (called BRS theorem in the sequel), which uses a partitioning argument as a means for reduction and is generic w.r.t. the underlying system model: Essentially, if failures and asynchrony allow for runs where the system partitions into k parts, the processes must decide on their own in every partition. By choosing distinct proposal values, solving k -set agreement in such runs requires solving consensus in every partition. Consequently, the impossibility of k -set agreement can be proved by showing that it is impossible to reach consensus in at least one of these partitions. Note that the BRS theorem actually works for fairly weak forms of partitioning, where some communication between partitions is still possible.

In [9], the authors applied the BRS theorem to various message-passing models, including purely asynchronous systems with crash failures, synchronous systems with omission failures, dynamic networks with omission failures, and even asynchronous systems with failure detectors.

Main contributions: In order to understand the impact of partitioning in distributed computing models, we present a topological version/interpretation of

the BRS theorem and, on the positive side, show that it can even be applied in the shared memory setting. On the negative side, we prove that one cannot show the impossibility of set agreement in the iterated immediate snapshot model based on a partitioning argument. In more detail:

- (1) We provide a topological variant of a slightly generalized version of the BRS theorem and its proof for message-passing systems, which reveals that partitioning is reflected in the protocol complex by a “color splitting” of the algorithm’s decision map that separates the sub-complexes representing the partitioned processes. This insight into the structure of the protocol complexes of partitionable models is of independent interest, as the second major contribution of our paper reveals.
- (2) We exploit the natural genericity of topological results to translate the BRS theorem to the shared memory model. First, we apply it to d -set agreement in the d -solo asynchronous read-write model, where up to d processes may run solo: We provide a simple and illustrative proof that this problem is not just wait-free impossible, as has been shown in [18] already, but even impossible if just a single process may crash. Second, we used our new insights obtained in (1) to prove that one cannot hope to show the impossibility of wait-free set agreement in the iterated immediate snapshot model using any form of partitioning arguments: For any set of partition-compatible runs (which do not contain runs where all processes always have a complete view), we provide a way to construct a simple algorithm that solves set agreement.

In a nutshell, the results of our paper reveal how the partitioning argument implemented by means of the BRS theorem actually works: It is the “color-split” structure of the resulting protocol complex, which effectively allows to avoid a complex global topological analysis and to perform a simple reduction to the solvability of consensus in some partition instead. In the case of set agreement, the existence of any such splitting already guarantees a solution algorithm. We conjecture that we will be able to come up with similar statements also for other problems, in particular, for general k -set agreement. It is important to note, though, that our results apply only to systems where some partitioning can occur. They cannot hence replace results like [11,20,22] in general.

Related work: We are not aware of much research that considers non-trivial reductions of k -set agreement to consensus. However, quite some papers, like [21], prove the impossibility of k -set agreement by partitioning the system into more than k sets of processes that decide independently.

For message-passing systems, besides [9], Biely et al. have employed reduction already in [8] to show that consensus is impossible in certain partially synchronous models, and to prove the tightness of the generalized loneliness failure detector $L(k)$ for k -set agreement. Similar reduction arguments were employed in [12] and, in particular, in [10], where certain k -set agreement runs with disjoint participants are pasted together in order to prove the necessity of the generalized quorum failure detector Σ_k for solving k -set agreement. In [4], a reduction to asynchronous set agreement is used to derive a lower bound on the minimum size of a “synchronous window” that is necessary for k -set agreement.

For shared memory systems, [1] shows the wait-free equivalence of k -set agreement and k -simultaneous consensus using read/write atomic registers. The latter problem allows processes to participate, with the same input, in k independent consensus instances that run simultaneously, but requires a decision only in one of those. Whereas this can be seen as an explicit form of partitioning, it is obviously much less general than the partitioning allowed by the BRS theorem.

Regarding different proof techniques for consensus and k -set agreement, to the best of our knowledge, the only alternatives to the celebrated impossibility proofs for k -set agreement [11,20,22], which are all based on algebraic topology resp. different proofs of Sperner’s lemma, are the alternative combinatorial impossibility proof for k -set agreement in wait-free environments provided in [6] and the counting-based impossibility of general wait-free colored and colorless tasks in [7]. The latter two results do not generalize to $k = f < n - 1$ crash failures, however.

In [3,2], Alistarh et al. described an approach for proving general impossibilities for wait-free tasks in the iterated immediate snapshot model by introducing *extension-based proofs*. The idea is to consider a game between a prover (constructing a schedule) and the protocol (specified by its decision map Δ , unknown to the prover). By constructing a protocol on the fly, via an adversarial strategy w.r.t. the prover, the authors could prove that the wait-free k -set agreement impossibility cannot be established by an extension-based proof. Part (2) of our work differs from [2] in that we restrict our attention to partitioning arguments, i.e., do not aim at general combinatorial proofs (not to speak of bivalence arguments) as does the latter. This restriction greatly reduces the effort needed to show that partitioning arguments are not sufficient for showing the wait-free set agreement impossibility: Rather than adversarially constructing a protocol following the prover’s strategy, the construction of our set agreement protocol just instantiates a simple generic algorithm and is hence relatively easy.

Paper organization: After a short introduction to topological modeling of distributed systems in Section 2, we translate the definitions and concepts underlying the original BRS theorem and prove some basic lemmas in Section 3. In Section 4, we develop our topological version of the BRS theorem, in Section 5, we prove that set agreement can be solved for any set of partition-compatible runs. We conclude in Section 6 with some open questions.

2 Topological Modeling of Distributed Systems

We now briefly describe the basics of modeling distributed systems using algebraic topology, see e.g. [17] for a comprehensive introduction. Whereas this powerful approach became particularly popular for asynchronous shared memory systems [20], it is well-suited for message passing systems as well [19].

We consider a set $\Pi = \{p_1, \dots, p_n\}$ of processes, each with its own unique identifier, which may suffer from crash or omission failures, i.e., may also lose messages. Except in Section 5, where we use the standard iterated snapshot asynchronous shared memory model, we will primarily deal with message passing

protocols, where each process $p_i \in \Pi$ has an individual *message buffer* m_j for received messages from every potential sender process $p_j \in \Pi$ in its local state (p_i, m_1, \dots, m_n) , where (m_1, \dots, m_n) is called p_i 's *view*. Note that m_i is assumed to contain p_i 's local variables. Valid messages are taken from a possibly infinite set M . Typically, we will consider deterministic full information protocols, where processes send messages that consist of the entire history of local states. We represent “ p_2 receives message $m \in M$ from p_1 ” by appending m to the message buffer of p_2 that corresponds to p_1 . We will assume that the initial input is given as a message from a process p_i to itself, and that the decision value is also appended to the message buffer reserved for itself. A global state of the system is a vector of local states, one for each $p_i \in \Pi$.

We define a run of a given protocol as a valid infinite sequence of global states, in which eventually each non-crashing process reaches a final decision state, i.e., a state where it has decided. Note carefully that, since we consider full-history protocols, the final decision state of a process contains its view of the entire run. Clearly, the decision of the process is determined by its local view at the first time it reached a decision state. We call such local views minimal final views. Note that we will assume that processes may still send messages after they have decided; messages that reach a process after a final decision state do not change the process' decision, however. Observe that if run α and run β have the same minimal final views for each process, then they have the same decisions. Therefore, we can restrict our attention to the equivalence classes of runs where two runs are equivalent if all processes have the same minimal final view.

