# Topological Characterization of Consensus Solvability in Directed Dynamic Networks*

Hugo Rincon Galeana[1], Ulrich Schmid[1], Kyrill Winkler[2], Ami Paz[3], and Stefan Schmid[4]

[1] TU Wien, Austria
[2] ITK Engineering, Austria
[3] CNRS, France
[4] TU Berlin, Germany

**Abstract.** Consensus is one of the most fundamental problems in distributed computing. This paper studies the consensus problem in a synchronous dynamic directed network, in which communication is controlled by an oblivious message adversary. The question when consensus is possible in this model has already been studied thoroughly in the literature from a combinatorial perspective, and is known to be challenging. This paper presents a topological perspective on consensus solvability under oblivious message adversaries, which provides interesting new insights.

Our main contribution is a topological characterization of consensus solvability, which also leads to explicit decision procedures. Our approach is based on the novel notion of a communication pseudosphere, which can be seen as the message-passing analog of the well-known standard chromatic subdivision for wait-free shared memory systems. We further push the elegance and expressiveness of the "geometric" reasoning enabled by the topological approach by dealing with uninterpreted complexes, which considerably reduce the size of the protocol complex, and by labeling facets with information flow arrows, which give an intuitive meaning to the implicit epistemic status of the faces in a protocol complex.

**Keywords:** Dynamic networks · message adversary · consensus · combinatorial topology · uninterpreted complexes

## 1 Introduction

Consensus is a most fundamental problem in distributed computing, in which multiple processes need to agree on some value, based on their local inputs. The problem has already been studied for several decades and in various different models, yet in many distributed settings the question of when and how fast consensus can be achieved continues to puzzle researchers.

This paper studies consensus in the fundamental setting where processes communicate over a synchronous dynamic directed network, where communication is controlled by an oblivious *message adversary* [2]. This model is appealing, because it is conceptually simple and still provides a highly dynamic network model. In this model, fault-free processes communicate in a lock-step synchronous fashion using message passing, and a message adversary may drop some messages sent by the processes in each round. Viewed more abstractly, the message adversary provides a sequence of directed communication graphs, whose edges indicate which process can successfully send a message to which other process in that round. An oblivious message adversary is defined by a set $\mathbf{D}$ of allowed communication graphs, from which it can pick one in each round [12], independently of its picks in the previous rounds.

The model is practically motivated, as the communication topology of many large-scale distributed systems is *dynamic* (e.g., due to mobility, interference, or failures) and its links are often *asymmetric* (e.g., in optical or in wireless networks) [31]. The model is also theoretically interesting, as solving consensus in general dynamic directed networks is known to be difficult [6, 12, 34, 36, 42].

Prior work primarily focused on the circumstances under which consensus is actually solvable under oblivious message adversaries [12]. Only recently, first insights have been obtained on the time

* This work is supported by the Austrian Science Fund (FWF) projects ADynNet (P28182) and ByzDEL (P33600), and the German Research Foundation (DFG) project FlexNets (470029389). Kyrill Winkler was supported by the Vienna Science and Technology Fund (WWTF) project WHATIF (ICT19-045) and Ami Paz by the Austrian Science Fund (FWF) and netIDEE SCIENCE project P 33775-N.

complexity of reaching consensus in this model [39], using a combinatorial approach. The present paper complements this by a topological perspective, which provides interesting new insights and results.

**Our contributions:** Our main contribution is a topological characterization of consensus solvability for synchronous dynamic networks under oblivious message adversaries. It provides not only intuitive ("geometric") explanations for the surprisingly intricate time complexity results established in [39], both for the decision procedure (which allows to determine whether consensus is solvable for a given oblivious message adversary or not) and, in particular, for the termination time of any correct distributed consensus algorithm.

To this end, we introduce the novel notion of a communication pseudosphere, which can be seen as the message-passing analog of the well-known standard chromatic subdivision for wait-free shared memory systems. Moreover, we use uninterpreted complexes [38], which considerably reduce the size and structure of our protocol complexes. And last but not least, following [19], we label the edges in our protocol complexes by the information flow that they carry, which give a very intuitive meaning to the the implicit epistemic status (regarding knowledge of initial values) of the vertices/faces in a protocol complex. Together with the inherent beauty and expressiveness of the topological approach, our tools facilitate an almost "geometric" reasoning, which provides simple and intuitive explanations for the surprising results of [39], like the sometimes exponential gap between decision complexity and consensus termination time. It also leads to a novel decision procedure for deciding whether consensus under a given oblivious message adversary can be achieved in some $k$ rounds.

In general, we believe that, unlike the combinatorial approaches considered in the literature so far, our topological approach also has the potential for the almost immediate generalization to other decision problems and other message adversaries, and may hence be of independent interest.

**Related work**: Consensus problems arise in various models, including shared memory architectures, message-passing systems, and blockchains, among others [1, 28, 33, 40]. The distributed consensus problem in the message-passing model, as it is considered in this paper, where communication occurs over a dynamic network, has been studied for almost 40 years [7, 10–12, 16, 22, 35, 36]. Already in 1989, Santoro and Widmayer [34] showed that consensus is impossible in this model if up to $n-1$ messages may be lost each round. Schmid, Weiss and Keidar [36] showed that if losses do not isolate the processes, consensus can even be solved when a quadratic number of messages is lost per round. Several other generalized models have been proposed in the literature [11, 17, 24], like the heard-of model by Charron-Bost and Schiper [11], and also different agreement problems like approximate and asymptotic consensus have been studied in these models [10, 16]. In many of these and similar works on consensus [5, 6, 9, 14, 32, 37, 42], a model is considered in which, in each round, a digraph is picked from a set of possible communication graphs. Afek and Gafni coined the term message adversary for this abstraction [2], and used it for relating problems solvable in wait-free read-write shared memory systems to those solvable in message-passing systems. For a detailed overview of the field, we refer to the recent survey by Winkler and Schmid [40].

An interesting alternative model for dynamic networks assumes a $T$-interval connectivity guarantee, that is, a common subgraph in the communication graphs of every $T$ consecutive rounds [29, 30]. In contrast to our directional model, solving consensus is relatively simple here, since the $T$-interval connectivity model relies on bidirectional links and always connected communication graphs. For example, 1-interval-connectivity, the weakest form of $T$-interval connectivity, implies that all nodes are able to reach all the other nodes in the system in each of the graphs. Solving consensus in undirected graphs that are always connected was also considered in the case of a given $(t+1)$-connected graph and at most $t$-node failures [8]. Using graph theoretical tools, the authors extend the notion of a radius in a graph to determine the consensus termination time in the presence of failures.

Coulouma, Godard, and Peters [12] showed an interesting equivalence relation, which captures the essence of consensus impossibility under oblivious message adversaries via the non-broadcastability of one of the so-called beta equivalence classes, hence refining the results of [35]. Building upon some of these insights, Winkler et al. [39] studied of the time complexity of consensus in this model. In particular, they presented an explicit decision procedure and analyzed both its decision time complexity and the termination time of distributed consensus. It not only turned out that consensus may take exponentially longer than broadcasting [13], but also that there is sometimes an exponential gap between decision time and termination time. Surprisingly, this gap is not caused by properties related to broadcastability of the beta classes, but rather by the number of those.

Whereas all the work discussed so far is combinatorial in nature, there is also some related topological research, see [20] for an introduction and overview. Using topology in distributed computing started out from wait-free computation in shared memory systems with immediate atomic snapshots (the IIS model), see e.g. [3, 4, 18, 23, 26, 27]. The evolution of the protocol complex in the IIS model is governed by the pivotal chromatic subdivision operation here. We will show that the latter can alternatively be viewed as a specific oblivious message adversary, the set $\mathbf{D}$ of which containins all transitively closed and unilaterally connected graphs.

Regarding topology in dynamic networks, Castañeda et al. [9] studied consensus and other problems in both static and dynamic graphs, albeit under the assumption that all the nodes know the graph sequence. That is, they focused on the question of information dissemination, and put aside questions of indistinguishability between graph sequences. In contrast, in our paper, we develop a topological model that captures both information dissemination and indistinguishability. An adversarial model that falls into "our" class of models has been considered by Godard and Perdereau [19], who studied general $k$-set agreement under the assumption that some maximum number of (bidirectional) links could drop messages in a round. The authors also introduced the idea to label edges in the protocol complex by arrows that give the direction of the information flow, which we adopted. Shimi and Castañeda [38] studied $k$-set agreement under the restricted class of oblivious message adversaries that are "closed-above" (with $\mathbf{D}$ containing, for every included graph, also all graphs with more edges).

One of the challenges of applying topological tools in distributed settings is that the simplicial complex representing the system grows dramatically with the number of rounds, as well as with the number of processes and possible input values. In the case of colorless tasks, such as $k$-set agreement, the attention can be restricted to colorless protocol complexes [20]. In the case of the IIS model, its evolution is governed by the barycentric subdivision, which results in much smaller protocol complexes than produced by the chromatic subdivision. Unfortunately, however, it is not suitable for tracing indistinguishability in dynamic networks under message adversaries. The same is true for the "local protocol complexes" introduced in [15]. By contrast, uninterpreted complexes, as introduced in [38], are effective here and are hence also used in our paper.

Apart from consensus being a special case of $k$-set agreement (for $k = 1$), consensus has not been the primary problem of interest for topology in distributed computing, in particular not for dynamic networks under message adversaries. However, a point-set topological characterization of when consensus is possible under general (both closed and non-closed) message adversaries has been presented by Nowak, Schmid and Winkler in [32]. The resulting decision procedure is quite abstact, though (it acts on infinite admissible executions), and so are some results on the termination time for closed message adversaries that confirm [41].

The topology of message-passing models in general has been considered by Herlihy, Rajsbaum, and Tuttle already in 2002 [22]. Herlihy and Rajsbaum [21] studied $k$-set agreement in models leading to shellable complexes.

**Paper organization:** We introduce our model of distributed computation and the oblivious message adversary in Section 2. In Section 3 we present a framework which will allow us to study consensus on dynamic networks from a topological perspective. Our characterization of consensus solvability/impossibility for the oblivious message adversary is presented in Section 4, where we also describe an explicit decision procedure. In Section 5 we further explore the relationship between the time complexity required by our decision procedure and the actual termination time of distributed consensus. We conclude our contribution and discuss future research directions in Section 6.

## 2 System Model

We consider a synchronous dynamic network consisting of a set of $n$ processes that do not fail, which are fully-connected via point-to-point links that might drop messages. We identify the processes solely by their unique ids, which are taken from the set $\Pi = \{p_1, \ldots, p_n\}$ and known to the processes. Let $[n] = \{1, \ldots, n\}$. Processes execute a deterministic full-information protocol $P$, using broadcast (send-to-all) communication. Their execution proceeds in a sequence of lock-step rounds, where every process simultaneously broadcasts a message to every other process, without getting immediately informed of a successful message reception, and then computes its next state based on its current local state and the messages received in the round. The rounds are communication-closed, i.e., messages not received in a specific round are lost and will not be delivered later.

Communication is hence unreliable, and in fact controlled by an oblivious *message adversary* (MA) with non-empty graph set $\mathbf{D} = \{D_1, \ldots, D_k\}$. All the graphs have $\Pi$ as their set of nodes, and an edge $p_i \to p_j$ represents a communication link from $p_i$ to $p_j$. For every round $r \geq 1$, the MA arbitrarily picks some communication graph $G_r$ from $\mathbf{D}$, and a message from a process $p_i$ arrives to process $p_j$ in this round if $G_r$ contains the edge $p_i \to p_j$, and otherwise it is lost. We assume processes have persistent memory, i.e., every graph in $\mathbf{D}$ contains all self-loops $p_i \to p_i$. An infinite graph sequence $\mathcal{G} = (G_r)_{r\geq 1}$ picked by the message adversary is called a feasible graph sequence, and $\mathbf{D}^\omega$ denotes the set of all *feasible graph sequences* for the oblivious message adversary with graph set $\mathbf{D}$. The processes know $\mathbf{D}$, but they do not have a priori knowledge of the graph $G_r$ for any $r$ (though they may infer it after the round occurred).

We consider a system where the global state is fully determined by the local states of each process. Therefore, a *configuration* is just the vector of the local states (also called *views*) of the processes. An admissible *execution* $a$ of $P$ is just the sequence of configurations $a = (a_r)_{r\geq 0}$ at the end of the rounds $r \geq 1$, induced by a feasible graph sequence $\mathcal{G} \in \mathbf{D}^\omega$ starting out from a given initial configuration $a_0$. Since we will restrict our attention to deterministic protocols $P$, the graph sequence $\mathcal{G}$ and the initial configuration $a_0$ uniquely determine $a$. The *view* of process $p_i$ in $a_r$ at the end of round $r \geq 1$ is denoted as $a_r(p_i)$; its initial view is denoted as $a_0(p_i)$.

We restrict our attention to deterministic protocols for the consensus problem, defined as follows:

**Definition 1 (Consensus).** *Every process $p_i \in \Pi$ has an input value $x_i \in \mathcal{V}_I$ taken from a finite input domain $\mathcal{V}_I$, which is encoded in the initial state, and an output value $y_i \in \mathcal{V}_I \cup \{\perp\}$, initially $y_i = \perp$. In every admissible execution, a correct consensus protocol $P$ must ensure the following properties:*

- ***Termination:*** *Eventually, every $p_i \in \Pi$ must decide, i.e., change to $y_i \neq \perp$, exactly once.*
- ***Agreement:*** *If processes $p_i$ and $p_j$ have decided, then $y_i = y_j$.*
- ***(Strong) Validity:*** *If $y_i \neq \perp$, then $y_i = x_j$ for some $p_j \in \Pi$, i.e., must be the input value of some process $p_j$.*

In any given admissible execution $a$ of $P$, induced by $\mathcal{G} \in \mathbf{D}^\omega$, for a process $p_i$, let $\mathrm{In}^{\mathcal{G}}(p_i, r)$ be the set of processes $p_i$ has *heard of* in round $r$ (see also [11]), i.e., the set of in-neighbors of process $p_i$ in $G_r$, and $\mathrm{In}^{\mathcal{G}}(p_i, 0) = \{p_i\}$. Since all graphs in $\mathbf{D}$ contain all self-loops, we have that $p_i \in \mathrm{In}^{\mathcal{G}}(p_i, r)$ for all $r \geq 0$ and $p_i \in \Pi$. If the round $r$ is clear from the context, we also abbreviate $\mathrm{In}^{\mathcal{G}}(p_i) = \mathrm{In}^{\mathcal{G}}(p_i, r)$.

The evolution of the local views of the processes in an admissible execution $a$, induced by $\mathcal{G}$ and the initial configuration $a_0$, can now be defined recursively as

$$a_r(p_i) = \big\{(p_j, r, a_{r-1}(p_j)) : p_j \in \mathrm{In}^{\mathcal{G}}(p_i, r)\big\} \quad \text{for } r > 0. \tag{1}$$

Note that we could drop the round number $r$ from $(p_j, r, a_{r-1}(p_j))$ in the above definition, since it is implicitly contained in the structure of $a_r(p_i)$; we included it explicitly for clarity only. The set of all possible round-$r$ views of $p_i$, including all the initial views for $r = 0$, in any admissible execution, is denoted by $A^r(p_i) = \{a_r(p_i) \mid \forall \text{ admissible executions } a \text{ under MA}\}$.

In any admissible execution $a$, every process must eventually reach a final view, where it can take a decision on an output value which will not be changed later. Consequently, there is some final round after which all processes have decided.

## 3 A Topological Framework for Consensus

In this section, we introduce the basic elements of combinatorial topology and specific concepts needed in our context of synchronous message-passing networks.

Combinatorial topology in distributed computing [20] rests on simplicial input and output complexes describing the feasible input and output values of a distributed decision task like consensus, and a carrier map that defines the allowed output value(s), i.e., output simplices, for a given input simplex. A protocol that solves such a task in some computational model gives rise to another simplicial complex, the protocol complex, which describes the evolution of the local views of the processes in any execution. Protocol complexes traditionally model full information protocols in round-based models, which ensures a well-organized structure: The processes execute a sequence of communication operations, which disseminate their complete views, until they are able to make a decision. Finally, a protocol induces a simplicial decision map, which maps each vertex in the protocol complex to an output vertex in a way compatible with the carrier map.