We define a task $T_\Pi = \langle \mathcal{I}, \mathcal{O}, \Delta \rangle$ as a tuple, where \mathcal{I} and \mathcal{O} are chromatic simplicial complexes, which model the valid inputs and outputs for a set Π of processes, and $\Delta : \mathcal{I} \rightarrow 2^{\mathcal{O}}$ is a valid decision function that maps valid input configurations to sets of valid output configurations. Both complexes are chromatic, with coloring χ (formally, a simplicial map from the complex to a simplex of matching dimension, i.e., one that maps simplices to simplices) that attaches a unique label (in fact, a process id from Π) to every vertex such that no two neighbors in any 1-simplex have the same label. Note that we will usually write T instead of T_Π for brevity.

More formally, the input complex $\mathcal{I} = \langle V(\mathcal{I}), F(\mathcal{I}) \rangle$ is given by its set of vertices $V(\mathcal{I})$ and its set of faces $F(\mathcal{I})$: $V(\mathcal{I}) = \{(p_i, v_i) \mid p_i \in \Pi, v_i \in V_i\}$ where V_i is the set of valid inputs for p_i , and $F(\mathcal{I}) = \{\sigma \subseteq V(\mathcal{I}) \mid \sigma \text{ is part of a valid input cfg.}\}$.

The output complex $\mathcal{O} = \langle V(\mathcal{O}), F(\mathcal{O}) \rangle$ is given by its set of vertices $V(\mathcal{O})$ and its set of faces $F(\mathcal{O})$: $V(\mathcal{O}) = \{(p_i, v_i) \mid p_i \in \Pi, v_i \in \hat{V}_i\}$ where \hat{V}_i is the set of valid outputs for p_i , $F(\mathcal{O}) = \{\sigma \subseteq V(\mathcal{O}) \mid \sigma \text{ is part of a valid output cfg.}\}$.

The decision function $\Delta : F(\mathcal{I}) \rightarrow 2^{F(\mathcal{O})}$ is a function with the property that, for every $\sigma \subseteq \rho$ and $\rho' \in \Delta(\rho)$, there is some $\sigma' \subseteq \rho'$ with $\sigma' \in \Delta(\sigma)$. Moreover, Δ is a chromatic map, i.e., $\chi(\Delta(\sigma)) \subseteq \chi(\sigma)$, where $\chi(\sigma)$ gives the set of colors of the vertices in σ and $\chi(S) = \bigcup_{\sigma \in S} \chi(\sigma)$ for every set S of simplices.

We define the protocol complex $\mathcal{P}_M = \langle V(\mathcal{P}_M), F(\mathcal{P}_M) \rangle$ for a given protocol P and model \mathcal{M} as $V(\mathcal{P}_M) = \{(p_i, v_i) \mid p_i \in \Pi, v_i \in \bar{V}_i\}$, where \bar{V}_i is the set of valid minimal final views for p_i in protocol P under a given model \mathcal{M} . The set of

faces is $F(\mathcal{P}_{\mathcal{M}}) = \{\sigma \subseteq V(\mathcal{P}_{\mathcal{M}})\}$, where σ corresponds to a valid configuration of minimal final views of a run. The chromatic function $\chi : V(\mathcal{P}_{\mathcal{M}}) \rightarrow \Pi$ is given by the id of each process, that is $\chi(p_i, v_i) = p_i$. The decision map for a protocol $\mu : V(\mathcal{P}_{\mathcal{M}}) \rightarrow 2^{V(\mathcal{O})}$ is a chromatic vertex map that maps final views of a process to valid outputs for a task.

Since the initial input values are self-messages in our model, each run is produced by a unique configuration of initial input values. Therefore, for each task T , there exists a chromatic simplicial map $i_T : F(\mathcal{P}_{\mathcal{M}}) \rightarrow F(\mathcal{I})$ that maps simplices to simplices of matching dimension, such that $i_T(\sigma)$ is the initial input configuration for each process in σ . A protocol P solves a task T in model \mathcal{M} if and only if the decision map for the protocol is a simplicial map that carries Δ , i.e., ensures $\mu(\sigma) \subseteq \Delta(i_T(\sigma))$. Note that, since the decision map μ needs to be chromatic, it is determined by the mapping values at the facets, i.e., the maximal faces. Therefore, the facet decision map $\hat{\mu} : \hat{F}(\mathcal{P}_{\mathcal{M}}) \rightarrow 2^{V(\mathcal{O})}$, which is just the restriction of μ to the facets $\hat{F}(\mathcal{P}_{\mathcal{M}})$ in $F(\mathcal{P}_{\mathcal{M}})$, fully determines μ .

3 BRS Basic Definitions

In this section, we recast the foundations of the BRS theorem introduced in [9] in our topological framework. Most of the concepts introduced here allow to relate the runs of different algorithms in different models. After all, our purpose is to reduce a run of a k -set agreement algorithm A in some model \mathcal{M} to a run of a consensus algorithm B in some model \mathcal{M}' , with a different set of processes and possibly different synchrony assumptions and failure models. The restricted model only requires that algorithm A is computationally compatible with \mathcal{M}' , i.e., that A can be executed in \mathcal{M}' . Since it is primarily the number of processes in \mathcal{M}' and \mathcal{M} that matter here, we will follow [9] and sloppily write $\mathcal{M} = \langle \Pi \rangle$ and $\mathcal{M}' = \langle D \rangle$ in the sequel. Note carefully, however, that the runs in \mathcal{M}' do not necessarily correspond to runs in \mathcal{M} and vice versa, due to $|\langle D \rangle| \neq |\langle \Pi \rangle|$ and the usually different synchrony assumptions and failure models.

A pivotal concept here is a restricted algorithm.

Definition 1 Restricted algorithm. *Let A be an algorithm for a model $\mathcal{M} = \langle \Pi \rangle$ that consists of the set of processes Π , and $D \subseteq \Pi$ a nonempty set of processes. Consider a restricted model $\mathcal{M}' = \langle D \rangle$. To restrict algorithm A for model \mathcal{M} to an algorithm for model \mathcal{M}' , we just drop all messages sent from D to the outside. We call the restricted algorithm $A|_D = B$.*

Restricted algorithms induce protocol complexes with specific properties. However, even if protocol B corresponds to a restriction of A to D , since message buffers for processes not in D are not present in \mathcal{B} (the protocol complex of B), \mathcal{B} is strictly different from any protocol complex that includes Π in its set of processes. Therefore, we need to define a way to extend the views in \mathcal{B} in a way that could possibly match a protocol complex with processes Π executed in \mathcal{M} . The natural way of doing this is by adding empty message buffers denoted by \perp for any process not in D .

Definition 2 Protocol complexes of restricted algorithms. *Given an algorithm A for \mathcal{M} and a restricted algorithm B for \mathcal{M}' , let \mathcal{B} be the protocol complex for B executed in \mathcal{M}' . We define the extended complex of B with respect to Π , \mathcal{A}_D , as follows: $V(\mathcal{A}_D) = \{(p, w, \perp, \dots, \perp) \mid (p, w) \in V(\mathcal{B})\}$ and \perp represents empty message buffers for processes in $\Pi \setminus D$, $F(\mathcal{A}_D) = \{\sigma \subseteq V(\mathcal{A}_D) \mid \exists \hat{\sigma} \in F(\mathcal{B}), (p, w, \perp, \dots, \perp) \in \sigma \Rightarrow (p, w) \in \hat{\sigma}\}$.*

The following Lemma 1 shows that \mathcal{A}_D is isomorphic to \mathcal{B} , i.e., \mathcal{A}_D is an “extended view copy” of \mathcal{B} . This isomorphic copy \mathcal{A}_D of the protocol complex \mathcal{B} will turn out to be essential, since it corresponds to a subcomplex of \mathcal{A} under certain conditions (Definition 4). Note carefully that, per se, this is not necessarily the case as, e.g., the synchrony model for \mathcal{M} may forbid empty message buffers.

Lemma 1 Isomorphic complex. *Let \mathcal{A}_D and \mathcal{B} be as defined above. Then there exists a chromatic bijective simplicial map $\mu : \mathcal{A}_D \rightarrow \mathcal{B}$.*

Proof. Consider the vertex map $\mu : V(\mathcal{A}_D) \rightarrow V(\mathcal{B})$, $(p, s_w, m_1, \dots, m_d, \perp, \dots, \perp) \mapsto (p, s_w, m_1, \dots, m_d)$. We define $\gamma : V(\mathcal{B}) \rightarrow V(\mathcal{A}_D)$ as $(p, s_w, m_1, \dots, m_d) \mapsto (p, s_w, m_1, \dots, m_d, \perp, \dots, \perp)$. Note that γ is well defined and is the inverse function of μ . This makes μ a bijection. It only remains to show that μ is simplicial and color preserving.

Consider $\sigma \in F(\mathcal{A}_D)$, from the definition of $F(\mathcal{A}_D)$, there exists $\hat{\sigma}$ such that $(p, w, \perp, \dots, \perp) \in \sigma \Rightarrow (p, w) \in \hat{\sigma}$. It is clear that μ maps $\sigma \mapsto \hat{\sigma}$. Since $\hat{\sigma} \in F(\mathcal{B})$, μ is indeed a simplicial map. Note that, by definition, μ preserves the process id, its local state and all message buffers in D . Therefore μ is a color preserving simplicial map.

The following definition captures the notion of indistinguishability of runs:

Definition 3 Indistinguishability of runs. *Runs α and β are indistinguishable for a process p if p has the same sequence of states in α and β until p decides. For a non-empty set D of processes, we say that $\alpha \stackrel{D}{\sim} \beta$ if α is indistinguishable from β until decision for all $p \in D$.*

Note that since p has the same sequence of states until decision for runs α and β , then the minimal final view for p (that contains the full history) is the same for both runs α and β . This means that the simplices σ_α and σ_β that correspond to runs α and β share vertex $(p, s_w, m_1, \dots, m_k)$, where (s_w, m_1, \dots, m_k) corresponds to the minimal final view of p at both runs α and β . This translates naturally to $D\text{-skel}(\sigma_\alpha) = D\text{-skel}(\sigma_\beta)$, where $D\text{-skel}(\sigma_\alpha) = \{(p, w) \in \sigma_\alpha : p \in D\}$.

Definition 4 Compatibility of Runs. *Let \mathcal{R} and \mathcal{R}' be sets of runs, possibly from system models with different synchrony assumptions and failure models. Runs \mathcal{R}' are compatible with runs \mathcal{R} for processes in D , denoted by $\mathcal{R}' \preceq_D \mathcal{R}$, if $\forall \alpha \in \mathcal{R}' \exists \beta \in \mathcal{R} : \alpha \stackrel{D}{\sim} \beta$.*

If \mathcal{R} and \mathcal{R}' are sets of runs for the same protocol A , and in the same model \mathcal{M} , then both induce subcomplexes of a common protocol complex \mathcal{A} . We will call those subcomplexes $\overline{\mathcal{R}}$ and $\overline{\mathcal{R}'}$ respectively. We define $D\text{-skel}(\overline{\mathcal{R}'})$ as the subcomplex of $\overline{\mathcal{R}'}$ where all vertices correspond to processes of D with views from $\overline{\mathcal{R}'}$. Given these definitions, it is clear that $\mathcal{R}' \preceq_D \mathcal{R}$ if and only if $D\text{-skel}(\overline{\mathcal{R}'}) \subseteq D\text{-skel}(\overline{\mathcal{R}})$. Note that the applicability of Definition 3 is limited, as it only works for sets of runs from the same protocol in the same model. However, we will give a more general definition below, which provides some correspondence between runs of *different* protocols in *different* models: Herein, \mathcal{M} and \mathcal{M}' only need to be computationally compatible, but may otherwise differ in the number of processes, synchrony assumptions, failure models, etc.

Definition 5 D -View embedding. *Let A and B be protocols with a non-empty set of common processes S , and $D \subseteq S$ with $s = |S|$, and let \mathcal{A} and \mathcal{B} be the protocol complexes corresponding to A 's runs in model \mathcal{M} (with $s + r$ processes, $r \geq 0$) and B 's runs in model \mathcal{M}' (with $s + k$ processes, $k \geq 0$), respectively. Consider sets of runs \mathcal{R} and \mathcal{R}' from protocol A in model \mathcal{M} and B in model \mathcal{M}' respectively, and the corresponding subcomplexes $\overline{\mathcal{R}} \subseteq \mathcal{A}$ and $\overline{\mathcal{R}'} \subseteq \mathcal{B}$. We say that $\overline{\mathcal{R}'}$ is D -view embedded in $\overline{\mathcal{R}}$, if for every $(p, w_s, m_1, \dots, m_k) \in D\text{-skel}(\overline{\mathcal{R}'})$ with w_s denoting the message buffers for the processes in S the following holds for every $1 \leq i \leq k$: (i) $p_i \notin S \Rightarrow m_i = \perp$; (ii) There exists $(p, w_s, m'_1, \dots, m'_r) \in V(\overline{\mathcal{R}})$ such that*

$$m'_j = \begin{cases} m_j & \text{if } p_j \in S, \\ \perp & \text{if } p_j \notin S; \end{cases}$$

(iii) $\mu : D\text{-skel}(\overline{\mathcal{R}'}) \rightarrow \overline{\mathcal{R}}$ defined by $(p, w_s, m_1, \dots, m_k) \mapsto (p, w_s, m'_1, \dots, m'_r)$ is a simplicial map. Note that μ is an embedding of $D\text{-skel}(\overline{\mathcal{R}'})$, i.e. an injective simplicial map.

If both sets of runs come from the same protocol ($B = A$) and the same synchrony model ($\mathcal{M} = \mathcal{M}'$), then the embedding is given by the inclusion $\iota : D\text{-skel}(\overline{\mathcal{R}'}) \rightarrow \overline{\mathcal{R}}$ with $\iota(p, w) = (p, w)$. This matches with the previous observation that $\mathcal{R}' \preceq_D \mathcal{R}$ if and only if $D\text{-skel}(\overline{\mathcal{R}'}) \subseteq D\text{-skel}(\overline{\mathcal{R}})$ in this case. More generally, we can formulate compatibility of runs in terms of D -view embeddings.

Lemma 2 Compatible runs are D -view embedded. *Let \mathcal{R} and \mathcal{R}' be sets of runs from algorithms A and B in model \mathcal{M} and \mathcal{M}' respectively. Let S be the set of common processes for \mathcal{R} and \mathcal{R}' and $D \subseteq S$. Then $\mathcal{R}' \preceq_D \mathcal{R} \Leftrightarrow \overline{\mathcal{R}'}$ is D -view embedded in $\overline{\mathcal{R}}$.*

Proof. We will show first that $\mathcal{R}' \preceq_D \mathcal{R} \Rightarrow \overline{\mathcal{R}'}$ is D -view embedded in $\overline{\mathcal{R}}$. For $(p, w_s, m_1, \dots, m_k) \in V(D\text{-skel}(\overline{\mathcal{R}'}))$, recall that $p \in D$ and (w_s, m_1, \dots, m_k) corresponds to a minimal final view for a run α in \mathcal{R}' . Note that since $\mathcal{R}' \preceq_D \mathcal{R}$, there exists a run $\beta \in \mathcal{R}$ with $\alpha \stackrel{D}{\sim} \beta$. Then, there exists $(p, w'_s, m'_1, \dots, m'_r) \in V(\overline{\mathcal{R}})$ that corresponds to the minimal final view of p for run β . Since $(w'_s, m'_1, \dots, m'_r)$ and (w_s, m_1, \dots, m_k) are indistinguishable for p , their local states are the same,