### 3.1 Basic topological definitions

We start with the definitions of the basic vocabulary of combinatorial topology:

**Definition 2 (Abstract simplicial complex).** *An* abstract simplicial complex $\mathcal{K}$ *is a pair* $\langle V(\mathcal{K}), F(\mathcal{K})\rangle$*, where* $V(\mathcal{K})$ *is a set,* $F(\mathcal{K}) \subseteq 2^{V(\mathcal{K})}$*, and for any* $\sigma, \tau \in 2^{V(\mathcal{K})}$ *such that* $\sigma \subseteq \tau$ *and* $\tau \in F(\mathcal{K})$*, then* $\sigma \in F(\mathcal{K})$*.* $V(\mathcal{K})$ *is called the set of* vertices*, and* $F(\mathcal{K})$ *is the set of* faces *or* simplices *of* $\mathcal{K}$*. We say that a simplex* $\sigma$ *is a* facet *if it is maximal with respect to containment, and a* proper face *otherwise. We use* $Fct(\mathcal{K})$ *to denote the set of all facets of* $\mathcal{K}$*, and note that for a given* $V(\mathcal{K})$ *we have that* $F(\mathcal{K})$ *uniquely define* $Fct(\mathcal{K})$ *and vice versa. A simplicial complex is* finite *if its vertex set is finite, which will be the case for all the complexes in this paper.*

All the simplicial complexes we consider in this work are abstract. For conciseness, we will usually sloppily write $\sigma \in \mathcal{K}$ instead of $\sigma \in F(\mathcal{K})$.

**Definition 3 (Subcomplex).** *Let* $\mathcal{K}$ *and* $\mathcal{L}$ *be simplicial complexes. We say that* $\mathcal{L}$ *is a* subcomplex *of* $\mathcal{K}$*, written as* $\mathcal{L} \subseteq \mathcal{K}$*, if* $V(\mathcal{L}) \subseteq V(\mathcal{K})$ *and* $F(\mathcal{L}) \subseteq F(\mathcal{K})$*.*

**Definition 4 (Dimension).** *Let* $\mathcal{K}$ *be a simplicial complex, and* $\sigma \in F(\mathcal{K})$ *be a simplex. We say that* $\sigma$ *has* dimension $k$*, denoted by* $\dim(\sigma) = k$*, if it has a cardinality of* $k+1$*. A simplicial complex* $\mathcal{K}$ *is of dimension* $k$ *if every facet has dimension at most* $k$*, and it is* pure *if all its facets have the same dimension.*

We sometimes denote a simplex as $\sigma^k$ in order to stress that its dimension is $k$.

**Definition 5 (Skeletons and boundary complex).** *The* $k$-skeleton $\mathrm{skel}^k(\mathcal{K})$ *of a simplicial complex* $\mathcal{K}$ *is the subcomplex consisting of all simplices of dimension at most* $k$*. The* boundary complex $\partial\sigma$ *of a simplex* $\sigma$*, viewed as a complex, is the complex made up of all proper faces of* $\sigma$*.*

**Definition 6 (Simplicial maps).** *Let* $\mathcal{K}$ *and* $\mathcal{L}$ *be simplicial complexes. We say that a vertex map* $\mu : V(\mathcal{K}) \to V(\mathcal{L})$ *is a* simplicial map *if, for any* $\sigma \in F(\mathcal{K})$*,* $\mu(\sigma) \in F(\mathcal{L})$*; here,* $\mu(\sigma) = \{\mu(v) \mid v \in \sigma\}$*.*

**Definition 7 (Colorings and chromatic simplicial complexes).** *We say that a simplicial complex* $\mathcal{K}$ *has a* proper $c$-coloring $\chi$*, if there exists* $\chi : V(\mathcal{K}) \to \{p_1, p_2, \dots, p_c\}$ *that is injective at every face of* $\mathcal{K}$*. If* $\mathcal{K}$ *has a proper* $(\dim(\mathcal{K})+1)$*-coloring, we say it is a* chromatic *simplicial complex.*

The range of $\chi$ is extended to sets of vertices $S$ by defining $\chi(S) = \{\chi(v) \mid v \in S\}$, which implies e.g. $\chi(\sigma) = \chi(V(\sigma))$.

**Definition 8 (Carrier Map).** *Let* $\mathcal{K}$ *and* $\mathcal{L}$ *be simplicial complexes and* $\Phi : F(\mathcal{K}) \to 2^{\mathcal{L}}$*. We say that* $\Phi$ *is a* carrier map*, if* $\Phi(\sigma)$ *is a subcomplex of* $\mathcal{L}$ *for any* $\sigma \in \mathcal{K}$*, and for any* $\sigma_1, \sigma_2 \in \mathcal{L}$*,* $\Phi(\sigma_1 \cap \sigma_2) \subseteq \Phi(\sigma_1) \cap \Phi(\sigma_2)$*.*

*We say that a carrier map is* rigid *if it maps every simplex* $\sigma \in \mathcal{K}$ *to a complex* $\Phi(\sigma)$ *which is pure of dimension* $\dim(\sigma)$*. It is said to be* strict *if that for any two simplices* $\sigma, \tau \in \mathcal{K}$*,* $\Phi(\sigma \cap \tau) = \Phi(\sigma) \cap \Phi(\tau)$*.*

*We say that a carrier map* $\Phi : \mathcal{K} \to 2^{\mathcal{L}}$ carries *a simplicial vertex map* $\mu : V(\mathcal{K}) \to V(\mathcal{L})$ *if for any* $\sigma \in \mathcal{K}$*,* $\mu(\sigma) \in \Phi(\sigma)$*.*

Having introduced our basic vocabulary, we can now define the main ingredients for the topological modeling of consensus in our setting.

Generally, a *distributed task* is defined by a tuple $T = \langle \mathcal{I}, \mathcal{O}, \Delta \rangle$ consisting of chromatic simplicial complexes $\mathcal{I}$ and $\mathcal{O}$ that model the valid input and output configurations respectively, for the set $\Pi$ of processes, and $\Delta : \mathcal{I} \to 2^{\mathcal{O}}$ is a carrier map that maps valid input configurations to sets of valid output configurations. Both complexes have vertices of the form $(p_i, x)$ with $p_i \in \Pi$, and they are chromatic with the coloring function $\chi((p_i, x)) = p_i$. All the simplicial maps we consider in this work are *color preserving*, in the sense that they map each vertex $(p_i, x)$ to a vertex $(p_i, y)$ with the same process id $p_i$.

Many interesting tasks have some degree of regularity (that is, symmetry) in the input complex. In the case of consensus, in particular, any combination of input values from $\mathcal{V}_I$ is a legitimate initial configuration. Consequently, the input complex for consensus in the classic topological modeling is a pseudosphere [22].

In this paper, we will exploit the fact that strong validity does not force us to individually trace the evolution of every possible initial configuration of the protocol complex. We will therefore restrict our attention to *uninterpreted complexes* [38]: Instead of providing different vertices for every possible value of $x_i$, we provide only one vertex labeled with $\{p_i\}$, carrying the meaning of "the actual input value $x_i$ of $p_i$". This way, we can abstract away the input domain $\mathcal{V}_I$ as well as the actual assignment of initial values $x_i \in \mathcal{V}_I$ to the processes. Topologically, uninterpreted complexes thus correspond to a "flattening" of the standard complexes with respect to all input and output values. The main advantages of resorting to uninterpreted protocol complexes is that they are exponentially smaller than the standard protocol complex, even in the case of binary consensus, and independent of the particular initial configuration. This can be compared with the study of *colorless tasks* [20, Ch. 4], where a different form of "flattening" of the complexes is done by omitting the process ids.

**Definition 9 (Uninterpreted input complex for consensus).** *The* uninterpreted input complex $\mathcal{I}$ *for consensus is just a single* initial simplex $\sigma_0 = \{(p_1, \{p_1\}), \dots, (p_n, \{p_n\})\}$ *and all its faces, with the set of vertices* $V(\mathcal{I}) = V(\sigma_0) = \{(p_i, \{p_i\}) \mid p_i \in \Pi\}$*, where the label* $\{p_i\}$ *represents the "uninterpreted" (i.e., fixed but arbitrary) input value of* $p_i$.

We use $\sigma_0$ throughout this paper to denote the above input simplex.

The *uninterpreted output complex* $\mathcal{O}$ for consensus just specifies the process whose input value will determine the decision value.

**Definition 10 (Uninterpreted output complex for consensus).** *The* uninterpreted output complex $\mathcal{O}$ *for consensus is the union of* $n$ *disjoint complexes* $\mathcal{O}(p_j)$, $p_j \in \Pi$*, each consisting of the simplex* $\{(p_1, \{p_j\}), \dots, (p_n, \{p_j\})\}$ *and all its faces. The label* $\{p_j\}$ *represents the "uninterpreted" (i.e., fixed but arbitrary) input value of* $p_j$.

The carrier map $\Delta$ for the consensus task maps any face $\rho$ of the initial simplex $\sigma_0 \in \mathcal{I}$ to $\dim(\rho)$-faces of $\mathcal{O}$ that all have a coloring equal to $\chi(\rho)$. Clearly, $\Delta$ is rigid and strict.

The *uninterpreted protocol complex* $\mathcal{P}_r^{\mathbf{D}^\omega}$ consists of vertices that are labeled by the heard-of histories the corresponding process has been able to gather so far.

**Definition 11 (Heard-of histories).** *For a feasible graph sequence* $\mathcal{G}$*, the* heard-of history $h_r^{\mathcal{G}}(p_i)$ *of a process* $p_i$ *at the end of round* $r$ *is defined as*

$$h_r^{\mathcal{G}}(p_i) = \{(p_j, h_{r-1}^{\mathcal{G}}(p_j)) \mid p_j \in \mathrm{In}^{\mathcal{G}}(p_i, r)\} \text{ for } r \geq 1, \tag{2}$$

$$h_0^{\mathcal{G}}(p_i) = \{p_i\}. \tag{3}$$

*The global heard-of history* $h_r^{\mathcal{G}}$ *at the end of round* $r$ *is just the tuple* $(h_r^{\mathcal{G}}(p_1), \dots, h_r^{\mathcal{G}}(p_n))$.

*The set of processes* $p_i$ *has ever heard of up to* $h_r^{\mathcal{G}}(p_i)$*, i.e., the end of round* $r$*, is denoted* $\cup h_r^{\mathcal{G}}(p_i) = \bigcup_{p_j \in \mathrm{In}^{\mathcal{G}}(p_i, r)} \cup h_{r-1}^{\mathcal{G}}(p_j)$ *and* $\cup h_0^{\mathcal{G}}(p_i) = h_0^{\mathcal{G}}(p_i) = \{p_i\}$.

*The set of all possible heard-of histories of* $p_i$ *(resp. the global ones) at the end of round* $r \geq 0$*, in every feasible graph sequence* $\mathcal{G} \in \mathbf{D}^\omega$*, is denoted by*

$$H^r(p_i) = \{h_r^{\mathcal{G}}(p_i) \mid \mathcal{G} \in \mathbf{D}^\omega\}, \tag{4}$$

$$H^r = \{(h_r^{\mathcal{G}}(p_1), \dots, h_r^{\mathcal{G}}(p_n)) \mid \mathcal{G} \in \mathbf{D}^\omega\}. \tag{5}$$

The uninterpreted protocol complex $\mathcal{P}_r^{\mathbf{D}^\omega}$, which does not depend on the initial configuration but only on $\mathbf{D}^\omega$, is defined as follows:

**Definition 12 (Uninterpreted protocol complex for $\mathbf{D}^\omega$).** *The* uninterpreted $r$-round protocol complex $\mathcal{P}_r^{\mathbf{D}^\omega} = \langle V(\mathcal{P}_r^{\mathbf{D}^\omega}), F(\mathcal{P}_r^{\mathbf{D}^\omega})\rangle$*,* $r \geq 0$*, for a given oblivious message adversary* $\mathbf{D}^\omega$*, is defined by its vertices and facets as follows:*

$$V(\mathcal{P}_r^{\mathbf{D}^\omega}) = \big\{(p_i, h_r(p_i)) \mid p_i \in \Pi,\ h_r(p_i) \in H^r(p_i)\big\},$$
$$Fct(\mathcal{P}_r^{\mathbf{D}^\omega}) = \big\{\{(p_1, h_r(p_1)), \dots, (p_n, h_r(p_n))\} | \forall 1 \leq i \leq n : p_i \in \Pi, (h_r(p_1), \dots, h_n(p_n)) \in H^r\big\}.$$

For conciseness, we will often omit the superscript $\mathbf{D}^\omega$ when the oblivious message adversary considered is clear from the context.

The *decision map* $\mu : V(\mathcal{P}_r^{\mathbf{D}^\omega}) \to V(\mathcal{O})$ is a chromatic simplicial map that maps a final view of a process $p_i$ at the end of round $r$ to an output value $p_j$ such that $p_j \in \cup h_r^{\mathcal{G}}(p_i)$; it is not defined for non-final views. Note that $\mu$ is uniquely determined by the images of the facets in $\mathcal{P}_r^{\mathbf{D}^\omega}$ after any round $r$ were all processes have final views. We say that consensus is solvable if such a simplicial map $\mu$ exists.

*Remark.* Standard topological modeling, which does not utilize uninterpreted complexes, also requires an execution carrier map $\Xi : \mathcal{I} \to 2^{\mathcal{P}}$, which defines the subcomplex $\Xi(\sigma)$ of the protocol complex $\mathcal{P}$ that arises when the protocol starts from the initial simplex $\sigma \in \mathcal{I}$. Solving a task requires $\mu \circ \Xi$ to be carried by $\Delta$, i.e., $\mu(\Xi(\sigma)) \in \Delta(\sigma)$ for all $\sigma \in \mathcal{I}$. In our setting, since we have only one (uninterpreted) facet in our input complex $\sigma_0$ and a protocol complex that can be written as $\bigcup_{r\geq 1} \mathcal{P}_r^{\mathbf{D}^\omega} = \bigcup_{r\geq 1} \mathcal{P}^r(\sigma_0)$ (i.e., the union of all iterated protocol complex construction operators $\mathcal{P}^r$ given in Definition 14 below), both the execution carrier map $\Xi$ and the carrier map $\Delta$ are independent of the actual initial values and hence quite simple: The former is just $\Xi = \bigcup_{r\geq 1} \mathcal{P}^r$ (with every $\mathcal{P}^r$ viewed as a carrier map), the latter has been stated after Definition 10.

### 3.2 Communication pseudospheres

Rather than directly using Definition 12 for $\mathcal{P}_r$, we will now introduce an alternative definition based on communication pseudospheres. The latter can be seen as the the message-passing analogon of the well-known standard chromatic subdivision (see Definition 16) for wait-free shared memory systems. Topologically, it can be defined as follows:

**Definition 13 (Communication pseudosphere).** *Let* $\mathcal{K}$ *be an* $(n-1)$*-dimensional pure simplicial complex with a proper coloring* $\chi : V(\mathcal{K}) \to \{p_1, \dots, p_n\}$*. We define the* communication pseudosphere $\mathrm{Ps}(\mathcal{K})$ *through its vertex set and facets as follows:*

$$V(\mathrm{Ps}(\mathcal{K})) = \big\{(p_i, \sigma) \mid \sigma \in F(\mathcal{K}), p_i \in \chi(\sigma)\big\}, \tag{6}$$
$$Fct(\mathrm{Ps}(\mathcal{K})) = \big\{\{(p_1, \sigma_1), (p_2, \sigma_2), \dots (p_n, \sigma_n)\} \mid \forall 1 \leq i \leq n :\ \sigma_i \in F(\mathcal{K}), p_i \in \chi(\sigma_i)\big\}. \tag{7}$$

Given an $(n-1)$-dimensional simplex $\sigma^{n-1} = \{(p_1, h_1), \dots, (p_n, h_n)\} \in \mathcal{K}$, the communication pseudosphere $\mathrm{Ps}(\sigma^{n-1})$ contains a vertex $(p_i, \sigma)$ for every subset $\sigma \subseteq \{(p_1, h_1), \dots, (p_n, h_n)\}$ that satisfies $\{(p_i, h_i)\} \in \sigma$. Intuitively, $\sigma$ represents the information of those processes $p_i$ could have heard of in a round (recall that $p_i$ always hears of itself). $\mathrm{Ps}(\sigma^{n-1})$ hence indeed matches the definition of a pseudosphere [22].

Since $\big|\{\sigma \subseteq \{(p_1, h_1), \dots, (p_n, h_n)\} \setminus \{(p_i, h_i)\}\}\big| = 2^{n-1}$ for every $p_i$, every communication pseudosphere $\mathrm{Ps}(\sigma^{n-1})$ consists of $|V(\mathrm{Ps}(\sigma^{n-1}))| = n2^{n-1}$ vertices: For every given vertex $(p_i, \sigma)$ and every $p_j \neq p_i$, there are exactly $2^{n-1}$ differently labeled vertices $(p_j, \cdot)$. Since $(p_i, \sigma)$ has an edge to each of those in the complex $\mathrm{Ps}(\sigma^{n-1})$, its degree must hence be $d = (n-1)2^{n-1}$.