$w_s = w'_s$, message buffers for common processes match, and message buffers for processes that are not common must be empty. From these observations it follows that $m_j = \perp$ for all $p_j \notin S$ and that

$$m'_j = \begin{cases} m_j & \text{if } p_j \in S, \\ \perp & \text{if } p_j \notin S. \end{cases}$$

Therefore $\mu : D\text{-skel}(\overline{\mathcal{R}'}) \rightarrow \overline{\mathcal{R}}$ is well defined. To show that it is an embedding of $D\text{-skel}(\overline{\mathcal{R}'})$ we need to show that it is injective and simplicial. Suppose that $(p, w_s, m_1, \dots, m_k)$ and $(p, w'_s, m'_1, \dots, m'_k)$ both map to $(p, \hat{w}_s, \hat{m}_1, \dots, \hat{m}_r)$. From the definition, then $w_s = \hat{w}_s = w'_s$. It is also required that the message buffers for processes not in S be empty, therefore, we only need to be concerned about message buffers for processes in S being the same. However, from the definition we also have that $m_j = \hat{m}_j = m'_j$ for all $p_j \in S$. Therefore μ is injective. It only remains to show that μ is a simplicial map. Since by definition μ preserves process ids (it is a chromatic vertex map), we only need to check that facets (maximal simplices) of $D\text{-skel}(\overline{\mathcal{R}'})$ are mapped inside a facet in $\overline{\mathcal{R}}$. Consider a facet σ' in $D\text{-skel}(\overline{\mathcal{R}'})$, from the definition of \mathcal{R}' , there exists a run α that corresponds to σ' . Since $\mathcal{R}' \preceq_D \mathcal{R}$, there exists a run β , and a facet $\sigma \in \overline{\mathcal{R}}$ such that $\mu(\sigma') \subseteq \sigma$. Therefore μ is a D -view embedding from $\overline{\mathcal{R}'}$ into $\overline{\mathcal{R}}$.

Conversely, assume that $\overline{\mathcal{R}'}$ is D -view embedded into $\overline{\mathcal{R}}$. Then there exists a D -view embedding $\mu : D\text{-skel}(\overline{\mathcal{R}'}) \rightarrow \overline{\mathcal{R}}$. It follows from the definition of a D -view embedding and from the definition of $D\text{-skel}(\overline{\mathcal{R}'})$ that $\mathcal{R}' \preceq_D \mathcal{R}$.

The following Definition 6 is crucial for expressing the consequences of partitioning in our topological setting. Essentially, it says that it is reflected by a splitting of the decision map.

Definition 6 Decision map split. *Let \mathcal{A} be the protocol complex for a given algorithm A on a model $\mathcal{M} = \langle \Pi \rangle$ and \mathcal{A}' a non-empty subcomplex of \mathcal{A} . Let $D \subseteq \Pi$ be a set of processes in \mathcal{M} , $\overline{D} = \Pi \setminus D$ and $B = A|_D$ the restriction of algorithm A to D in a given model $\mathcal{M}' = \langle D \rangle$ with possibly different synchrony assumptions and failure models, resulting in protocol complex \mathcal{B} . We say that D splits the decision map of A at \mathcal{A}' with respect to \mathcal{M}' if \mathcal{B} is D -view embedded in \mathcal{A}' and*

$$\mu|_{\mathcal{A}'}(p, w) = [\mu_D * \mu_{|\overline{D}}](p, w) := \begin{cases} \mu_D(p, w) & \text{if } p \in D, \\ \mu_{|\overline{D}}(p, w) & \text{if } p \in \overline{D}. \end{cases}$$

Herein, $\mu|_{\mathcal{A}'}$ is the decision map μ of \mathcal{A} restricted to \mathcal{A}' , μ_D is the decision map for restricted algorithm B at the extended complex \mathcal{A}_D , and $\mu_{|\overline{D}}$ is the decision map μ restricted to $\overline{D}\text{-skel}(\mathcal{A}')$.

An illustration of decision map splitting for a fixed $D = \{r, q\}$, $\overline{D} = \{p\}$ can be found in Fig. 1b. Note that the subcomplex for D (at the bottom) is the full 1-round subdivided complex for 2 processes, whereas the subcomplex for \overline{D} represents a process that only hears from itself. Fig. 1a shows the analogous splitting in the full one-round complex.




Fig. 1: 1-round protocol complex of d -solo runs for $n = 3, d = 2$: (a) full complex, (b) zoom-in for $D = \{r, q\}, \bar{D} = \{p\}$. Nodes represent local process states (views), with the color encoding the respective process and with the 3-tuple encoding (m_p, m_q, m_r) .

The following Lemma 3 shows that decision map splitting is equivalent to the extended complex of the restricted algorithm being equal to the corresponding D -skeleton:

Lemma 3 Decision map splitting condition. *Let \mathcal{A} be the protocol complex for a given algorithm A in a model $\mathcal{M} = \langle \Pi \rangle$ and \mathcal{A}' a non-empty subcomplex of \mathcal{A} . Let $D \subseteq \Pi$ be a set of processes in \mathcal{M} , $B = A|_D$ a restriction of algorithm A to D in a model \mathcal{M}' with possibly different synchrony assumptions and failure assumptions, \mathcal{A}_D the extended complex of B with respect to Π and μ be the decision map for \mathcal{A} , then D splits μ at \mathcal{A}' with respect to $\mathcal{M}' \Leftrightarrow \mathcal{A}_D = D\text{-skel}(\mathcal{A}')$.*

Proof. Suppose that D splits μ at \mathcal{A}' . Then $\mu|_{\mathcal{A}'} = \mu_D * \mu_{|\bar{D}}$. Therefore,

$$\mu|_{\mathcal{A}'}(p, w) = \begin{cases} \mu_D(p, w) & \text{if } p \in D, \\ \mu(p, w) & \text{otherwise.} \end{cases}$$

By definition, if $(p, w) \in V(\mathcal{A}')$ and $p \in D$, since μ_D can be applied to (p, w) , then $(p, w) \in \text{Dom}(\mu_D) = \mathcal{A}_D$. Therefore, $D\text{-skel}(\mathcal{A}') \subseteq \mathcal{A}_D$. Conversely, by Definition 5, \mathcal{B} is D -view embedded in \mathcal{A}' . Therefore, by Lemma 2, \mathcal{A}_D is a subcomplex of $D\text{-skel}(\mathcal{A}')$. Consequently $D\text{-skel}(\mathcal{A}') = \mathcal{A}_D$.

Now let's assume that $D\text{-skel}(\mathcal{A}') = \mathcal{A}_D$. Therefore in \mathcal{A}' , each process from D reached a decision before receiving messages from \bar{D} . Note that, since \mathcal{A}_D is isomorphic to \mathcal{B} , the processes from D at \mathcal{A}' decide in the same way as the restricted algorithm. Therefore,

$$\mu(p, w) = \begin{cases} \mu_D(p, w) & \text{if } p \in D, \\ \mu_{|\bar{D}}(p, w) & \text{if } p \in \bar{D}. \end{cases}$$

This shows that $\mu = \mu_D * \mu_{\overline{D}}$, and since $D\text{-skel}(\mathcal{A}') = \mathcal{A}_D$, the protocol complex of the restricted algorithm \mathcal{B} is D -view embedded in \mathcal{A}' . Hence D splits μ at \mathcal{A}' with respect to \mathcal{M}' .

We will show in the next section that decision map splitting is equivalent to finding a partition of Π into D, \overline{D} and a set of runs where D decides independently from \overline{D} . However, notice that \overline{D} could use information from D to decide.

4 Topological BRS Theorem

Since we have already established an equivalence between topological conditions and run compatibility, we can proceed to state a slightly more general version of the BRS theorem from a topological perspective. The BRS theorem requires that conditions (A)–(D), as stated in Theorem 1 below, hold, in order to guarantee the k -set agreement impossibility. In the following lemmas, we will state topological properties that are slightly weaker than the original conditions (A)–(D).

Let $\mathcal{M} = \langle \Pi \rangle$ be a model and A an algorithm that runs in \mathcal{M} . Assume that each $p \in \Pi$ starts with an input value taken from a set of at least k different values.³ Also assume that there is a distinguished set $D \subseteq \Pi$, and a partition of $\Pi \setminus D = \overline{D}$ given by D_1, \dots, D_{k-1} . Let $\{v_1, \dots, v_{k-1}\}$ be a fixed set of different values. Let \mathcal{M}_A denote the runs of algorithm A in model \mathcal{M} .

Theorem 1 Original BRS theorem [9]. *We consider the following runs of algorithm A in model \mathcal{M} and a restricted model $\mathcal{M}' = \langle D \rangle$:*

- dec-D** *Any $p_j \in D$ receives no messages from any process in \overline{D} until every process in D has decided.*
- dec- \overline{D}** *For every D_i , there is some $q \in D_i$ that decides v_i , which was proposed by some $p \in \overline{D}$.*

$\mathcal{R}_{(D)}$ denotes the set of runs from \mathcal{M} where **dec-D** holds, $\mathcal{R}_{(D, \overline{D})}$ denotes the set of runs from \mathcal{M} where both **dec-D** and **dec- \overline{D}** hold. If all of the following conditions (A) $\mathcal{R}_{(D)}$ is nonempty, (B) $\mathcal{R}_{(D)} \preceq_D \mathcal{R}_{(D, \overline{D})}$, (C) Consensus is not solvable in \mathcal{M}' , (D) $\mathcal{M}'_{A|_D} \preceq_D \mathcal{M}_A$, hold, then A does not solve k -set agreement.

We generalize the statement of the BRS theorem by slightly relaxing conditions (A)–(D):

Lemma 4 Condition equivalence. *Let $\mathcal{M}, \mathcal{M}', A, D, \overline{D}$ and D_i be as defined for the BRS theorem. Then, (A)–(D) imply (A') $\mathcal{R}_{(D)}$ is nonempty, (B') Consensus is not solvable in \mathcal{M}' , (C') $\mathcal{M}'_{A|_D} \preceq_D \mathcal{R}_{(D, \overline{D})}$.*

³ Note that the original BRS theorem actually assumed that every process starts with a unique value.

Proof. Note that (A') and (A) are the same, also (B') and (C). Therefore it is enough to show that (B) and (D) imply (C'). Let γ be a run in $\mathcal{M}'_{A|D}$. From condition (D), there exists a run $\gamma' \in \mathcal{M}_A$ such that $\gamma \stackrel{D}{\sim} \gamma'$. Since γ' is indistinguishable from γ until decision for D , each process $p \in D$ in run γ' decides before receiving a message from \bar{D} . Therefore, $\gamma' \in \mathcal{R}_{(D)}$. Also note that from (B) we get that there exists $\gamma'' \in \mathcal{R}_{(D, \bar{D})}$ such that $\gamma' \stackrel{D}{\sim} \gamma''$. Since $\stackrel{D}{\sim}$ is transitive, we have $\gamma \stackrel{D}{\sim} \gamma''$. This proves that (C') holds.

The following Lemma 5 is the key technical lemma for the topological equivalence of a slightly stronger version of the BRS theorem established in Theorem 2 below.

Lemma 5 BRS equivalence. *Let $\mathcal{M}, \mathcal{M}', A, D, \bar{D}$ and D_i be as defined for the BRS theorem. Then (A')–(C') are equivalent to the following :*

- (1) *There exists a non empty subcomplex \mathcal{A}' of \mathcal{A} such that $D\text{-skel}(\mathcal{A}') = \mathcal{A}_D$ (the extended complex for $A|D$ with respect to \mathcal{M}).*
- (2) *For each D_i , the decision map $\mu|_{\mathcal{A}'}$ maps every view from D_i into a decision configuration that includes v_i as a decision value.*
- (3) *Each v_i is the input value for some process $p \in \bar{D}$ at subcomplex \mathcal{A}' .*
- (4) *Consensus is not solvable in $\mathcal{M}' = \langle D \rangle$.*

Proof. First we will show that (A')–(C') \Rightarrow (1)–(4): It follows from (C'), that \mathcal{B} is D -view embedded in $\bar{\mathcal{R}}_{(D, \bar{D})}$, where $\bar{\mathcal{R}}_{(D, \bar{D})}$ is the subcomplex of \mathcal{A} that corresponds to runs $\mathcal{R}_{(D, \bar{D})}$. This implies that $\mathcal{A}_D \subseteq D\text{-skel}(\bar{\mathcal{R}}_{(D, \bar{D})})$. We define $F(\mathcal{A}') = \{\sigma \in F(\bar{\mathcal{R}}_{(D, \bar{D})}) : D\text{-skel}(\sigma) \in F(\mathcal{A}_D)\}$, $V(\mathcal{A}') = \{(p, w) \in V(\mathcal{A}) : (p, w) \in \sigma \in F(\mathcal{A}')\}$. Note that this trivially implies $\mathcal{A}_D \subseteq \mathcal{A}'$, so $\mathcal{A}' \neq \emptyset$. It also follows from the definition of \mathcal{A}' that $\mathcal{A}' \subseteq \bar{\mathcal{R}}_{(D, \bar{D})}$, since $F(\mathcal{A}') \subseteq F(\bar{\mathcal{R}}_{(D, \bar{D})})$. Therefore condition (1) holds for \mathcal{A}' . Conditions (2) and (3) follow from $\mathcal{A}' \subseteq \bar{\mathcal{R}}_{(D, \bar{D})}$, condition (4) is the same as (B').

We will now show that (1)–(4) \Rightarrow (A')–(C'): From (1) we get that \mathcal{B} is D -view embedded in \mathcal{A}' since $\mathcal{A}_D = D\text{-skel}(\mathcal{A}')$. From conditions (2) and (3) we get that $\mathcal{A}' \subseteq \bar{\mathcal{R}}_{(\bar{D})}$, therefore $\mathcal{A}_D \subseteq \bar{\mathcal{R}}_{(\bar{D})}$, where $\bar{\mathcal{R}}_{(\bar{D})}$ is the complex of all runs where $\text{dec-}\bar{D}$ holds. Note that since $\mathcal{A}_D = D\text{-skel}(\mathcal{A}')$ is the extended view complex of the restricted algorithm, for each run in \mathcal{A}' , the processes from D reached a decision before receiving messages from \bar{D} . Since $\mathcal{A}' \subseteq \bar{\mathcal{R}}_{(D)}$, we therefore get $\mathcal{A}' \subseteq \bar{\mathcal{R}}_{(D, \bar{D})}$. Since \mathcal{A}' is non-empty, it follows that $\emptyset \neq \mathcal{R}_{(D, \bar{D})} \subseteq \mathcal{R}_{(D)}$. This shows that (A') holds. Condition (4) is the same as condition (B'). Finally, since $\mathcal{A}_D = D\text{-skel}(\mathcal{A}') \subseteq \bar{\mathcal{R}}_{(D, \bar{D})}$, condition (C') holds.

The following Theorem 2 provides the main result of this section, the topological version of a generalization of the BRS theorem.