In the case of $n = 2$ or $n = 3$, let $v = |V(\mathrm{Ps}(\sigma^{n-1}))|$, $e = |E(\mathrm{Ps}(\sigma^{n-1}))|$ and $f = |Fct(\mathrm{Ps}(\sigma^{n-1}))|$ denote the numbers of vertices, edges and facets in $\mathrm{Ps}(\sigma^{n-1})$, respectively. It obviously holds that $v \cdot d = 2e$ and $v \cdot d = nf$. Therefore, $e = vd/2 = n(n-1)2^{2(n-1)-1}$ and $f = vd/n = (n-1)2^{2(n-1)}$. For $n = 2$, we thus get $v = 4$, $f = e = 4$, $d = 2$ and hence the following communication pseudosphere $\mathrm{Ps}(\sigma_0^1)$ for the initial simplex $\sigma_0^1 = \{(p_r, \{p_r\}), (p_w, \{p_w\})\}$:

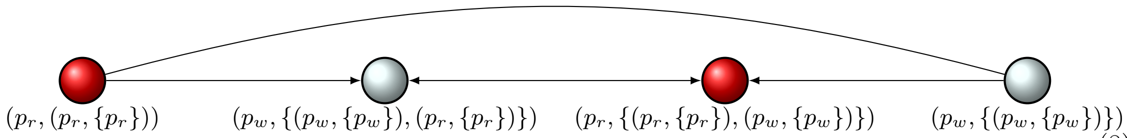

(8)

In the above figure, and throughout this paper, we use the labeling convention of the edges proposed in [19], which indicates the information flow between the vertices in a simplex. For example, in the middle simplex (connected with edge $\leftrightarrow$), both processes have heard from each other in round 1, so the connecting edge is denoted by $\leftrightarrow$. An edge without any arrow means that the two endpoints do not hear from each other. Note carefully that we will incorporporate these arrows also when talking about facets and faces that are *isomorphic*: Throughout this paper, two faces $\sigma$ and $\kappa$ arising in our protocol complexes will be considered isomorphic only if $\chi(\sigma) = \chi(\kappa)$ and if all edges have the same orientation.

We note also that the labeling of the vertices with the faces of $\sigma_0^1$ is highly redundant. We will hence condense vertex labels when we need to refer to them explicitly, and e.g. write $(p_r, \{p_r, p_w\})$ instead of $(p_r, \{(p_r, \{p_r\}), (p_w, \{p_w\})\})$.

The communication pseudosphere $\mathrm{Ps}(\sigma_0^2)$ for the initial simplex $\sigma_0^2 = \{(p_r, \{p_r\}), (p_g, \{p_g\}), (p_w, \{p_w\})\}$ for $n = 3$ is depicted in Fig. 1. It also highlights two facets, corresponding to the graphs $G_1$ (grey) and $G_2$ (yellow):

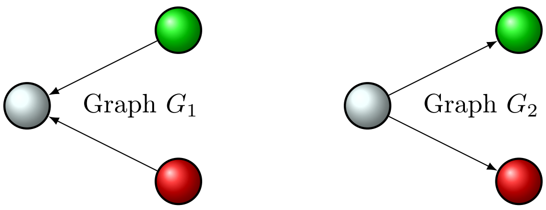

(9)

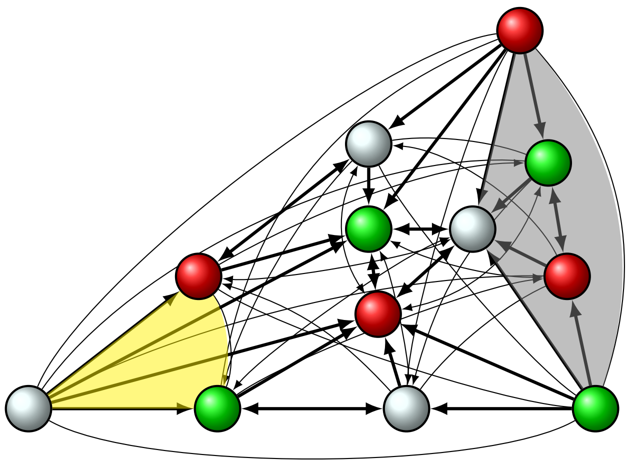

**Fig. 1.** Communication pseudosphere $\mathrm{Ps}(\sigma^{n-1})$ for $n = 3$ (where $L = 4$, $V = 12$, $E = 48$, $F = 32$ and $d = 8$), with the communication graphs of Eq. (9) highlighted. Thick edges represent the standard chromatic subdivision $\mathrm{Ch}(\sigma^2)$.

We will now recast the definition of the uninterpreted protocol complex for a given oblivious message adversary $\mathbf{D}^\omega$ in terms of a communication pseudosphere. Recall from Definition 12 that the uninterpreted initial protocol complex $\mathcal{P}_0 = \mathcal{P}_0^{\mathbf{D}^\omega}$ only consists of the single initial simplex $\sigma_0 = \sigma_0^{n-1} = \{(p_1, \{p_1\}), \ldots, (p_n, \{p_n\})\}$ and all its faces. It represents the uninterpreted initial state, where every process has heard only from itself. Here is an example for $n = 3$ and $\mathit{\Pi} = \{p_w, p_r, p_g\}$:

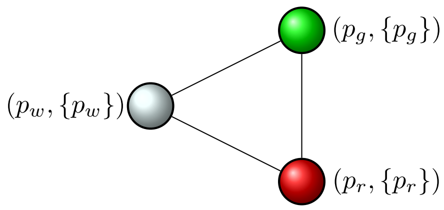


(10)

Consequently, the single-round protocol complex $\mathcal{P}_1 = \mathcal{P}_1^{\mathbf{D}^\omega}$ is just the subcomplex of the communication pseudosphere $\mathrm{Ps}(\sigma_0)$ induced by the set $\mathbf{D}$ of possible graphs. For example, $\mathcal{P}_1$ for $\mathbf{D} = \{G_1, G_2\}$ is the subcomplex of $\mathrm{Ps}(\sigma_0)$ made up by the two highlighted facets corresponding to the graphs $G_1$ and $G_2$ in Fig. 1. That is, rather than labeling the vertices of $\mathrm{Ps}(\sigma_0)$ with *all* the possible subsets of faces of $\sigma_0$ as in Definition 13, only those faces that are communicated via one of the graphs in $\mathbf{D}$ are used by the *protocol complex construction operator* $\mathcal{P} = \mathcal{P}^{\mathbf{D}^\omega}$ for generating $\mathcal{P}_1 = \mathcal{P}(\sigma_0)$. Conversely, if one interprets $\mathbf{D}$ as an $(n-1)$-dimensional simplicial complex $\mathbf{D}(\sigma_0)$, consisting of one facet (and all its faces) per graph $G \in \mathbf{D}$ according to (7), one could write $\mathcal{P}^{\mathbf{D}^\omega}(\sigma_0) = \mathrm{Ps}(\sigma_0) \cap \mathbf{D}(\sigma_0)$.

This can be compactly summarized in the following definition:

**Definition 14 (Protocol complex construction pseudosphere).** *Let $\mathcal{K}$ be an $(n-1)$-dimensional pure simplicial complex with a proper coloring $\chi : V(\mathcal{K}) \to \{p_1, \dots, p_n\}$, and $\mathrm{In}^G(p_i)$ be the set of processes that $p_i$ hears of in the communication graph $G \in \mathbf{D}$. We define the* protocol complex construction pseudosphere *$\mathcal{P}(\mathcal{K})$ for the message adversary $\mathbf{D}^\omega$, induced by the operator $\mathcal{P} : Fct(\mathcal{K}) \to \mathcal{P}(\mathcal{K})$ that can be applied to the facets of $\mathcal{K}$, through its vertex set and facets as follows:*

$$V(\mathcal{P}(\mathcal{K})) = \big\{(p_i, \sigma) \in \Pi \times F(\mathcal{K}) \mid \exists G \in \mathbf{D} : \mathrm{In}^G(p_i) = \chi(\sigma)\big\}, \tag{11}$$

$$Fct(\mathcal{P}(\mathcal{K})) = \big\{\{(p_1, \sigma_1), \dots, (p_n, \sigma_n)\} \mid \exists G \in \mathbf{D}, \forall 1 \leq i \leq n : \mathrm{In}^G(p_i) = \chi(\sigma_i)\big\}. \tag{12}$$

According to Definition 14, our operator $\mathcal{P}$ (as well as Ps) is actually defined only for the facets in $\mathcal{K}$, i.e., the dimension $n-1$ is actually implicitly encoded in the operator. We will establish below that this is sufficient for our purposes, since every $\mathcal{P}$ is *boundary consistent*: This property will allow us to uniquely define $\mathcal{P}$ for proper faces in $\mathcal{K}$ as well. We will use the following simple definition of boundary consistency, which makes use of the fact that the proper coloring of the vertices of a chromatic simplicial complex defines a natural ordering of the vertices of any of its faces.

**Definition 15 (Boundary consistency).** *We say that a protocol complex construction operator $\mathcal{P}$ according to Definition 14 is* boundary consistent*, if for all possible choices of three facets $\sigma$, $\kappa$ and $\tau$ from every simplicial complex on which $\mathcal{P}$ can be applied, it holds that*

$$\sigma \cap \kappa = \sigma \cap \tau \implies \mathcal{P}(\sigma) \cap \mathcal{P}(\kappa) = \mathcal{P}(\sigma) \cap \mathcal{P}(\tau). \tag{13}$$

The following Lemma 1 shows that every $\mathcal{P}$ is boundary consistent and that one can uniquely define $\mathcal{P}(\rho)$ also for a non-maximal simplex $\rho$ (taken as a complex). Moreover, it reveals that $\mathcal{P}$, viewed as a carrier map, is strict (but not necessarily rigid):

**Lemma 1 (Boundary consistency of $\mathcal{P}$).** *Every protocol construction operator $\mathcal{P}$ according to Definition 14 is boundary consistent. It can be applied to any simplex $\rho \in \mathcal{K}$, viewed as a complex, and produces a unique (possibly impure) chromatic complex $\mathcal{P}(\rho)$ with dimension at most $\dim(\rho)$. Moreover,*

$$\mathcal{P}(\sigma \cap \kappa) = \mathcal{P}(\sigma) \cap \mathcal{P}(\kappa) \tag{14}$$

*for any $\sigma, \kappa \in \mathcal{K}$.*

*Proof.* Using the notation from Definition 15, assume $\rho = \sigma \cap \kappa = \sigma \cap \tau$ for $0 \leq \dim(\rho) < n-1$; for the remaining cases, Eq. (14) holds trivially. Consider the facet $F_\sigma = \{(p_1, \sigma_1), (p_2, \sigma_2), \dots (p_n, \sigma_n)\}$ resp. $F_\kappa = \{(p_1, \kappa_1), (p_2, \kappa_2), \dots (p_n, \kappa_n)\}$ caused by the same graph $G \in \mathbf{D}$ in $\mathcal{P}(\sigma)$ resp. $\mathcal{P}(\kappa)$ according to Eq. (12). Recall that $\sigma_i$ resp. $\kappa_i$ is a face of $\sigma$ resp. $\kappa$ that represents the information $p_i$ receives from the processes in $\chi(\sigma_i) = \chi(\kappa_i)$ via $\mathrm{In}^G(p_i)$.

A vertex $(p_i, \kappa_i)$ appears in $\mathcal{P}(\sigma) \cap \mathcal{P}(\kappa)$ if and only if $\kappa_i = \sigma_i$, which, in turn, holds only if $\chi(\kappa_i) \subseteq \chi(\rho)$. Indeed, if $\kappa_i$ would contain just one vertex $v \in V(\kappa)$ with $\chi(v) \in \chi(\kappa \setminus \rho)$, then $\sigma_i$ would contain the corresponding vertex $v' \in V(\sigma)$ with $\chi(v') = \chi(v)$ satisfying $v' \neq v$ since $(\kappa \setminus \rho) \cap (\sigma \setminus \rho) = \emptyset$, by the definition of $\rho$. This would contradict $\kappa_i = \sigma_i$, however. Note that, since $p_i \in \chi(\sigma_i)$ for every $i$, this also implies $p_i \in \chi(\rho)$.

Consequently, it is precisely the maximal face in $F_\sigma$ (and in $F_\kappa$) consisting only of identical vertices $(p_i, \kappa_i) = (p_i, \sigma_i)$ that appears in $\mathcal{P}(\sigma) \cap \mathcal{P}(\kappa)$. Since this holds for all graphs $G \in \mathbf{D}$, it follows that the subcomplex $\mathcal{P}(\sigma) \cap \mathcal{P}(\kappa)$, as the union of the resulting identical maximal faces, has dimension at most $\dim(\rho)$. Now, since exactly the same reasoning also applies when $\kappa$ is replaced by $\tau$, we get $\mathcal{P}(\sigma) \cap \mathcal{P}(\kappa) = \mathcal{P}(\sigma) \cap \mathcal{P}(\tau)$, so Eq. (13) and hence boundary consistency of $\mathcal{P}$ holds.

We can now just *define* $\mathcal{P}(\rho) = \mathcal{P}(\sigma \cap \kappa) := \mathcal{P}(\sigma) \cap \mathcal{P}(\kappa)$, which secures Eq. (14) for facets $\sigma, \kappa \in \mathcal{K}$. For general simplices, assume for a contradiction that there are $\sigma$, $\kappa$ with $\rho = \sigma \cap \kappa \neq \emptyset$ but $\mathcal{P}(\rho) \neq \mathcal{P}(\sigma) \cap \mathcal{P}(\kappa)$. Choose facets $\sigma'$, $\kappa'$ and $\sigma''$, $\kappa''$ satisfying $\rho = \sigma' \cap \kappa'$, $\rho = \sigma'' \cap \kappa''$, $\sigma = \sigma' \cap \sigma''$ and $\kappa = \kappa' \cap \kappa''$, which is always possible. Applying Eq. (14) to all these pairs results in $\mathcal{P}(\rho) = \mathcal{P}(\sigma') \cap \mathcal{P}(\kappa') = \mathcal{P}(\sigma'') \cap \mathcal{P}(\kappa'')$, $\mathcal{P}(\sigma) = \mathcal{P}(\sigma') \cap \mathcal{P}(\sigma'')$ and $\mathcal{P}(\kappa) = \mathcal{P}(\kappa') \cap \mathcal{P}(\kappa'')$. We hence find

$$\mathcal{P}(\rho) \neq \mathcal{P}(\sigma) \cap \mathcal{P}(\kappa) = \mathcal{P}(\sigma') \cap \mathcal{P}(\sigma'') \cap \mathcal{P}(\kappa') \cap \mathcal{P}(\kappa'') = \mathcal{P}(\rho) \cap \mathcal{P}(\sigma'') \cap \mathcal{P}(\kappa'') = \mathcal{P}(\rho),$$

which is a contradiction. □

Note that $\mathcal{P}$ can hence indeed be interpreted as a carrier map, according to Definition 8, which is strict. It is well known that strictness implies that, for any simplex $\rho \in \mathcal{P}(\mathcal{K})$, there is a unique simplex $\sigma$ with smallest dimension in $\mathcal{K}$, called the *carrier* of $\rho$, such that $\rho \in \mathcal{P}(\sigma)$.

A comparison with Definition 12 reveals that $\mathcal{P}_1 = \mathcal{P}(\sigma_0)$ as given in Definition 14 is indeed just the uninterpreted 1-round protocol complex. The general $r$-round uninterpreted protocol complex $\mathcal{P}_r$, $r \geq 1$, is defined as $\mathcal{P}(\mathcal{P}_{r-1})$, i.e., as the union of $\mathcal{P}$ applied to every facet $\sigma$ of $\mathcal{P}_{r-1}$, formally $\mathcal{P}_r = \bigcup_{\sigma \in \mathcal{P}_{r-1}} \mathcal{P}(\sigma)$. Boundary consistency ensures that $\mathcal{P}_r = \mathcal{P}^r(\sigma_0)$ for the initial simplex $\sigma_0 = \{(p_1, \{p_1\}), \dots, (p_n, \{p_n\})\}$ is well-defined for any $r \geq 0$. An example for $r = 2$ can be found in the bottom part of Fig. 3. Note that the arrows of the in-edges of a vertex $(p_i, h_r(p_i))$ in a facet in $\mathcal{P}_r$ represent the outermost level in Definition 11; the labeling of the in-edges of $p_i$ in earlier rounds $< r$ is no longer visible here. However, given the simplex $\rho = \{(p_1, h_r(p_1)), \dots, (p_n, h_r(p_n))\} \in \mathcal{P}_r$, the labeling of the vertices $(p_j, h_{r-1}(p_j)) \in V(\sigma)$ of the carrier $\sigma \in \mathcal{P}_{r-1}$ of $\rho$, i.e., the unique simplex satisfying $\rho \in \mathcal{P}_1(\sigma)$, can be used to recover the arrows for round $r - 1$.