Theorem 2 Decision split theorem. *Let $\mathcal{M} = \langle \Pi \rangle$ be a model and A an algorithm that runs in \mathcal{M} . Let \mathcal{A} be the protocol complex of A , $\mathcal{M}' = \langle D \rangle$ be a model with D a subset of Π , and μ the decision map for \mathcal{A} . Assume that the following conditions hold:*

- (a) *There exists a non-empty subcomplex \mathcal{A}' , such that D splits μ at \mathcal{A}' with respect to \mathcal{M}' .*
- (b) *Consensus is not solvable in \mathcal{M}' .*
- (c) *Processes in $\overline{D} = \Pi \setminus D$ always decide at least $k - 1$ input values from \overline{D} for runs in \mathcal{A}' .*

Then, A does not solve k -set agreement.

Proof. Assume by contradiction that A solves k -set agreement. In particular $\mu|_{\mathcal{A}'}$ always decides on at most k input values. Note that since μ always decides at least $k - 1$ input values from \overline{D} for runs in \mathcal{A}' , it follows that $\mu|_{\overline{D}}$ also decides on at least $k - 1$ input values from \overline{D} . From Condition (a) and Definition 6, we have that $\mu|_{\mathcal{A}'} = \mu_D * \mu|_{\overline{D}}$. Since all possible decision values from μ_D are disjoint from at least $k - 1$ decision values in $\mu|_{\overline{D}}$, and $\mu|_{\mathcal{A}'}$ solves k -set agreement, μ_D always decides on at most 1 decision value. Therefore μ_D solves consensus at $D\text{-skel}(\mathcal{A}') = \mathcal{A}_D$. Since $\mathcal{A}_D \cong \mathcal{B}$ by Lemma 1, where \mathcal{B} is the protocol complex for algorithm B in model \mathcal{M}' , consensus is solvable in \mathcal{M}' . This contradicts Condition (b), however.

We conclude this section with some corollaries of Theorem 2.

Corollary 1 . *Let \mathcal{M} , \mathcal{M}' , A , D , \overline{D} and D_i be as defined for the original BRS theorem (Theorem 1). If conditions (1)–(4) given in Lemma 5 hold, then A does not solve k -set agreement.*

Proof. First, the equivalence of decision splitting and $\mathcal{A}_D = D\text{-skel}(\mathcal{A}')$ established in Lemma 3 implies that (a) in Theorem 2 is equivalent to

- (a') There exists a non-empty subcomplex \mathcal{A}' , such that $D\text{-skel}(\mathcal{A}') = \mathcal{A}_D$.

Note that Conditions (1) and (4) correspond to (a') and (b) of Theorem 2, respectively; (2) and (3) imply condition (c). Therefore, Conditions (1)–(4) imply (a),(b) and (c) in Theorem 2, which shows that A does not solve k -set agreement.

Corollary 2 . *Let \mathcal{M} , \mathcal{M}' , A , D , \overline{D} and D_i be as defined for the original BRS theorem (Theorem 1). If conditions (A')–(C') in Lemma 4 hold, then A does not solve k -set agreement.*

Proof. It follows from Lemma 5 and Corollary 1 that A does not solve k -set agreement.

5 Partition Compatibility in Shared Memory and Set Agreement

In the previous sections, we developed a topological version of the BRS theorem [9], which allows to reduce the impossibility of k -set agreement to the impossibility of consensus in a wide variety of message-passing systems. In this section, we will transfer some of the resulting insights to the standard *asynchronous shared*

memory (ASM) model. We thus consider a set of processes $\Pi = \{p_1, \dots, p_n\}$ and a shared snapshot memory $M = (e, m_1, \dots, m_n)$, where e is a buffer for global shared variables [not used in our protocols], and each m_i corresponds to the local memory portion of process p_i in the snapshot memory. To simplify our reasoning, we will assume that the protocols are full information immediate snapshot layered protocols, which does not change the computability power of our algorithms [20,17].

In this model, each process executes a predefined number r of asynchronous rounds, called layers. Each layer i consists of concurrent write-read snapshots. A write-read snapshot at layer i consists of writing the full view (the complete local state) of a process into its corresponding part of the memory, and immediately after writing, taking a snapshot of the views from processes at the same layer i . The initial view of a process consists only of its input value; therefore, during the first round, each process only writes its input value to the shared memory. Note that each process' current view contains the history of previous views, so we need not be concerned with overwriting previous views in the shared memory. The final view of a process in an r -round protocol is its view after round r , and the protocol's decision map μ maps the process's final view to some output value.

Definition 7 Views. *The view for a given process p at round k is defined as follows: If $k = 0$, the view consists of a tuple (p, s, v) , where p is the process id, s is the initial local state of p and v is the input value for p . If $k > 0$, the view consists of a tuple (p, s, v_1, \dots, v_n) , where each v_j corresponds to either the view of process p_j at the end of round $k - 1$ if the write-read execution for round k happened before or at the same time as the write-read execution for process p , or else \perp , which represents that p finished its write-read in round k before p_j .*

Note that this definition of the processes' views is extremely useful, since it has a nice combinatorial structure: Herlihy and Shavit showed in [20] that there is an isomorphism between the standard chromatic subdivision of the input complex and the protocol complex for a general 1-layer immediate snapshot protocol. Since each layer i is only determined by the previous layers and the scheduling of layer i , it follows by induction that a $k + 1$ layered protocol complex corresponds to the $k + 1$ -th chromatic subdivision of the input complex.

We start our considerations by using the BRS theorem to show that d -set agreement cannot be implemented in the d -solo model introduced in [18] if at most d processes may crash. In the d -solo model for asynchronous shared memory, up to d processes may run solo, i.e., have no *other* process in their view in every round of a run. The wait-free 1-solo model is equivalent to asynchronous read-write shared memory, and any d -solo model can be simulated in the 1-solo model. Since d -set agreement for $d < n$ cannot be implemented in the wait-free read-write model [20], it cannot be implemented in the wait-free d -solo model either. The following Theorem 3 shows that this does not change if one strengthens the model by allowing only up to a single crash.

Theorem 3 d -set agreement impossibility in d -solo model. *In the d -solo model, it is impossible to solve d -set agreement if just a single process may crash, not even with a colored, i.e., non-anonymous, algorithm.*

Proof. We find it instructive to present the proofs of this result both using the standard version of the BRS theorem (Corollary 2 and Theorem 1) and our topological variant (Theorem 2).

We use Corollary 2 with $D_1 = \{p_1\}, \dots, D_{d-1} = \{p_{d-1}\}$ and $D = \{p_d, \dots, p_n\}$, and assume that the solo processes are p_1, \dots, p_d and start with input value $x_i = i$ (constituting p_i 's initial view v_i) in all runs. The conditions required in Theorem 1 are verified as follows: (A') $\mathcal{R}_{(D)}$ is nonempty, since the processes in $\bar{D} = D_1 \cup \dots \cup D_{d-1}$ may take their first step only after all the processes in D have decided. (B') Consensus is not solvable in the model \mathcal{M}' of asynchronous read-write shared memory (which allows one solo process) where one process may crash. To prove (C') $\mathcal{M}'_{A|D} \preceq_D \mathcal{R}_{(D, \bar{D})}$, consider a run $\rho' \in \mathcal{M}'_{A|D}$. We can find the following run $\rho \in \mathcal{R}_{(D, \bar{D})}$, satisfying $\rho' \stackrel{D}{\sim} \rho$: The processes in D take their steps exactly as in ρ' , and all the $d - 1$ processes in \bar{D} start to run solo only after all processes in D have decided. Clearly, every process $p_i \in \bar{D}$ must decide on its own value i , as a solo process never has another process in its view in any round. As all conditions (A')–(C') of Lemma 4 are satisfied, Corollary 2 proves that d -set agreement is impossible. Since this holds for colored algorithms, the impossibility also extends to colorless ones.