We note that $\mathcal{P}_r = \mathcal{P}^{r-1}(\mathcal{P}(\sigma_0)) = \mathcal{P}(\mathcal{P}^{r-1}(\sigma_0))$ allows to view the construction of $\mathcal{P}_r$ equivalently as applying the one-round construction $\mathcal{P}$ to every facet $F_{r-1}$ of $\mathcal{P}_{r-1}$ or else as applying the $(r-1)$-round construction $\mathcal{P}^{r-1}$ to every facet $F$ of $\mathcal{P}_1$. Boundary consistency of $\mathcal{P}$ again ensures that this results in exactly the same protocol complex. Our decision procedure for consensus solvability/impossibility provided in Section 4 will benefit from the different views provided by this construction.

In the remainder of this section, we will discuss some consequences of the fact that the carrier map corresponding to a general protocol complex construction operator $\mathcal{P}$ is always strict but need not be rigid (recall Lemma 1). This is actually a consequence of the asymmetry in the protocol complex construction caused by graphs $\mathbf{D}$ that do not treat all processes alike.

Consider the complete communication pseudosphere shown in Fig. 1, which corresponds to $\mathbf{D}$ containing *all* possible graphs with $n$ vertices (recall Definition 13). It does treat all processes alike, which also implies that its outer border, which is defined by $\mathcal{P}(\partial\sigma_0)$ (see Definition 17 below), has a very regular structure: For example, the four white and green vertices aligned on the bottom side of the outer triangle of Fig. 1 are actually an instance of the 2-process communication pseudosphere shown in (8). Its corresponding carrier map is rigid. By contrast, the protocol complex for the message adversary $\mathbf{D} = \{G_1, G_2\}$ depicted by the two highlighted facets corresponding to the graphs $G_1$ and $G_2$ in Fig. 1 has a very irregular border shown in Fig. 2.

It is worth mentioning, though, that there are other instances of protocol complex construction operators that also have a rigid equivalent carrier map. One important example is the popular standard chromatic subdivision [20, 25], which characterizes the iterated immediate snapshot (IIS) model of shared memory [23]:

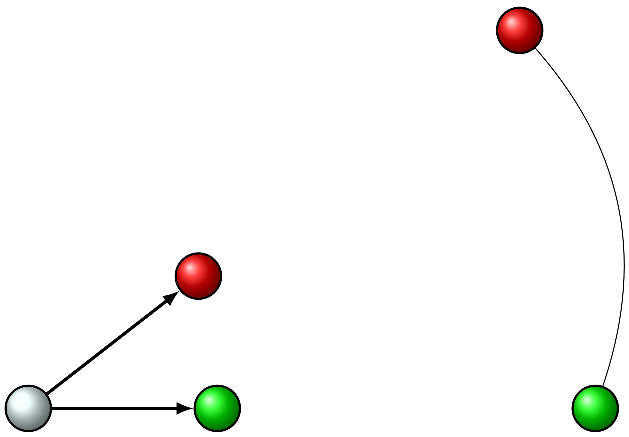

**Fig. 2.** Border $\mathrm{Bd}(\mathcal{P}_1)$ of the simple message adversary shown in Fig. 1.

**Definition 16 (Chromatic subdivision).** *Let $\mathcal{K}$ be an $(n-1)$-dimensional simplicial complex with a proper coloring $\chi : V(\mathcal{K}) \to \{p_1, \ldots, p_n\}$. We define the chromatic subdivision through its vertex set and facets as follows:*

$$V(\mathrm{Ch}(\mathcal{K})) = \big\{(p_i, \sigma) \in \Pi \times F(\mathcal{K}) \mid p_i \in \chi(\sigma)\big\}, \tag{15}$$

$$Fct(\mathrm{Ch}(\mathcal{K})) = \big\{\{(p_1, \sigma_1), \ldots, (p_n, \sigma_n)\} \mid \exists \pi : [n] \to [n] \text{ permutation } \sigma_{\pi(1)} \subseteq \ldots \subseteq \sigma_{\pi(n)}, \\ \forall 1 \le i, j \le n : \chi(\sigma_i) \wedge (p_i \in \chi(\sigma_j) \Rightarrow \sigma_i \subseteq \sigma_j)\big\}. \tag{16}$$

It is immediately apparent from comparing Definition 13 and Definition 16 that $V(\mathrm{Ps}(\sigma^{n-1})) = V(\mathrm{Ch}(\sigma^{n-1}))$ and $\mathrm{Ch}(\sigma^{n-1}) \subseteq \mathrm{Ps}(\sigma^{n-1})$, i.e., $\mathrm{Ch}(\sigma^{n-1})$ is indeed a subcomplex of $\mathrm{Ps}(\sigma^{n-1})$. In Fig. 1, we have highlighted, via thick edges and arrows, the protocol complex $\mathrm{Ch}(\sigma_0)$ for the corresponding message adversary. In fact, the chromatic subdivision and hence the IIS model is just a special case of our oblivious message adversary, the set $\mathbf{D}$ of which consists of all the directed graphs that are unilaterally connected ($\forall G \in \mathbf{D}, a, b \neq a \in V(G)$ : $\exists$ directed path from $a$ to $b$ or from $b$ to $a$ in $G$) and transitively closed ($\forall G \in \mathbf{D} : (a, b), (b, c) \in E(G) \Rightarrow (a, c) \in E(G)$).

**Lemma 2 (Equivalent message adversary for chromatic subdivision).** *Let $\sigma_0$ be the uninterpreted input complex with process set $\Pi = \{p_1, \ldots, p_n\}$, and $\mathbf{D}$ be the set of all unilaterally connected and transitively closed graphs on $\Pi$. Then, $\mathcal{P}(\sigma_0) = \mathrm{Ch}(\sigma_0)$.*

*Proof.* Notice first that for any face $\sigma \in \sigma_0$ such that $p_i \in \chi(\sigma)$, there exists a graph $G_\sigma \in \mathbf{D}$ such that $\mathrm{In}^{G_\sigma}(p_i) = \chi(\sigma)$: simply consider $E(G_\sigma) = \{(u, v) \mid u \in \chi(\sigma) \wedge v \in \Pi\} \cup \{(w, y) \mid w, y \notin \chi(\sigma)\}$. By construction, $G_\sigma$ is both transitively closed and unilaterally connected. Therefore, $V(\mathrm{Ch}(\sigma_0)) \subseteq V(\mathcal{P}(\sigma_0))$. On the other hand, from Definition 14 of the protocol complex pseudosphere construction, it follows that $V(\mathcal{P}(\sigma_0)) \subseteq V(\mathrm{Ch}(\sigma_0))$. Consequently, $V(\mathcal{P}(\sigma_0)) = V(\mathrm{Ch}(\sigma_0))$.

Let $\sigma = \{(p_1, \sigma_1), \ldots, (p_n, \sigma_n)\}$ be a facet of $\mathrm{Ch}(\sigma_0)$, and consider the graph $G_\sigma$ with edges $E(G_\sigma) = \{(p_j, p_i) \mid p_j \in \chi(\sigma_i)\}$. Assume that $(p_i, p_j)$, $(p_j, p_k) \in E(G_\sigma)$. By Definition 16, it holds that $p_i \in \chi(\sigma_j)$, which implies $\sigma_i \subseteq \sigma_j$. Analogously, $\sigma_j \subseteq \sigma_k$ and therefore $\sigma_i \subseteq \sigma_k$. It hence follows that $p_i \in \chi(\sigma_k)$, and by construction of $G_\sigma$, that $(p_i, p_k) \in E(G_\sigma)$. This shows that $G_\sigma$ is transitively closed.

Now consider $p_i, p_j \in \Pi$. Since $\pi$ is a permutation, $p_i = p_{\pi(i')}$ and $p_j = p_{\pi(j')}$ for some $i', j' \in [n]$. Let us assume w.l.o.g that $i' \le j'$. Then $\sigma_i = \sigma_{\pi(i')} \subseteq \sigma_{\pi(j')} = \sigma_j$, which implies that $p_i \in \chi(\sigma_j)$. From the definition of $G_\sigma$, it follows that $(p_i, p_j) \in E(G_\sigma)$. This shows that $G_\sigma$ is also unilaterally connected. Therefore, $\sigma$ must also be a facet of $\mathcal{P}(\sigma_0)$, i.e., $Fct(\mathrm{Ch}(\sigma_0)) \subseteq Fct(\mathcal{P}(\sigma_0))$.

Conversely, let $\sigma = \{(p_1, \sigma_1), \ldots, (p_n, \sigma_n)\}$ be a facet of $\mathcal{P}(\sigma_0)$. Let $G_\sigma$ be the graph from $\mathbf{D}$ that induces $\sigma$. Recall that $G_\sigma$ is unilaterally connected and transitively closed. Let $S_i$ denote the strongly connected component containing $p_i$. Since $G_\sigma$ is transitively closed, $S_i$ is in fact a directed clique. Therefore, $S_i \subseteq \chi(\sigma_i) = \mathrm{In}^{G_\sigma}(p_i)$. Consider the component graph $G^*$ where $V(G^*) = \{S_i \mid i \in [n]\}$, and $E(G^*) = \{(S_i, S_j) \mid (p_i, p_j) \in E(G_\sigma)\}$. Since $G_\sigma$ is transitively closed and unilaterally connected, $G^*$ is a transitive tournament (where $(a, b)$ or $(b, a)$ must be present for all $a, b \in V(G^*)$). Therefore, $G^*$ has a directed Hamiltonian path $S_{\pi(1)}, \dots, S_{\pi(s)}$ for $s = |V(G^*)|$; note that $s \leq n$ since $S_i = S_j$ may be the same for different processes $p_i$ and $p_j$.

Clearly, the permutation from the Hamiltonian path of connected components, extended by ordering processes leading to the same connected component according to their ids, induces a complete ordering of the process indices: $i \preceq j$ if $S_i = S_{\pi(i')}$ and $S_j = S_{\pi(j')}$ with $i' \leq j'$ and $i \leq j$, i.e., first we order each index $i$ according to the order of their connected component in the Hamiltonian path in $G^*$, and break ties according to their process ids. Therefore, $\preceq$ is a total ordering on $[n]$, and thus induces a permutation $\pi'$ with the property that if $i \leq j$, then either $S_{\pi'(i)} = S_{\pi'(j)}$, or there exists an edge from $S_{\pi'(i)}$ to $S_{\pi'(j)}$.

From the transitive closure of $G_\sigma$ and the construction of $\pi'$, we get $\mathrm{In}^{G_\sigma}(\pi'(p_i)) = \bigcup_{j=1}^{i} S_{\pi'(i)}$. Therefore, $\pi'$ is also a permutation of the $\sigma_i$ in $\sigma$ that satisfies the conditions for being a facet of $\mathrm{Ch}(\sigma_0)$. It follows that $\sigma \in Fct(\mathrm{Ch}(\sigma_0))$. Therefore $Fct(\mathcal{P}(\sigma_0)) \subseteq Fct(\mathrm{Ch}(\sigma_0))$, which completes the proof that $Fct(\mathcal{P}(\sigma_0)) = Fct(\mathrm{Ch}(\sigma_0))$ and thus $\mathcal{P}(\sigma_0) = \mathrm{Ch}(\sigma_0)$. □

For any pair of simplices $\sigma, \kappa \in \mathcal{K}$, it hence holds by Eq. (14) that $\mathrm{Ch}(\sigma) \cap \mathrm{Ch}(\kappa) = \mathrm{Ch}(\sigma \cap \kappa)$, i.e., subdivided simplices that share a face intersect precisely in the subdivision of that face in $\mathrm{Ch}(\mathcal{K})$. Lemma 1 thus ensures that the iterated standard chromatic subdivision $\mathrm{Ch}^r(\mathcal{K})$ is well-defined.

Thanks to its regular structure, the equivalent carrier map is also rigid. As is the case for the communication pseudosphere in Fig. 1, the four white and green vertices aligned on the bottom side of the outer triangle connected by thick arrows are actually an instance of a 2-process chromatic subdivision. Indeed, the standard chromatic subdivision $\mathrm{Ch}(\sigma^\ell)$ of a simplex $\sigma^\ell$ of dimension $\ell$ can be constructed iteratively [20]: Starting out from the vertices $V(\sigma^\ell)$, i.e., the 0-dimensional faces $\sigma^0$ of $\sigma^\ell$, where $\mathrm{Ch}(\sigma^0) = \sigma^0$, one builds $\mathrm{Ch}(\sigma^1)$ for the edge $\sigma^1$ by placing 2 new vertices in its interior and connecting them to each other and to the vertices of $\sigma^1$. For constructing $\mathrm{Ch}(\sigma^3)$, one places 3 new vertices in its interior and connects them to each other and to the vertices constructed before, etc.

**Corollary 1.** *Let $\mathcal{K}$ be an arbitrary simplicial complex, then* $\mathrm{Ch}(\mathcal{K}) = \mathcal{P}(\mathcal{K})$ *with* $\mathbf{D}$ *as the set of allowed graphs.*

*Proof.* Follows immediately from Lemma 2 and boundary consistency of $\mathcal{P}(\mathcal{K})$.

### 3.3 Classification of facets of protocol complexes

We first define the important concept of the border of a protocol complex.

**Definition 17 (Border).** *The* border $\mathrm{Bd}(\mathcal{P}_1)$ *of a 1-round protocol complex* $\mathcal{P}_1 = \mathcal{P}(\sigma_0)$ *is defined as* $\mathrm{Bd}(\mathcal{P}_1) = \mathcal{P}(\partial\sigma_0)$. *The* border $\mathrm{Bd}(\mathcal{P}_r)$ *(resp. the border* $\mathrm{Bd}(\mathcal{C})$ *of some subcomplex* $\mathcal{C} \subseteq \mathcal{P}_r$*) of the general r-round complex* $\mathcal{P}_r = \mathcal{P}^r(\sigma_0)$ *is* $\mathrm{Bd}(\mathcal{P}_r) = \mathcal{P}^r(\partial\sigma_0)$.

Due to the boundary consistency property of $\mathcal{P}$ (Lemma 1), the border is just the "outermost" part of $\mathcal{P}_r$, i.e., the part that is carried by $\partial\sigma_0$; the dimension of every facet $F \in \mathrm{Bd}(\mathcal{P}_r)$ is at most $\dim(\sigma_0) - 1 = n - 2$. Recall that it may also be smaller than $n - 2$, since $\mathcal{P}$ viewed as a carrier map need not be rigid. Obviously, however, $F$ is always a face of some facet in $\mathcal{P}_r$. In the case of Fig. 1, where $\mathcal{P}_1 = \{G_1, G_2\}$ with the graphs $G_1, G_2$ given in Eqn. (9), $\mathrm{Bd}(\mathcal{P}_1)$ only consists of the three edges and the vertices shown in Fig. 2. Observe that the processes of the vertices $V(\rho)$ of a face $\rho \in \mathrm{Bd}(\mathcal{P}_r)$ may possibly have heard from each other, but not from processes in $\Pi \setminus V(\rho)$, in any round $1, \dots, r$.

For a more elaborate running example, consider the RAS message adversary shown in Fig. 3: Its 1-round uninterpreted complex $\mathcal{P}_1^{RAS}$ (top left part) is reminiscent of the well-known radioactivity sign, hence its name. Its 2-round uninterpreted complex $\mathcal{P}_2^{RAS}$ is shown in the bottom part of the figure. It is constructed by taking the union of the 1-round uninterpreted complexes $\mathcal{P}(F)$ for every facet $F \in \mathcal{P}_1$. Its border $\mathrm{Bd}(\mathcal{P}_2^{RAS})$ is formed by all the vertices and edges of the faces that lie on the (dotted and partly dash-dotted) borders of the outermost triangle.

For classifying the facets in a protocol complex, the *root components* of the graphs in $\mathbf{D}$ will turn out to be crucial.

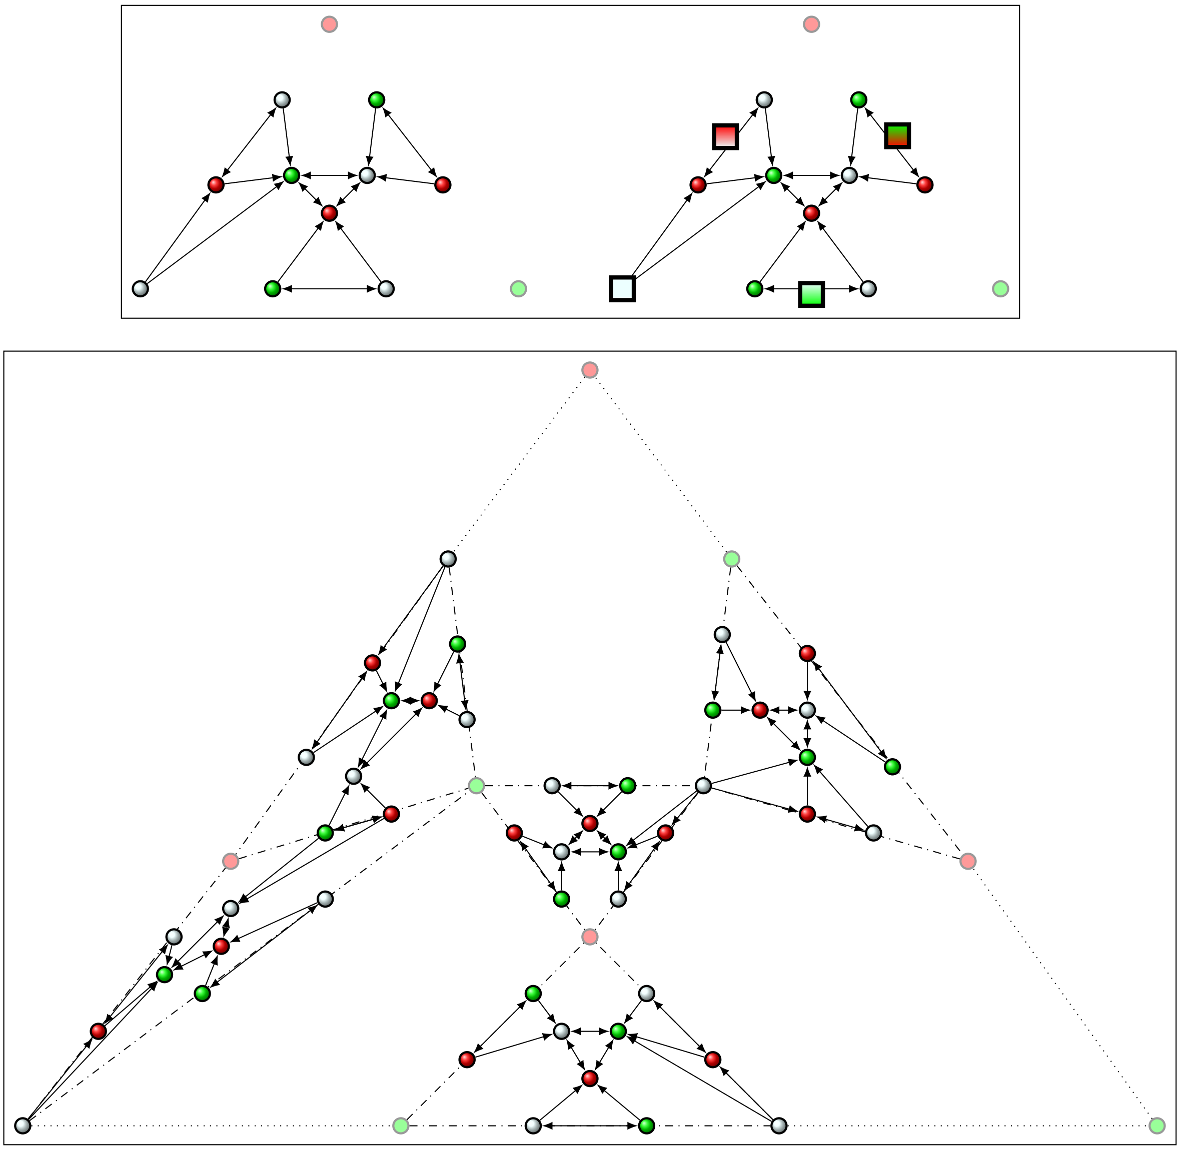

**Fig. 3.** Protocol complex $\mathcal{P}_1^{RAS}$ for one round (top) and $\mathcal{P}_2^{RAS}$ for two rounds (bottom) of the RAS message adversary. The top right figure also shows the border root components of $\mathcal{P}_1^{RAS}$.

**Definition 18 (Root components).** *Given any facet $F$ in the protocol complex $\mathcal{P}_r$, $r \geq 1$, let $\sigma \in \mathcal{P}_{r-1}$ be its carrier, i.e., the unique facet such that $F \in \mathcal{P}(\sigma)$, and $G \in \mathbf{D}$ be the corresponding graph leading to $F$ in $\mathcal{P}(\sigma)$. A* root component *$R(F)$ of $F$ is the face of $F$ corresponding to a strongly connected component $R$ in $G$ without incoming edges from $G \setminus R$.*

It is well-known that every directed graph with $n$ vertices has at least one and at most $n$ root components, and that every process in $G$ is reachable from every member of at least one root component via some directed path in $G$. Graphs with a single root component are called rooted, and it is easy to see that just one graph in $\mathbf{D}$ that is not rooted makes consensus trivially impossible: The adversary might repeat this graph forever, preventing the different root components from coordinating the output value. In the sequel, we will therefore restrict our attention to message adversaries where $\mathbf{D}$ is made up of rooted graphs only, and will denote by $R(G) = R(F)$ the face representing the root component of $F$. Note that $R(G)$ is a face and hence includes the edges of the interconnect and their orientation; its set of vertices is denoted by $V(R(F)) = \{(p_i, h_r(p_i)) \mid p_i \in \chi(R(F))\}$. Recall from Definition 11 that the set of processes that $p_i$ has actually heard of in some vertex $v = (p_i, h_r(p_i)) \in V(\rho)$ is denoted $\cup h_r(p_i)$.

**Definition 19 (Border facets).** *A facet $F \in \mathcal{P}_r$ is a* border facet*, if the subcomplex $F \cap \mathrm{Bd}(\mathcal{P}_r)$ is non-empty. The subcomplex $F \cap \mathrm{Bd}(\mathcal{P}_r)$ will be called* facet borders *of $F$. A border facet $F$ is* proper *if the members of the root component did not collectively hear from all processes, i.e., $\bigcup_{(p_i,h_r(p_i))\in V(R(F))} \cup h_r(p_i) \neq \Pi$.*

Intuitively, a border facet $F \in \mathcal{P}_r$ has at least one vertex $v \in \mathrm{Bd}(\mathcal{P}_r)$. It is immediately apparent that $v$ may have heard at most from processes in some face $\rho \in \mathrm{Bd}(\mathcal{P}_r)$, which has dimension at most $n-2$, but not from processes outside $\rho$ (so, in particular, not from all processes).

The facet borders $F \cap \mathrm{Bd}(\mathcal{P}_r)$ of a border facet $F$ form indeed a subcomplex in general, rather than a single face, as is the case in, e.g., the left part of Fig. 2 (generated by $F$ that represents the graph $G_2$) shows. Moreover, $F \cap \mathrm{Bd}(\mathcal{P}_r)$ does not even need to be connected. For example, if the message adversary of Fig. 2 would also include the graph $G_3 = \{r \to g \to w\}$, i.e., a chain (with root component $r$), we observe $F_3 \cap \mathrm{Bd}(\mathcal{P}_1) = \{r \to g, w\}$ for the corresponding facet $F_3$. Finally, it need not even be the case that $F \cap \mathrm{Bd}(\mathcal{P}_r)$ contains the entire root component $R(F)$: Since $\dim(\mathrm{Bd}(\mathcal{P}_r)) = n-2$, this is inevitable if $F$ is not a proper border facet, i.e., if the members of $R(F)$ have collectively heard from all processes. For instance, if the message adversary of Fig. 2 also contained the cycle $G_4 = \{r \to g \to w \to r\}$ (with root component $R(F_4) = F_4 = \{r \to g \to w \to r\}$ consisting of all processes), then the (improper) border facet $F_4 \cap \mathrm{Bd}(\mathcal{P}_1) = \{r, g, w\}$ obviously cannot contain $R(F_4)$.

**Definition 20 (Border components and border root components).** *For every proper border facet $F \in \mathcal{P}_r$, the* border component *$B(F)$ is the smallest face of $F$ whose members did not hear from processes outside of $B(F)$, that is, $\bigcup_{(p_i,h_r(p_i))\in V(B(F))} \cup h_r(p_i) \subseteq \chi(B(F))$. For a facet $F$ that is not a proper border facet, we use the convention $B(F) = F$ for completeness. The set of all proper border components of $\mathcal{P}_r$ is denoted as $\mathrm{BdC}(\mathcal{P}_r)$ (with the appropriate restriction $\mathrm{BdC}(\mathcal{C})$ for a subcomplex $\mathcal{C} \subseteq \mathcal{P}_r$).*

*The root component $R(F)$ of a proper border facet $F$ is called* border root component*; it necessarily satisfies $R(F) \neq F$. The set of all border root components of $\mathcal{P}_r$ resp. a subcomplex $\mathcal{C} \subseteq \mathcal{P}_r$ is denoted $\mathrm{BdR}(\mathcal{P}_r)$ resp. $\mathrm{BdR}(\mathcal{C})$.*

Lemma 3 below will assert that the border component of a facet is unique and contains its root component.

**Definition 21 (Border component carrier).** *The* border component carrier *$\beta(F)$ of a proper border facet $F$ is the smallest face of the initial simplex $\sigma_0 = \{(p_1, \{p_1\}), \ldots, (p_n, \{p_n\})\}$ that carries $B(F)$. For a facet $F$ that is not proper, we use the convention $\beta(F) = \sigma_0$ for consistency.*

Since $\chi\big(\beta(F)\big) = \chi(B(F))$, it is apparent that $\beta(F)$ implicitly also characterizes the heard-of sets of the processes in $B(F)$: According to Definition 20, its members may have heard from processes in $\beta(F)$ but not from other processes. Note carefully that this also tells something about the knowledge of the processes regarding the initial values of other processes, as the members of $B(F)$ can at most know the initial values of the processes in $\beta(F)$.

For an illustration, consider the top right part of Fig. 3, which shows the border root components of border facets of $\mathcal{P}_1^{RAS}$ for the RAS message adversary, represented by square nodes with fat borders. $\mathrm{BdR}_2^{RAS}$ depicted in the bottom part of Fig. 3 consists of all faces $B(F)$ of border facets $F$ touching the outer border: Going in clockwise direction, starting with the bottom-leftmost border face, we obtain the following pairs (border root component, border component carrier) representing $B(F)$ of a border facet: $(w, \{w\})$, $(r \leftrightarrow w, \{r, w\})$, $(r \leftrightarrow w, \{r, w\})$, $(w, \{w, r\})$, $(r \leftrightarrow g, \{r, g\})$, $(w, \{w, g\})$, $(w \leftrightarrow g, \{w, g\})$. It is apparent that the members of the border root component $w \leftrightarrow g$ in the last pair $(w \leftrightarrow g, \{w, g\})$ (representing the border facet on the bottom) only know their own initial values, but not the initial value of the red process.

**Lemma 3 (Properties of border component of a proper border facet).** *The border component $B(F)$ of a proper border facet $F \in \mathcal{P}_r$ satisfies the following properties:*

*(i)* $B(F)$ *is unique,*
*(ii)* $R(F) \subseteq B(F) \subseteq F \cap \mathrm{Bd}(\mathcal{P}_r)$*, which also implies* $B(F) \neq F$*,*
*(iii)* $B(F) = R(F)$ *for* $r = 1$*, but possibly* $R(F) \subsetneq B(F)$ *for* $r \geq 2$*.*

*Proof.* As for (i), assume for a contradiction that there is some alternative $B'(F)$ of the same size. Due to Definition 18, both $R(F) \subseteq B(F)$ and $R(F) \subseteq B'(F)$ must hold, since some process in $B(F)$ would have heard from a process in $R(F) \setminus B(F)$ otherwise, and, analogously, for $B'(F)$. As $B(F) \neq B'(F)$, there is a $v' = (p'_i, h_r(p'_i)) \in B'(F) \setminus B(F)$ that is present in $B'(F)$ because some $v \in B(F) \cap B'(F)$ has heard from $p'_i$ earlier. But then, $v'$ is also present in $B(F)$, a contradiction. For (ii), besides $R(F) \subseteq B(F)$, we also have $\sigma = R(F) \cap \mathrm{Bd}(\mathcal{P}_r) \neq \emptyset$, since $R(F)$ of a proper border facet according to Definition 19 does not encompass all processes. Now assume first that $R(F) \not\subseteq F \cap \mathrm{Bd}(\mathcal{P}_r)$, i.e., $R(F) \not\subseteq \mathrm{Bd}(\mathcal{P}_r)$ (since $R(F) \subseteq F$ obviously always holds). For every facet $\sigma \in \mathrm{Bd}(\mathcal{P}_r)$, there is hence some $v = (p_i, h_r(p_i)) \in R(F) \setminus \sigma \neq \emptyset$. However, by the properties of root components, some process $p_j \in \chi(\sigma)$ must have heard from $p_i \notin \chi(\sigma)$, which would contradict $p_j \in \chi(\sigma)$. Therefore, we must have $R(F) \subseteq \mathrm{Bd}(\mathrm{Pr})$. For the final contradiction, by the same token, assume that $B(F) \setminus \mathrm{Bd}(\mathcal{P}_r) \neq \emptyset$, i.e., for any facet $\sigma \in \mathrm{Bd}(\mathcal{P}_r)$, there is some $v = (p_i, h_r(p_i)) \in B(F) \setminus \sigma$. By the definition of border components according to Definition 20, however, such a $v$ exists only if some process $p_j \in \chi(\sigma)$ has already heard from $p_i \notin \chi(\sigma)$, which would contradict $p_j \in \chi(\sigma)$. Thus, $B(F) \subseteq F \cap \mathrm{Bd}(\mathcal{P}_r)$. Finally, $B(F) \neq F$ follows from the fact that $B(F) \subseteq \mathrm{Bd}(\mathcal{P}_r)$ imposes a maximum dimension of $n-2$ for $B(F)$.

As for (iii), $B(F) = R(F)$ for $r = 1$ follows immediately from Definition 20. For $r \geq 2$, it is of course possible that some process in $R(F)$ has already heard from a process outside $R(F)$ in some earlier round, see the bottom-left facet $F_2^l$ in Fig. 4 for an example. □

We point out that the facet borders $F \cap \mathrm{Bd}(\mathcal{P}_r)$ of a proper border facet $F \in \mathcal{P}_r$ for $r \geq 2$ may also contain (small) faces in $\mathrm{Bd}(\mathcal{P}_r)$ that are disjoint from $B(F)$, albeit they are of course always contained in a larger face of $F$ that also contains $B(F)$. Examples can be found in the left part of Fig. 2, where $B(F) = \{w\}$ is disjoint from the single vertex $\{g\} \in F \cap \mathrm{Bd}(\mathcal{P})$, or in the top-left facet $F_1^l$ in Fig. 4, where $B(F_1^l) = w$ is disjoint from the single vertex $\{g\} \in F_1^l \cap \mathrm{Bd}(\mathcal{P}_1)$.

We conclude this section by stressing that border facets and border (root) components in $\mathcal{P}_r$ for $r \geq 2$ implicitly represent *sequences* of faces: A single border facet $F$ and the corresponding $R(F) \in \mathrm{BdR}(\mathcal{P}_r)$ resp. $B(F) \in \mathrm{BdC}(\mathcal{P}_r)$ actually represent all the border facets $F_1, \dots, F_r$ in rounds $1, \dots, r$ that carry $F = F_r$, and their corresponding $R(F_1), \dots, R(F_r)$ resp. $B(F_1), \dots, B(F_r)$. Note carefully that, although $\beta(F_{i+k})$, $k \geq 1$, is typically not smaller than $\beta(F_i)$, in particular, if $\chi\big(R(F_{i+k})\big) = \chi\big(R(F_i)\big)$, this need not always be the case, since a process present in $B(F_i)$ may have heard from all other processes in $F_{i+k}$, so that it is no longer present in $\mathrm{Bd}(F_{i+k})$ and hence in $\beta(F_{i+k})$. Fig. 4 provides some additional illustrating examples of border component carriers and border root components.

## 4 Consensus Solvability/Impossibility

In this section, we will characterize consensus solvability/impossibility under an oblivious message adversary $\mathbf{D}^\omega$ by means of the topological tools introduced in Section 3. Due to its "geometrical" flavor, our topological view not only provides interesting additional insights, but also prepares the ground for additional results provided in Section 5.