On the other hand, the conditions required in Theorem 2 are verified as follows: (a) holds, since a d -solo execution splits the decision map μ into $\mu_{\bar{D}} * \mu_D$ by definition; herein, $\mu_{\bar{D}}$ and μ_D correspond to the decision maps at \bar{D} and D respectively. Fig. 1a shows the full 1-round subcomplex for $n = 2, d = 2$, Fig. 1a zooms into the subcomplex for $D = \{r, q\}, \bar{D} = \{p\}$. Note carefully that the subcomplex at the bottom of the latter is the full subdivided complex for the two processes in $D = \{r, q\}$. Condition (b) is the same as (B'), and (c) is implied by the fact that the processes in \bar{D} run solo and must hence decide on their own value, as already argued. Therefore, d -set agreement is impossible by Corollary 2.

The remaining part of this section is devoted to the very different result of translating the insights gained from the topological version of the BRS theorem to the way of how partitioning is reflected in protocol complexes. More specifically, we show that set agreement can always be solved in ASM systems with *partition-compatible runs*, which are runs that allow some processes to hide their information from others. Our key insight is that partition-compatible runs “pierce” a hole at the center of the protocol complex, which allows the border to be retracted continuously to the center, thereby allowing to solve set agreement.

One implication of this result is that partitioning arguments cannot be used to prove the set-agreement impossibility in general: Assume that there is such a proof, which necessarily relies on some set of partition-compatible runs. After all, any partitioning argument rules out runs where every process has a complete view of all other processes in all iterations. Since we can construct a correct set agreement protocol for this set of runs, however, we have a contradiction.

The central idea of partitioning arguments, which exploit limited communication between sets of processes, stimulated the notion of partition-compatible views and (sets of) runs:

Definition 8 Partition compatibility. *A view $(p_i, s, v_1, \dots, v_n)$ of a process p_i at the end of some round $k \geq 1$ is called partition-compatible, if p_i did not get information from some p_j during round k , i.e., when $v_j = \perp$. A set of runs S is partition-compatible if, for every run $\alpha \in S$, there exists some participating process p_i and a round k at the end of which p_i 's view is partition-compatible.*

Note that it is the presence of the complete-view run, i.e., the presence of the complete-view simplex in the corresponding protocol complex, that makes a set of runs *not* partition compatible. This, once made, obvious observation gave us the idea for the following simple set agreement protocol:

We define the following 1-layer protocol P for the immediate snapshot model with a set of processes $\Pi = \{p_0, \dots, p_n\}$.

Definition 9 1-layer protocol. *Let $p_i \in \Pi = \{p_0, \dots, p_n\}$ be a process and (m_0, \dots, m_n) its view. Since we are considering the iterated immediate snapshot model, the protocol is determined just by the number of communication rounds (1 in this case) and the decision map. We define*

$$\mu(p_i, m_0, \dots, m_n) = \begin{cases} m_i & \text{if } \perp = m_j \text{ for some } j \in \{1, \dots, n\}, \\ m_{(i+1) \bmod n+1} & \text{otherwise.} \end{cases}$$

Obviously, since μ always chooses the input value for p_i unless all other input values have been observed, μ satisfies the validity condition. In fact, we can prove the following result:

Lemma 6 Correctness of the 1-round protocol. *Let S be a set of partition compatible runs for a 1-round immediate snapshot protocol. Then, μ solves set agreement in S .*

Proof. Let $\alpha \in S$ be a partition compatible run for a 1-round immediate snapshot protocol. Note that if the set of participating processes of α has cardinality less than $n + 1$, then the decision map μ trivially solves n -set agreement. For the sake of readability, we will assume for the rest of the proof that any increase or decrease of a process index is considered to be modulo $n + 1$.

Assume that α has $n + 1$ participating processes. Let p_j be one of the processes that executed their round in the last step of the protocol. Since α has $n + 1$ participating processes, p_j observed all of the input values. From the definition of μ , then p_j decides the input value for p_{j+1} . Assume for a contradiction that μ does not solve n -set agreement for α . Since p_j did not decide on its input value, and μ does not solve n -set agreement, this implies that p_{j-1} decides the input value for p_j . Since p_{j-1} decided the input value for p_j , then p_{j-1} also was among the processes that executed their round in the last step. It follows from a simple inductive argument that all of Π executed their round in their last step.

Therefore, α is an execution where all processes of Π executed concurrently. It follows from the immediate snapshot model specification that all processes hear from each other in α . This contradicts the fact that α was a partition compatible run, however.

An immediate consequence of Lemma 6 is that a partitioning argument cannot be used for showing n -set agreement impossibility for 1-round immediate snapshot protocols. However, in order to show that a partitioning argument cannot be used for a general shared memory protocol, we need to show this result for any number of layers. In order to do so, we define a k -round set agreement protocol for any value of k :

Definition 10 k -layer protocol. Consider a general k -layered immediate snapshot protocol. Let $p_i \in \Pi$ be a process, and (m_0, \dots, m_n) its final view. We denote $L^\ell(m_0, \dots, m_n)$ as the view for p_i at layer ℓ . Alternatively, if α is a run of an r -layered protocol and $s < r$ then L_α^s denotes the s -layered protocol run induced by α . We say that a view $v = (m_0, \dots, m_n)$ of a process at a layer s is incomplete if either $\perp \in v$ or if there exists $\ell < s$ and $0 \leq r \leq n$ such that $L^\ell(m_r)$ is an incomplete view; recall that m_r is p_r 's view in round $s - 1$ or \perp . Let

$$\begin{aligned} \mu^1 &= \mu, \\ \mu^{k+1}(p_i, m_0, \dots, m_n) &= \begin{cases} \mu^k(p_i, L^k(m_0, \dots, m_n)) & \text{if } (m_0, \dots, m_n) \text{ is incomplete,} \\ \mu^k(p_{i+1}, L^k(m_0, \dots, m_n)) & \text{otherwise.} \end{cases} \end{aligned} \tag{1}$$

We can prove the following result:

Lemma 7 Correctness of the k -round protocol. Let S be a set of partition compatible runs for a k -round immediate snapshot protocol. Then, the decision map μ^k solves n -set agreement in S .

Proof. We will proceed by induction over the number of rounds. The base of the induction is already given by Lemma 6. For the induction step, we assume that μ^k solves n -set agreement for a general k -layered immediate snapshot protocol. Now consider a general $k + 1$ -layered protocol, and let α be a set of partition compatible runs for this protocol. Assume that α has a full participating set Π (otherwise n -set agreement is trivial). Note that L_α^k is either partition compatible, or it is the run where every process heard from each other during all rounds up to k . We distinguish two cases:

(i) Assume that L_α^k is partition compatible. Let (p_i, m_0, \dots, m_n) be a process and its view in the $k + 1$ layer protocol. If $\perp \in v = (m_0, \dots, m_n)$, then v is an incomplete view, and $\mu^{k+1}(p_i, v) = \mu^k(p_i, v)$. If $\perp \notin v$, since L_α^k is partition compatible, there exists a pair of processes p_s and p_t and a round $u \leq k$ such that p_s does not hear from p_t during round u . Note that since $\perp \notin v$, then m_s corresponds to the view of process p_s until layer k . From the previous assumptions, it follows that $\perp \in L^u(m_s)$. Therefore m_s is an incomplete view and thus v . It follows that $\mu^{k+1}(p_i, v) = \mu^k(p_i, v)$. By the induction hypothesis, we have that

μ^k solves n -set agreement, therefore if L_α^k is partition compatible, then μ^{k+1} solves n -set agreement.