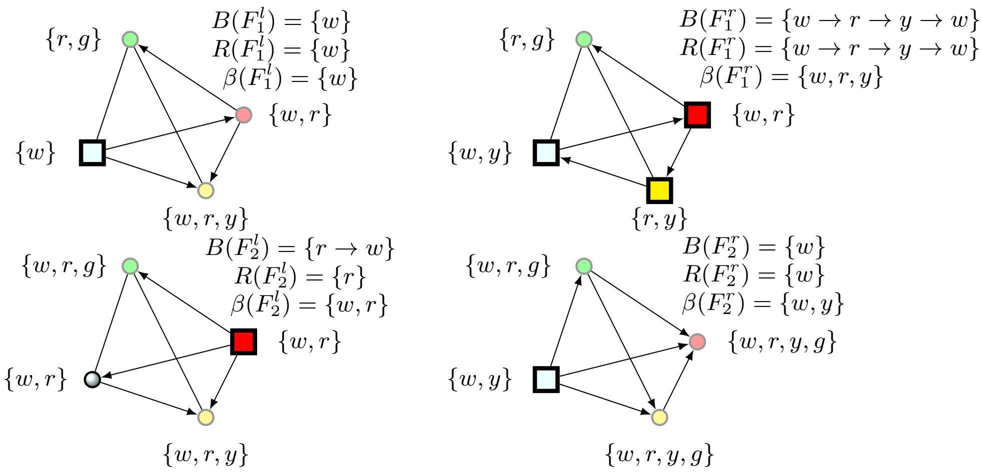


**Fig. 4.** Illustration of border components and border root components, for two different examples (left, right column) for $n = 4$ processes. Faded nodes represent vertices outside $B(F)$; squared nodes represent members of the border root component. The first row shows the respective border facets in $\mathcal{P}_1$, the second row shows the border facets in $\mathcal{P}_2$. The labels provide the heard-of set of the nearby process, assuming that the round-2 facet is applied to the round-1 facet atop of it. Observe that the red and yellow process in the bottom-right facet $F_2^r$ have heard from everybody and are hence removed from $\mathrm{Bd}(F_2^r)$ and $\beta(F_2^r)$.

The key insight of Section 4.1 is that one cannot solve consensus in $r$ rounds if the $r$-round protocol complex $\mathcal{P}_r$ comprises a connected component that contains *incompatible* proper border facets (defined as having a set of border components with an empty intersection). In Section 4.2, we focus on paths connecting pairs of incompatible proper border facets in $\mathcal{P}_{r-1}$, and exhaustively characterize what happens to such a path when $\mathcal{P}_{r-1}$ evolves to $\mathcal{P}_r$: It may either break, in which case consensus might be solvable in $\mathcal{P}_r$ (unless some other path still prevents it), or it may be lifted, in which it still prohibits consensus. In Section 4.3, we recast our path-centric characterization in terms of its effect on the connected components in the evolution from $\mathcal{P}_{r-1}$ to $\mathcal{P}_r$. A suite of examples in Section 4.4 illustrates all the different cases. Finally, Section 4.5 presents a alternative (and sometimes more efficient) formulation of the consensus decision procedure given in [39], which follows right away from the topological characterization of consensus solvability developed in the previous subsections.

### 4.1 Incompatibility of border components

Consensus is impossible to solve in $r$ rounds if the $r$-round protocol complex $\mathcal{P}_r$ has a connected component $\mathcal{C}$ that contains $k \geq 1$ proper border facets $\hat{F}_1, \dots, \hat{F}_k$ with incompatible border components $B(\hat{F}_1), \dots, B(\hat{F}_k) \in \mathrm{Bd}(\mathcal{P}_r)$, where incompatibility means $\bigcap_{i=1}^{k} \beta(\hat{F}_i) = \emptyset$: Since no vertex of $B(\hat{F}_i)$ could have had incoming edges from processes outside $B(\hat{F}_i)$, in any of the rounds $1, \dots, r$, their corresponding processes cannot decide on anything but one of their own initial values. As all vertices in a connected component must decide on the same value, however, this is impossible.

Incompatible border components occur, in particular, when $\hat{F}_1, \dots, \hat{F}_k$ have incompatible border root components $R(\hat{F}_1), \dots, R(\hat{F}_k) \in \mathrm{BdR}(\mathcal{P}_r)$. An instance of this situation can be seen in the top right part of Fig. 3: Since there is a path from the bottom-left white vertex (shown as a fat squared node that represents the border root component consisting only of this vertex) to the border root component consisting of the red and green square on the right edge of the outer triangle, consensus cannot be solved in one round.

### 4.2 Characterizing solvability via paths connecting incompatible border components

In this subsection, we will characterize the possible evolutions of a path that connects facets with incompatible border components in some protocol complex $\mathcal{P}_{r-1}$, for some $r \geq 1$, which may either break or may lead to a *lifted* path connecting incompatible border components in $\mathcal{P}_r$.

Consider two border facets $\hat{F}_x \neq \hat{F}_y$ taken from a set of $k \geq 2$ incompatible proper border facets $\hat{F}_1, \ldots, \hat{F}_k \in \mathcal{C} \subseteq \mathcal{P}_{r-1}$, $r \geq 2$, i.e., belonging to the connected component $\mathcal{C}$ and having incompatible border components $B(\hat{F}_1), \ldots, B(\hat{F}_k)$ (see Fig. 5 for an illustration). Since $\mathcal{P}_{r-1}$ is the result of repeatedly applying $\mathcal{P}$ to the single facet $\sigma_0$, there must be some smallest round number $1 \leq \overline{r} \leq r-1$ and two facets $\overline{\tau}_x \neq \overline{\tau}_y$ with $B(\overline{\tau}_x) \neq B(\overline{\tau}_y)$ in $\mathcal{P}_{\overline{r}}$ that carry $\hat{F}_x$ and $\hat{F}_y$, respectively, i.e., $\hat{F}_x \in \mathcal{P}_{r-1-\overline{r}}(\overline{\tau}_x)$ and $\hat{F}_y \in \mathcal{P}_{r-1-\overline{r}}(\overline{\tau}_y)$. Note that, as $\overline{r}$ is minimal, $\overline{\tau}_x$ and $\overline{\tau}_y$ are facets obtained by applying $\mathcal{P}$ to the same facet $\overline{F} \in \mathcal{P}_{\overline{r}-1}$ (see Fig. 6). For simplicity of exposition, we will assume below that $\overline{r} = 1$, as otherwise we would have to introduce the definition of a "generalized border" that does not start from $\mathcal{P}_1$ but rather from $\mathcal{P}_{\overline{r}}$. We will hence subsequently just write $\mathcal{P}_1$, $\mathcal{P}_{r-1}$ and $\mathcal{P}_{r-2}$ instead of $\mathcal{P}_{\overline{r}}$, $\mathcal{P}_{r-\overline{r}}$, and $\mathcal{P}_{r-1-\overline{r}}$, respectively. Fortunately, this assumption can be made without loss of generality, as all the scenarios that can occur in the case of $\overline{r} > 1$ will also occur when $\overline{r} = 1$.

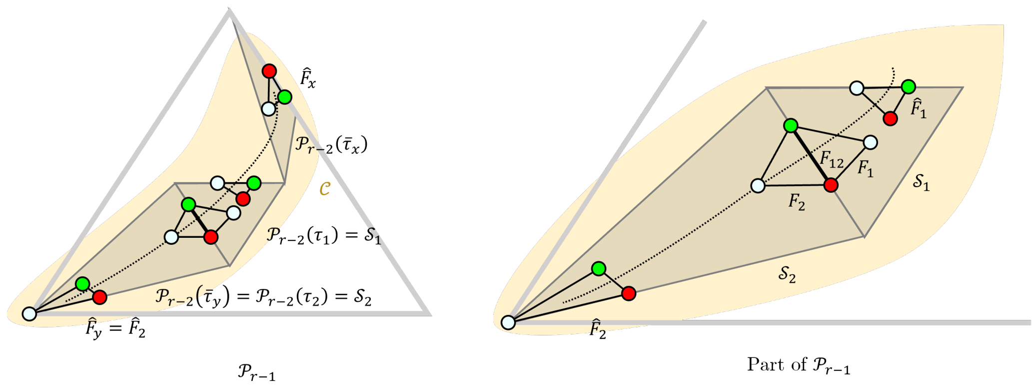


**Fig. 5.** $P_{r-1}$

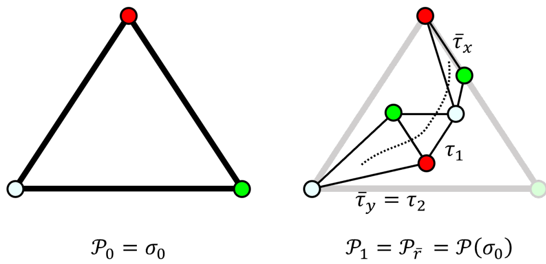


**Fig. 6.** $\mathcal{P}_0$ and $P_1$

Since $\hat{F}_x$ and $\hat{F}_y$ are connected in $\mathcal{C} \subseteq \mathcal{P}_{r-1}$, $\overline{\tau}_x$ and $\overline{\tau}_y$ must be connected via one or more paths of adjacent facets in $\mathcal{P}_1$ as well. Consider an arbitrary, fixed path connecting the proper border facets $\hat{F}_x$ and $\hat{F}_y$ in $\mathcal{P}_{r-1}$, and its unique corresponding path connecting $\overline{\tau}_x$ and $\overline{\tau}_y$ in $\mathcal{P}_1$. Let $\tau_1$ and $\tau_2$ be any two adjacent facets on the path in $\mathcal{P}_1$, and $\tau_{12} = \tau_1 \cap \tau_2 \neq \emptyset$. In $\mathcal{P}_{r-1}$, the facets $\tau_1$ and $\tau_2$ induce connected subcomplexes $\mathcal{S}_1 = \mathcal{P}^{r-2}(\tau_1)$ and $\mathcal{S}_2 = \mathcal{P}^{r-2}(\tau_2)$ with a non-empty intersection $\mathcal{S}_1 \cap \mathcal{S}_2 \neq \emptyset$. The path from $\hat{F}_x$ to $\hat{F}_y$ in $\mathcal{P}_{r-1}$ must enter $\mathcal{S}_1$ at some facet $\hat{F}_1$ and exit $\mathcal{S}_2$ at some facet $\hat{F}_2$, i.e., there is a path connecting $\hat{F}_x$ to $\hat{F}_1$ and a path connecting $\hat{F}_y$ to $\hat{F}_2$, and cross the border between $\mathcal{S}_1$ and $\mathcal{S}_2$ via adjacent facets $F_1 \in \mathcal{S}_1 \subseteq \mathcal{P}_{r-1}$ and $F_2 \in \mathcal{S}_2 \subseteq \mathcal{P}_{r-1}$ with $\emptyset \neq F_{12} = F_1 \cap F_2 \in \mathcal{S}_1 \cap \mathcal{S}_2$. Note that $F_{12}$, as the intersection of two facets in the protocol complex $\mathcal{P}_{r-1}$, is of course a face.

Now consider the outcome of applying $\mathcal{P}$ again to all the facets in $\mathcal{P}_{r-1}$, which of course gives $\mathcal{P}_r$ (see Fig. 7). In particular, this results in the subcomplexes $\mathcal{S}'_1 = \mathcal{P}\big(\mathcal{P}^{r-2}(\tau_1)\big) = \mathcal{P}^{r-1}(\tau_1)$ and analogously $\mathcal{S}'_2 = \mathcal{P}^{r-1}(\tau_2)$, which may or may not have a non-empty intersection. We will be

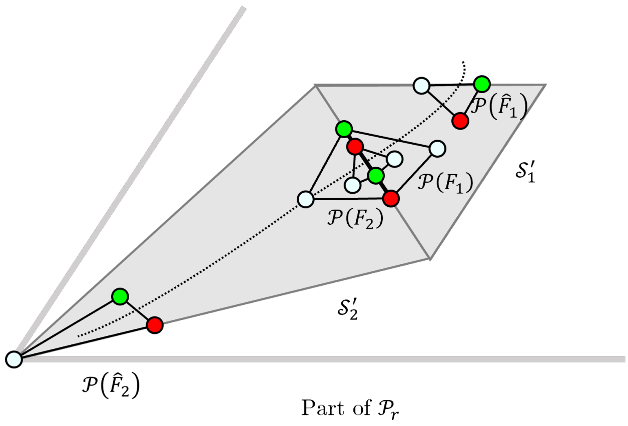


**Fig. 7.** Part of $\mathcal{P}_r$

interested in the part of this possible intersection created by the application of $\mathcal{P}$ to $F_1$ and $F_2$, i.e., in $\mathcal{P}(F_1) \cap \mathcal{P}(F_2) \subseteq \mathcal{S}'_1 \cap \mathcal{S}'_2$. Note that both $\mathcal{S}'_1$ and $\mathcal{S}'_2$ are isomorphic to $\mathcal{P}_{r-1}$. Clearly, the application of $\mathcal{P}$ to the facets $F_1, F_2 \in \mathcal{P}_{r-1}$ typically creates many pairs of intersecting border facets $F'_1 \in \mathcal{P}(F_1) \subseteq \mathcal{S}'_1 \subseteq \mathcal{P}_r$ and $F'_2 \in \mathcal{P}(F_2) \subseteq \mathcal{S}'_2 \subseteq \mathcal{P}_r$, such that each pair shares some non-empty face $\emptyset \neq F'_{12} = F'_1 \cap F'_2 \subseteq \mathcal{S}'_1 \cap \mathcal{S}'_2$. The shared faces $F'_{12}$ together form the subcomplex $FC'_{12} \subseteq \mathcal{P}(F_1) \cap \mathcal{P}(F_2)$ (see Fig. 8, left part, for two different examples).

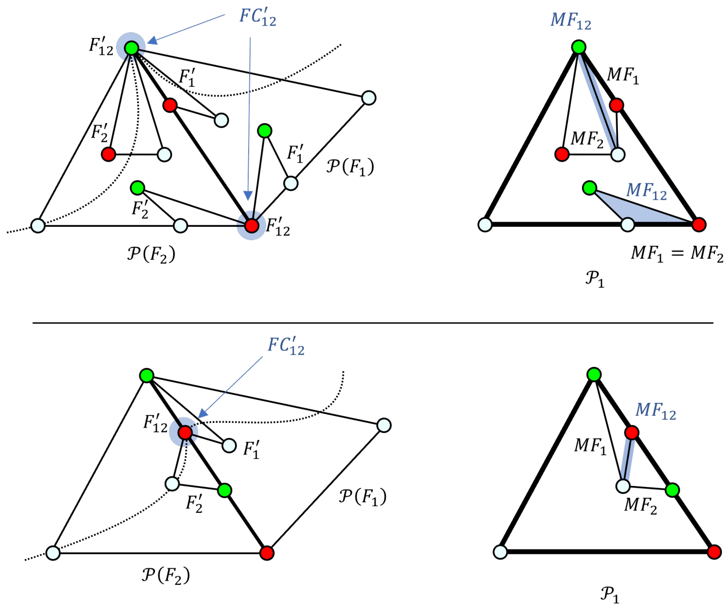


**Fig. 8.** The subcomplex $FC'_{12} \subseteq \mathcal{P}(F_1) \cap \mathcal{P}(F_2)$ in $\mathcal{P}_r$ and the corresponding subcomplex $MF_{12}$ in $\mathcal{P}_1$. Case (1) top and Case (2) bottom.

Any such $F'_{12}$ is not completely arbitrary, though: First of all, since $FC'_{12} \subseteq \mathcal{P}(F_1) \cap \mathcal{P}(F_2)$ implies that its colors can only be drawn from $\chi(F_{12})$ due to boundary consistency, we have

$$\chi(F'_{12}) \subseteq \chi(FC'_{12}) \subseteq \chi(F_{12}). \tag{17}$$

Moreover, every pair of properly intersecting facets $F'_1$ and $F'_2$ is actually created by two unique *matching border facets* $MF_1, MF_2 \in \mathcal{P}_1$: the adjacent facets $F'_1 \in \mathcal{P}(F_1)$ and $F'_2 \in \mathcal{P}(F_2)$ are isomorphic to some intersecting border facets $MF_1 \in \mathcal{P}_1$ and $MF_2 \in \mathcal{P}_1$, respectively, which match at the boundary (see Fig. 8, right part). This actually leaves only two possibilities for their intersection $MF_{12} = MF_1 \cap MF_2 \neq \emptyset$:

(1) $MF_1$ and $MF_2$ are proper border facets with the same root component $R(MF_1) = R(MF_2) \in MF_{12}$ (possibly $MF_{12} \setminus R(MF_1) \neq \emptyset$, though). Two instances are shown in the top part of Fig. 8.
(2) $MF_1$ and $MF_2$ are proper border facets with different root components, or improper border facets, with $R(MF_1) \cup R(MF_2) \not\subseteq MF_{12}$ (taken as a complex). An instance is shown in the bottom part of Fig. 8.

Note that these are all possibilities, since our single-rootedness assumption rules out $R(MF_1) \cup R(MF_2) \subseteq MF_{12}$: After all, every $v \in R(MF_1) \setminus R(MF_2) \neq \emptyset$ (w.l.o.g.) would need to have an outgoing path to every vertex in $R(MF_2)$, which is not allowed for the root component $R(MF_2)$ by Definition 18. Keep in mind, for case (2) below, that every vertex in $MF_{12}$ must have an incoming path from every member of $R(MF_1)$ in $MF_1$ and from every member of $R(MF_2)$ in $MF_2$.

Now, given any pair of facets $F_1'$ and $F_2'$, we will consider conditions ensuring the lifting/breaking of paths that run over their intersection $F_{12}'$. Not surprisingly, we will need to distinguish the two cases (1) and (2) introduced above. To support the detailed description of the different situations that can happen here, we recall the path in $\mathcal{C} \subseteq \mathcal{P}_{r-1}$ that forms our starting point (Fig. 5): It starts out from the proper border facet $\hat{F}_x$ and leads to $\hat{F}_1$, where it enters the subcomplex $\mathcal{S}_1 = \mathcal{P}^{r-2}(\tau_1)$. Within $\mathcal{S}_1$, the path continues to $F_1$. The latter has a non-empty intersection $F_{12}$ with $F_2$, which belongs to the subcomplex $\mathcal{S}_2 = \mathcal{P}^{r-2}(\tau_2)$. The path continues within $\mathcal{S}_2$ and exits it at $\hat{F}_2$, from where it finally leads to the proper border facet $\hat{F}_y$.

(0) If $F_{12}' = \emptyset$, i.e., the border facets $F_1'$ and $F_2'$ do not intersect at all, there obviously cannot be any path in $\mathcal{P}_r$ that connects $\hat{F}_1'$ and $\hat{F}_2'$ via $F_{12}'$.
(1) If $F_{12}' \neq \emptyset$ is caused by proper border facets $F_1'$ and $F_2' \in \mathcal{P}_r$ with the same border root component $R(F_1') = R(F_2')$, then both $F_1' \in \mathcal{P}_r$ and $MF_1 \in \mathcal{P}_1$ are isomorphic to some proper border facet $BF_1 \in \mathcal{P}_{r-1}$; analogously, $F_2' \in \mathcal{P}_r$ and $MF_2 \in \mathcal{P}_1$ are both isomorphic to some proper border facet $BF_2 \in \mathcal{P}_{r-1}$. This holds since $\mathcal{S}_1' = \mathcal{P}^{r-1}(\tau_1)$ and $\mathcal{S}_2' = \mathcal{P}^{r-1}(\tau_2)$ are both isomorphic to $\mathcal{P}_{r-1}$. Note carefully, though, that $BF_{12} = BF_1 \cap BF_2$ is isomorphic to $F_{12}'$ (but not necessarily to $MF_{12}$, as can be seen in the top of Fig. 8), so $\chi(R(BF_1)) = \chi(R(BF_2)) \subseteq \chi(BF_{12}) \subseteq \chi(F_{12}') \subseteq \chi(F_{12})$ by Eq. (17). Note that Definition 20 immediately implies $B(BF_1) = B(BF_2) \subseteq BF_{12}$ for the respective border components as well.
   Depending on $BF_1$ and $BF_2$ (actually, depending on $BF_{12}$ and, ultimately, on $R(BF_1) = R(BF_2)$, which we will say to *protect* $F_{12}$), all paths running over $R(F_1') = R(F_2')$ will either (a) be lifted or else (b) break:
   (a) We say that $R(BF_1) = R(BF_2)$ *successfully* protects $F_{12}$ (see Fig. 9) if $BF_{12} \in \mathcal{C}$, i.e., if both $BF_1$ and $BF_2$ are within the same connected component $\mathcal{C}$ as $F_1$ and $F_2$ (which also implies that there are paths in $\mathcal{C}$ connecting $\hat{F}_1$ to $BF_1$ and $\hat{F}_2$ to $BF_2$). In this case, there is a lifted path in $\mathcal{P}_r$ connecting some border facets $\hat{F}_1' \in \mathcal{P}(\hat{F}_1)$ and $\hat{F}_2' \in \mathcal{P}(\hat{F}_2)$ via $BF_{12}$, carried by the one in $\mathcal{P}_{r-1}$ that connected the proper border facets $\hat{F}_1$ and $\hat{F}_2$ via $F_{12}$: It exists, since both $\mathcal{S}_1' = \mathcal{P}^{r-1}(\tau_1)$ and $\mathcal{S}_2' = \mathcal{P}^{r-1}(\tau_2)$ are isomorphic to $\mathcal{P}_{r-1}$. By applying $\mathcal{P}$ to all remaining facets on the path that connected $\hat{F}_x$ and $\hat{F}_y$ in $\mathcal{P}_{r-1}$ as well, a path in $\mathcal{P}_r$ may be created that connects some incompatible proper border facets $\hat{F}_x', \hat{F}_y' \in \mathcal{P}_r$; of course, this requires a successful lifting everywhere along the original path, not just at the intersection between $F_1$ and $F_2$.
   (b) We say that $R(BF_1) = R(BF_2)$ *unsuccessfully* protects $F_{12}$ if $BF_{12} \notin \mathcal{C}$. In this case, there cannot be a path connecting the border facets $\hat{F}_1'$ and $\hat{F}_2'$ in $\mathcal{P}_r$ running via $BF_{12}$, i.e., the connecting path in $\mathcal{P}_{r-1}$ cannot be lifted to $\mathcal{P}_r$ and thus breaks!
(2) If $F_{12}' \neq \emptyset$ is not caused by proper border facets $F_1' \in \mathcal{P}_r$ and $F_2' \in \mathcal{P}_r$ with common border root components $R(F_1') = R(F_2')$, we know from (2) above that $R(MF_1) \cup R(MF_2) \not\subseteq MF_{12}$, i.e., at least one of the root components, say, $R(F_1')$, has a vertex $v_1' \in F_1'$ outside $F_{12}'$. Clearly, the corresponding vertex $v_2' \in F_2'$ with $\chi(v_1') = \chi(v_2')$ is also outside $F_{12}'$ and hence different from $v_1'$. Since there is a path from $v_1'$ to every vertex in $F_1'$, at least one member in the intersection $F_{12}$ will be gone in $F_{12}'$, so
$$|V(F_{12}')| < |V(F_{12})|. \tag{18}$$

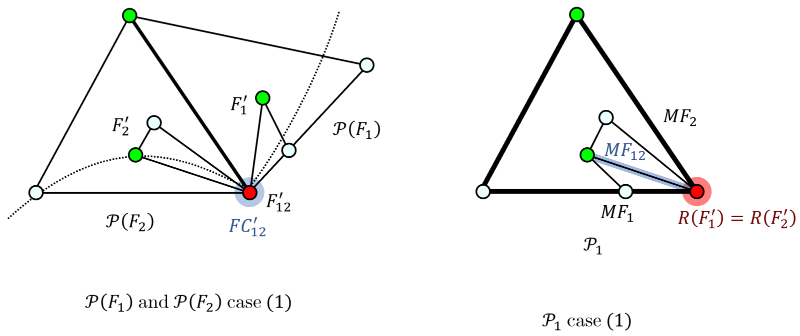


$\mathcal{P}(F_1)$ and $\mathcal{P}(F_2)$ in $\mathcal{P}_r$, and the corresponding subcomplexes $MF_1$ and $MF_2$ in $\mathcal{P}_1$

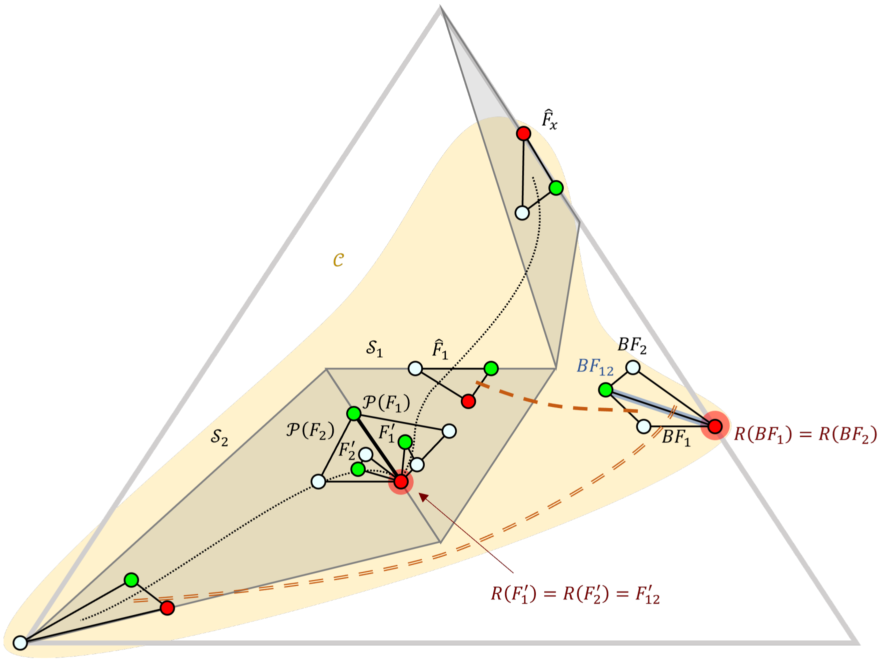


$\mathcal{P}_{r-1}$

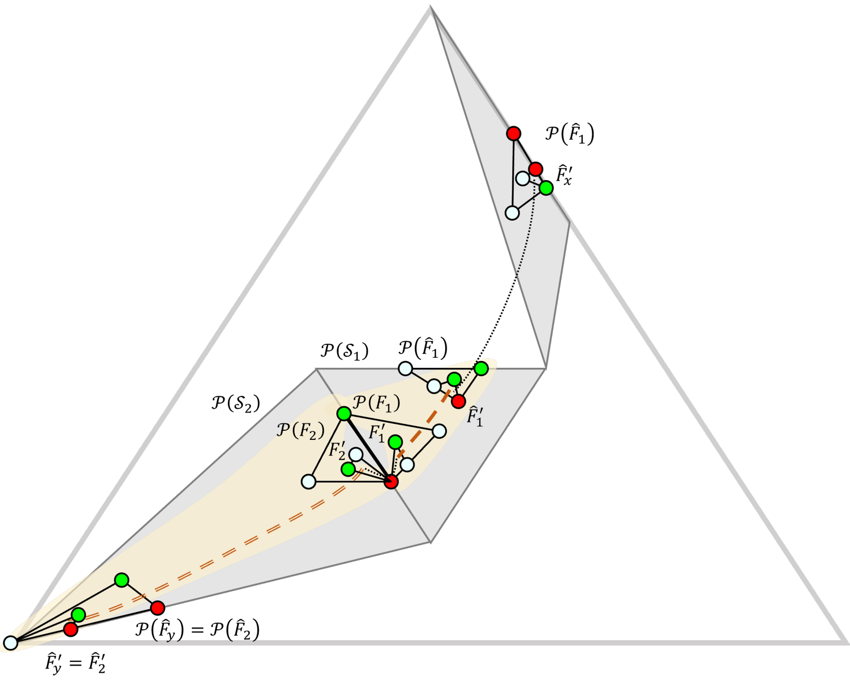


$\mathcal{P}_r$ with the images of $\mathcal{C}$ in $\mathcal{S}_1$ and $\mathcal{S}_2$ marked

**Fig. 9.** Lifting — Case 1(a)

This completes the (exhaustive) list of cases that need to be considered w.r.t. a single pair of facets $F_1'$ and $F_2'$. Clearly, in order for a lifted path connecting $\hat{F}_x'$ and $\hat{F}_y'$ in $\mathcal{P}_r$ to break, it suffices that it breaks for just one pair of adjacent facets. On the other hand, for a given pair $F_1, F_2 \in \mathcal{P}_{r-1}$, many paths are potentially created simultaneously in $\mathcal{P}_r$, each corresponding to a possible selection of $F_1'$ and $F_2'$ and the particular intersection facet $F_{12}'$, which all need to break eventually. Moreover, there are different paths in $\mathcal{P}_{r-1}$ connecting $\hat{F}_x$ and $\hat{F}_y$ via different pairs $F_1, F_2$ that need to be considered. In Section 5, we will show that there is even another subtle complication caused by case (1.b), the case where there is *no* lifted path in $\mathcal{P}_r$: It will turn out that "bypassing" may create a *new* path connecting some incompatible proper border facets in $\mathcal{P}_r$ when the path in $\mathcal{P}_{r-1}$ breaks; see Fig. 16 for an example.

Finally, for consensus to be solvable, *no* connected component $\mathcal{C}$ containing all border facets $\hat{F}_1, \dots, \hat{F}_k \in \mathcal{C}$ with incompatible border components $\hat{B}_1, \dots, \hat{B}_k$ may exist. In other words, there must be some $r$ such that *none* of the connected components of $\mathcal{P}_r$ contains facets with incompatible border components. If this is ensured, the processes in any facet $F \in C \subseteq \mathcal{P}_r$ can eventually decide on the initial value of a deterministically chosen process in $\bigcap_{F \in \mathcal{C}, B(F) \neq \emptyset} \beta(F) \neq \emptyset$. Note carefully, however, that this also requires that all connections between incompatible borders that are caused by facet borders different from the border component $B(F)$ have disappeared. Since this is solely a matter of case (2), Eq. (18) reveals that this must happen after at most $n-1$ additional rounds.

### 4.3 Characterizing consensus solvability via connected components

It is enlightening to view cases (0)–(2) introduced before w.r.t. the effect that they cause on the connected component $\mathcal{C} \subseteq \mathcal{P}_{r-1}$ that connects incompatible border facets: Reconsider the two adjacent facets $F_1, F_2 \in \mathcal{C}$ with intersection $F_{12}$, and assume, for clarity of the exposition, that $\mathcal{C}$ would fall apart if the path running over $F_{12}$ would break. We will now discuss what happens w.r.t. the connected component(s) in $\mathcal{P}_r$ when going to $F_1' \in \mathcal{P}(F_1)$ and $F_2' \in \mathcal{P}(F_2)$, under the assumption that $F_{12}'$ is the only intersecting facet in $\mathcal{P}_r$, according to our three cases:

(0) If $F_{12}' = \emptyset$, then $\mathcal{P}_r$ would contain two separate connected components $\mathcal{C}_1'$ and $\mathcal{C}_2'$ with $F_1' \in \mathcal{C}_1'$ and $F_2' \in \mathcal{C}_2'$, i.e., the connected component(s) in $\mathcal{P}_r$ resulting from $\mathcal{C}$ are separated by what is generated from $F_{12}$, namely, $F_{12}' = \emptyset$. A nice example is the (shaded) green center node in $\mathcal{P}_2^{RAS}$ of the RAS message adversary shown in Fig. 3.

(1) If $F_{12}' \neq \emptyset$ is caused by proper border facets $F_1' \in \mathcal{P}_r$ and $F_2' \in \mathcal{P}_r$ (isomorphic to $BF_1$ resp. $BF_2$) with the same border root component $R(F_1') = R(F_2') = F_{12}'$ that successfully protects $F_{12}$, we have our two subcases:

   (a) If $BF_{12} \in \mathcal{C}$, then $\mathcal{P}_r$ would contain a single connected component $\mathcal{C}'$ (resulting from $\mathcal{C}$) that connects incompatible border facets. An instance of such a successful protection can be found in Fig. 10 of $\mathcal{P}_2^{iRAS}$ for the iRAS message adversary, where consensus is impossible. Note that just one communication graph has been added to RAS here, namely, the additional triangle that connects the bottom left white vertex to the central triangle in the 1-round uninterpreted complex $\mathcal{P}_1^{iRAS}$ in the top left part of Fig. 10. Consider the left border of the dash-dotted central triangle, for example, where two adjacent facets intersect in the common root $r \leftrightarrow g$. It results from the fact that the border root component $r \leftrightarrow g$ of the proper border facet on the right outermost border of the protocol complex $\mathcal{P}_1^{iRAS}$ protects the intersection $r \leftrightarrow g$ of the central facet and the facet left to it in $\mathcal{P}_1^{iRAS}$.