(ii) Assume that L_α^k is not partition compatible. Therefore, any view for any process in α is complete up until layer k , thus, for any pair of process views v, v' we have that $L^k(v) = L^k(v')$. Consider a process p_s that executed its $k + 1$ round last. Since α has a full participating set, then p_s has a complete view v' . From the definition, then $\mu^{k+1}(p_s, v') = \mu^k(p_{s+1}, L^k(v'))$. Since α is partition compatible, then there exists a minimal t such that p_{s+t} has an incomplete view. From the definition, it follows that $\mu^{k+1}(p_{s+t-1}, v') = \mu^k(p_{s+t}, L^k(v')) = \mu^{k+1}(p_{s+t}, L^k(v''))$ where v'' is the view from process p_{s+t} . It follows from a simple pigeonhole principle that μ^{k+1} solves n -set agreement in this case.

This completes the induction and the proof of Lemma 7.

We can thus state the main result of this section:

Theorem 4 . *Let S be any set of partition compatible runs in the iterated immediate snapshot model. Then, there exists a protocol P that solves set agreement for any run in S .*

Proof. Let k be the number of rounds executed in S . From Lemma 7, it follows that the generalized k -layered full information protocol with decision map μ^k solves set agreement.

6 Conclusions

We developed a topological version of a generalization of the BRS theorem. Our findings reveal that partitioning is reflected by a “color splitting” of the algorithm’s decision map, which separates the sub-complexes representing the partitioned processes. We used these insights to show that the impossibility of wait-free set agreement in the layered immediate snapshot model cannot be proved using partitioning arguments: For any set of partition compatible runs it is possible to construct a simple protocol that solves set agreement.

Extending the latter to k -set agreement and investigating the applicability of the BRS theorem to alternative shared memory systems remain as open questions.

References

1. Afek, Y., Gafni, E., Rajsbaum, S., Raynal, M., Travers, C.: The k -simultaneous consensus problem. *Distrib. Comput.* **22**(3), 185–195 (Mar 2010). <https://doi.org/10.1007/s00446-009-0090-8>, <https://doi.org/10.1007/s00446-009-0090-8>
2. Alistarh, D., Aspnes, J., Ellen, F., Gelashvili, R., Zhu, L.: Why extension-based proofs fail. *CoRR* **abs/1811.01421** (2018), <http://arxiv.org/abs/1811.01421>

3. Alistarh, D., Aspnes, J., Ellen, F., Gelashvili, R., Zhu, L.: Why extension-based proofs fail. In: Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23–26, 2019. pp. 986–996 (2019). <https://doi.org/10.1145/3313276.3316407>, <https://doi.org/10.1145/3313276.3316407>
4. Alistarh, D., Gilbert, S., Guerraoui, R., Travers, C.: Brief announcement: New bounds for partially synchronous set agreement. In: DISC. pp. 404–405 (2010)
5. Attiya, H., Bar-Noy, A., Dolev, D.: Sharing memory robustly in message-passing systems. *J. ACM* **42**(1), 124–142 (Jan 1995). <https://doi.org/10.1145/200836.200869>, <http://doi.acm.org/10.1145/200836.200869>
6. Attiya, H., Castañeda, A.: A non-topological proof for the impossibility of k -set agreement. *Theor. Comput. Sci.* **512**, 41–48 (2013)
7. Attiya, H., Paz, A.: Counting-based impossibility proofs for renaming and set agreement. In: Aguilera, M.K. (ed.) Proceedings, 26th International Symposium on Distributed Computing (DISC’12). pp. 356–370. Springer Berlin Heidelberg, Berlin, Heidelberg (2012)
8. Biely, M., Robinson, P., Schmid, U.: Weak synchrony models and failure detectors for message passing k -set agreement. In: Proceedings of the International Conference on Principles of Distributed Systems (OPODIS’09). pp. 285–299. LNCS, Springer Verlag, Nimes, France (Dec 2009), <http://www.vmars.tuwien.ac.at/documents/extern/2657/paper.pdf>
9. Biely, M., Robinson, P., Schmid, U.: Easy impossibility proofs for k -set agreement in message passing systems. In: Proceedings 15th International Conference on Principles of Distributed Systems (OPODIS’11), Springer LNCS 7109. pp. 299–312 (2011)
10. Bonnet, F., Raynal, M.: On the road to the weakest failure detector for k -set agreement in message-passing systems. *Theoretical Computer Science* **412**(33), 4273 – 4284 (2011). <https://doi.org/DOI:10.1016/j.tcs.2010.11.007>
11. Borowsky, E., Gafni, E.: Generalized FLP impossibility result for t -resilient asynchronous computations. In: STOC ’93: Proceedings of the twenty-fifth annual ACM symposium on Theory of computing. pp. 91–100. ACM, New York, NY, USA (1993). <https://doi.org/http://doi.acm.org/10.1145/167088.167119>
12. Bouzid, Z., Travers, C.: (anti- $\Omega^x \times \Sigma_z$)-based k -set agreement algorithms. In: Proceedings of the 14th international conference on Principles of distributed systems. pp. 189–204. OPODIS’10, Springer-Verlag, Berlin, Heidelberg (2010), <http://dl.acm.org/citation.cfm?id=1940234.1940256>
13. Brewer, E.A.: Towards robust distributed systems (abstract). In: Proceedings of the Nineteenth Annual ACM Symposium on Principles of Distributed Computing. PODC ’00, ACM, New York, NY, USA (2000). <https://doi.org/10.1145/343477.343502>
14. Fich, F., Ruppert, E.: Hundreds of impossibility results for distributed computing. *Distributed Computing* **16**, 121–163 (2003). <https://doi.org/10.1007/s00446-003-0091-y>
15. Fischer, M.J., Lynch, N.A., Paterson, M.S.: Impossibility of distributed consensus with one faulty process. *Journal of the ACM* **32**(2), 374–382 (Apr 1985)
16. Gilbert, S., Lynch, N.: Brewer’s conjecture and the feasibility of consistent, available, partition-tolerant web services. *SIGACT News* **33**(2), 51–59 (Jun 2002). <https://doi.org/10.1145/564585.564601>
17. Herlihy, M., Kozlov, D.N., Rajsbaum, S.: Distributed Computing Through Combinatorial Topology. Morgan Kaufmann (2013), <https://store.elsevier.com/product.jsp?isbn=9780124045781>

18. Herlihy, M., Rajsbaum, S., Raynal, M., Stainer, J.: From wait-free to arbitrary concurrent solo executions in colorless distributed computing. *Theoretical Computer Science* **683**, 1 – 21 (2017). <https://doi.org/https://doi.org/10.1016/j.tcs.2017.04.007>, <http://www.sciencedirect.com/science/article/pii/S0304397517303298>
19. Herlihy, M., Rajsbaum, S., Tuttle, M.R.: An overview of synchronous message-passing and topology. *Electronic Notes in Theoretical Computer Science* **39**(2), 1 – 17 (2000). [https://doi.org/https://doi.org/10.1016/S1571-0661\(05\)01148-5](https://doi.org/https://doi.org/10.1016/S1571-0661(05)01148-5), <http://www.sciencedirect.com/science/article/pii/S1571066105011485>
20. Herlihy, M., Shavit, N.: The topological structure of asynchronous computability. *J. ACM* **46**(6), 858–923 (Nov 1999). <https://doi.org/10.1145/331524.331529>, <http://doi.acm.org/10.1145/331524.331529>
21. de Prisco, R., Malkhi, D., Reiter, M.: On k-set consensus problems in asynchronous systems. *IEEE Trans. Parallel Distrib. Syst.* **12**(1), 7–21 (Jan 2001). <https://doi.org/10.1109/71.899936>, <http://dx.doi.org/10.1109/71.899936>
22. Saks, M., Zaharoglou, F.: Wait-free k-set agreement is impossible: The topology of public knowledge. *SIAM J. Comput.* **29**(5), 1449–1483 (2000). <https://doi.org/http://dx.doi.org/10.1137/S0097539796307698>