   (b) If $BF_{12} \notin \mathcal{C}$, then $\mathcal{P}_r$ would contain two connected components $\mathcal{C}_1'$ and $\mathcal{C}_2'$ with $F_1' \in \mathcal{C}_1'$ and $F_2' \in \mathcal{C}_2'$. Unlike in case (0), however, they are separated by a third connected component $\mathcal{C}_{12}'$ that contains $F_1'$ and $F_2'$. It can be viewed as an "island" that develops around $F_{12}'$. A nice example of such an unsuccessful protection is the connected component containing the red process in the center of Fig. 11 for $\mathcal{P}_2^{2C}$ for the two-chain message adversary, which now separates the single connected component containing this process in $\mathcal{P}_1^{2C}$.
   We note that the bypassing effect already mentioned (and discussed in detail in Section 5) can also be easily explained via this view: It could happen that the "island" $C_{12}'$ is such that it connects some *other* incompatible proper border facets in $\mathcal{P}_r$ (see Fig. 16 for an example). So whereas it successfully separates $\mathcal{C}_1'$ and $\mathcal{C}_2'$, it creates a new path that prohibits the termination of consensus in round $r$.

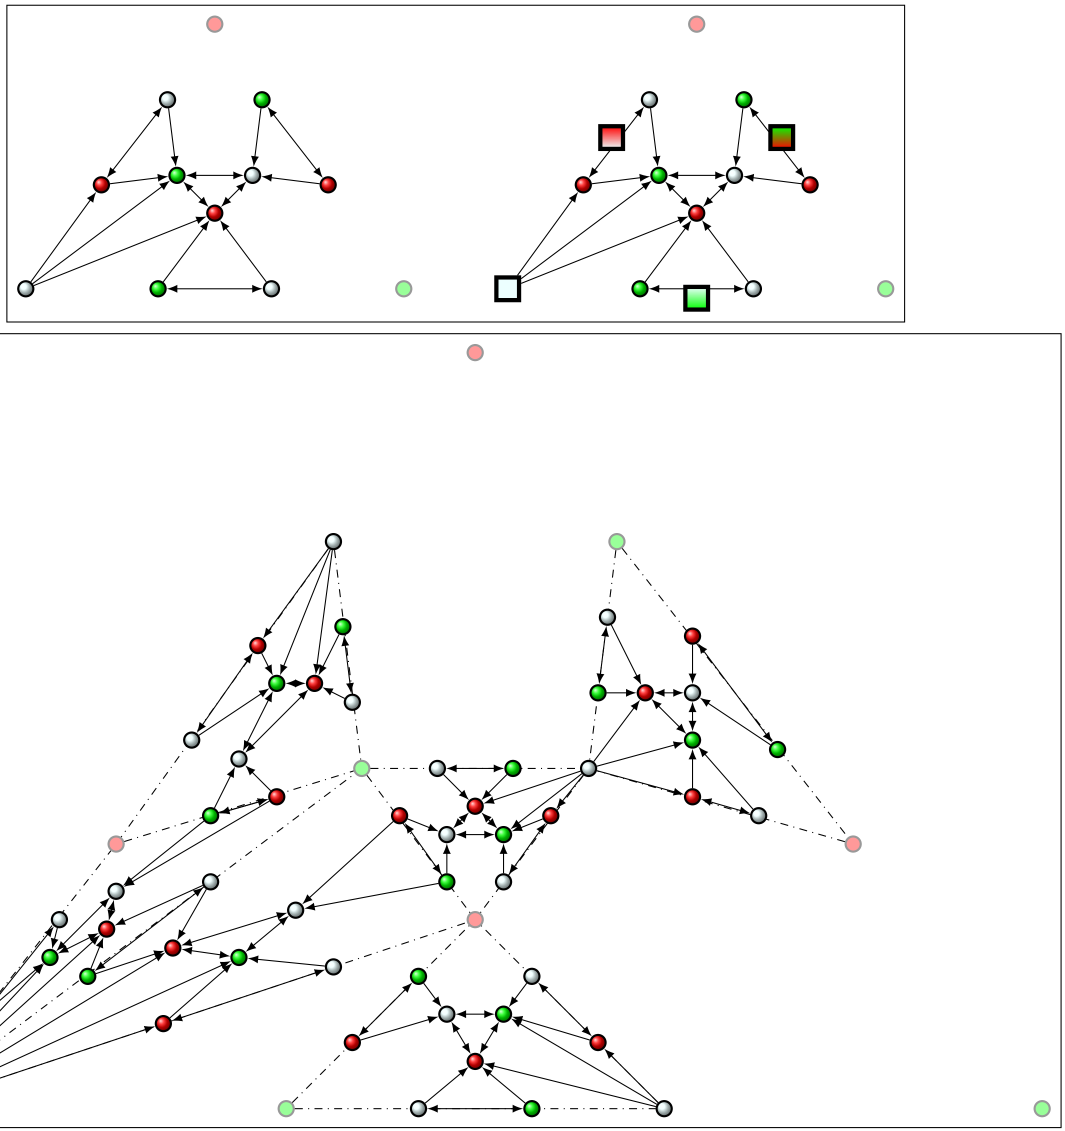

**Fig. 10.** Protocol complex for one round ($\mathcal{P} = \mathcal{P}_1^{iRAS}$, top) and two rounds ($\mathcal{P}_2^{iRAS}$, bottom) of the iRAS message adversary. The top right figure also shows the border root components of $\mathcal{P}$.

(2) If $F'_{12} \neq \emptyset$ is not caused by a proper border facet $F'_1 \in \mathcal{P}_r$ and $F'_2 \in \mathcal{P}_r$ with common border root component $R(F'_1) = R(F'_2) = F'_{12}$, $\mathcal{P}_r$ would contain a single connected component $\mathcal{C}'$ (resulting from $\mathcal{C}$) that still connects incompatible border facets.

### 4.4 Examples

An example of an unsuccessful protection (a breaking path, case (1.b)) can be found in the 1-round uninterpreted complex $\mathcal{P}_1^{RAS}$ for the RAS message adversary in the top-right part of Fig. 3, where the facets $\hat{F}_x$ and $\hat{F}_y$ containing the border root components $\hat{R}_x$ (the single white vertex in the bottom-left corner) and $\hat{R}_y$ (the bidirectionally connected red and green vertices on the right border) are connected by a path that runs over the bottom leftmost triangle $F_1 = \hat{F}_x$ and the central triangle $F_2$, in a joint connected component $\mathcal{C} \subseteq \mathcal{P}_1^{RAS}$. Note that $F_1$ and $F_2$ intersect in the single green central vertex $F_{12} = F_1 \cap F_2 = \{g\}$, and that there is no facet with a border (root) component consisting only of the green vertex in $\mathcal{P}_1$ and hence in $\mathcal{C}$. Consequently, it follows that $\mathcal{P}_2^{RAS}$ cannot contain a corresponding path connecting $\hat{F}'_x$ with border root component $\hat{R}'_x$ (the single white vertex in the bottom-left corner in the bottom part of Fig. 3) and $\hat{R}'_y$ (the bidirectionally connected red and green vertices on the right outer border) running over the faded central vertex $F'_{12} = \{g\}$, as is confirmed by our figure.

For an example of a successful protection (a non-breaking path, case (1.a)), consider the path connecting the border facets $\hat{F}_x$ and $\hat{F}_z$ containing the border root components $\hat{R}_x$ (the single white vertex in the bottom-left corner) and $\hat{R}_z$ (the bidirectionally connected red and white vertices on the left border) in the top-right part of Fig. 3. This path only consists of the bottom leftmost triangle $F_1 = \hat{F}_x$ and the triangle $F_2 = \hat{F}_z$ in a joint connected component $\mathcal{C} \subseteq \mathcal{P}_1^{RAS}$. Note that $F_1$ and $F_2$ intersect in a red-green edge $F_{12} = F_1 \cap F_2 = \{r \to g\}$ here, and that there is the border facet $\hat{F}_y \in \mathcal{C}$ with a border root component $\hat{R}_y = \{g \leftrightarrow r\}$ on the rightmost outer border. According to our considerations above, $\mathcal{P}_2^{RAS}$ contains a corresponding path connecting $\hat{F}'_x$ with border root component $\hat{R}'_x$ (the single white vertex in the bottom-left corner in the bottom part of Fig. 3) and the border facet $\hat{F}'_z$ with border root component $\hat{R}'_z$ (the bidirectionally connected red and white vertices on the leftmost outer border) running via $F'_{12} = \{g \leftrightarrow r\}$, as is confirmed by our figure.

To further illustrate the issue of successful/unsuccessful protection, consider the modified RAS message adversary iRAS depicted in Fig. 10, where consensus is impossible. The border facets $\hat{F}_w$ (the additional triangle) resp. $\hat{F}_y$ containing the border root component $\hat{R}_w$ (the single white vertex in the bottom-left corner) resp. $\hat{R}_y$ (the bidirectionally connected red and green vertices on the right border) are connected by a path that runs over the central bidirectional red-green edge $F_{12} = F_1 \cap F_2 = \{g \leftrightarrow r\}$ in $\mathcal{C}$ here. In sharp contrast to RAS, the border facet $\hat{F}_y$ with the border root component $\hat{R}_y = \{g \leftrightarrow r\}$ on the right outer border is now also in $\mathcal{C}$, however. Consequently, $\mathcal{P}_2^{RAS}$ contains a corresponding path connecting $\hat{F}'_w$ with border root component $\hat{R}'_w$ (the single white vertex in the bottom-left corner in the bottom part of Fig. 10) and $\hat{F}'_y$ with border root component $\hat{R}'_y$ (the bidirectionally connected red and green vertices on the right outer border) running via $F'_{12} = \{g \leftrightarrow r\}$, as is confirmed by our figure. Note that this situation recurs also in all further rounds, making consensus impossible.

To illustrate the issue of delayed path breaking (case (2)), consider another message adversary, called the *2-chain message adversary* (2C), shown for $n = 4$ processes in Fig. 11 (top part). It consists of three graphs, a chain $G_1 = y \to g \to w \to r$, another chain $G_2 = g \to y \to w \to r$, and a star $G_3 = r \to \{y, w, g\}$. In $\mathcal{P}_1$, the facets $F_1$ and $F_2$, corresponding to $G_1$ and $G_2$, respectively, are connected by a path running over the intersection $F_{12} = \{r\}$ in a joint connected component $\mathcal{C}$. There is also a border root component $R = \{r\}$ in the facet $F_3$ resulting from $G_3$, which, however, lies in a *different* connected component $\mathcal{C}' \neq \mathcal{C}$ in $\mathcal{P}$. According to our considerations (case (1.b), the path (potentially) connecting $F'_1$ (the border facet representing $G_1$ both in round 1 and 2) and $F'_2$ (the border facet representing $G_2$ both in round 1 and 2) via $F'_{12} = \{r\}$ in $\mathcal{P}_2^{2C}$ breaks: As is apparent from the bottom part of Fig. 11, there is no single red vertex shared by these two facets.

If one adds another process $p$ (pink) to 2C for $n = 5$, denoted by the message adversary 2C+, such that $G_1 = y \to g \to w \to p \to r$, $G_2 = g \to y \to w \to p \to r$, and $G_3 = r \to \{y, w, g, p\}$, then $F_{12} = \{p \to r\}$ is in $\mathcal{P}_1^{2C+}$. Now there is a path in $\mathcal{P}_2^{2C+}$ connecting $F'_1$ (the border facet representing $G_1$ both in round 1 and 2) and $F'_2$ (the border facet representing $G_2$ both in round 1 and 2) running via $F'_{12} = F'_1 \cap F'_2 = \{r\}$: Whereas the pink vertex has learned in round 2 where it belongs to, i.e., either $F'_1$ and $F'_2$, from the respective root component, this is not (yet) the case for the red vertex. However, whereas the corresponding path did not break in $\mathcal{P}_2^{2C+}$, it will finally break in $\mathcal{P}_3^{2C}$ since the red vertex will also learn where it belongs to.

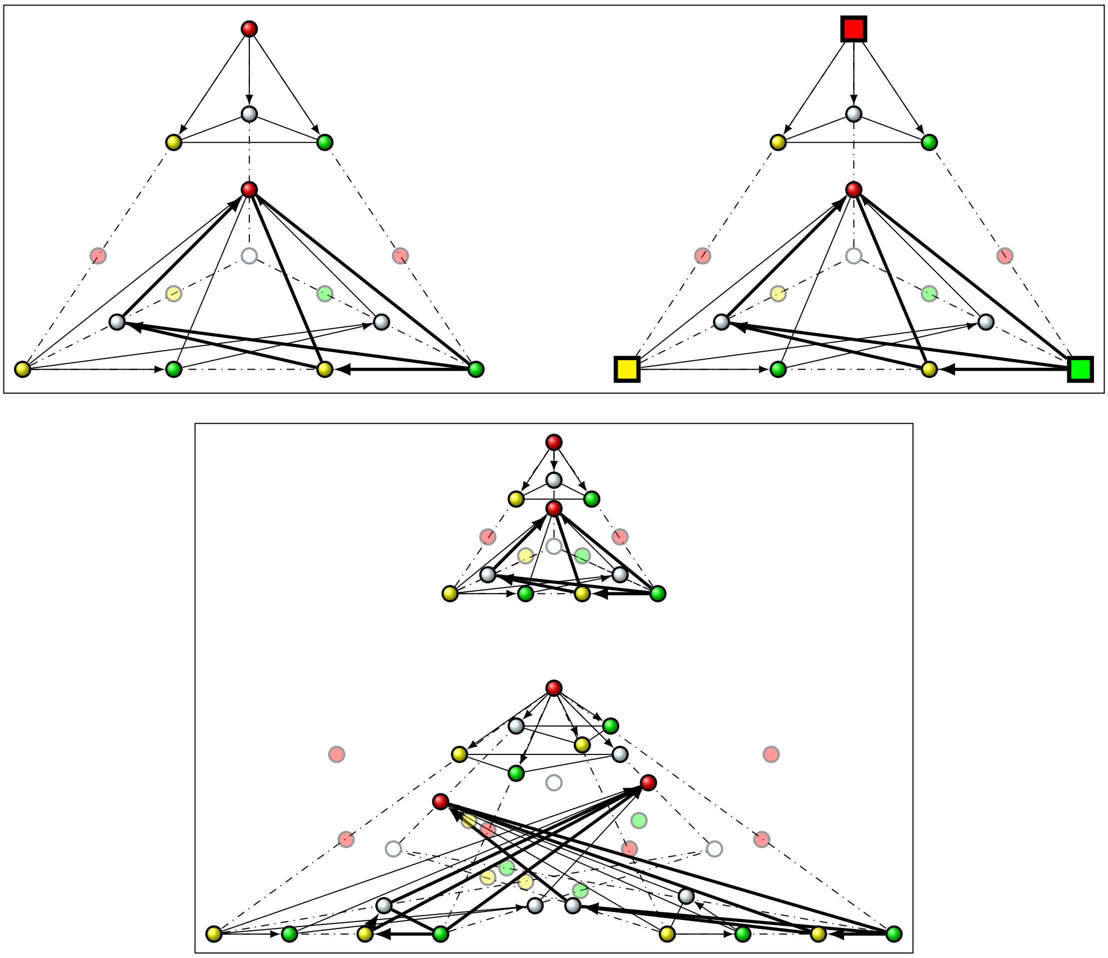

**Fig. 11.** Protocol complex for one round ($\mathcal{P} = \mathcal{P}_1^{2C}$, top) and two rounds ($\mathcal{P}_2^{2C}$, bottom) of the two-chain message adversary for $n = 4$ processes. The top right figure also shows the border root components of $\mathcal{P}$.

### 4.5 A decision procedure for consensus solvability

Revisiting the different cases (0)–(2) that can occur w.r.t. lifting/breaking a path connecting incompatible border facets in $\mathcal{P}_{r-1}$ to $\mathcal{P}_r$, it is apparent that the only case that might lead to a path that never breaks, i.e., in no round $r \geq 1$, is case (1.a): In case (0) and (1.b), there cannot be a lifted path running via $F'_{12}$ in $\mathcal{P}_r$, i.e., the path in $\mathcal{P}_{r-1}$ breaks immediately. In case (2), it follows from Eq. (18) that this type of lifting could re-occur in at most $n-2$ consecutive rounds after a path running over $F_{12}$ is lifted to a path running over $F'_{12}$ in $\mathcal{P}_r$ for the first time. Since these are all possibilities, after the "exhaustion" of case (2), $F'_{12} = \emptyset$ and hence case (0) necessarily applies.

In order to decide whether consensus is solvable for a given message adversary $\mathbf{D}^\omega$ at all, it hence suffices to keep track of case (1.a) over rounds $1, 2, \ldots$. If one finds that case (1.a) does not occur for any path in $\mathcal{P}_{r-1}$ for some $r$, there is no need for iterating further. On the other hand, if one finds that case (1.a) re-occurs for some path forever, consensus is impossible. Since the facets $BF_1$ and $BF_2$, where the common root component $R(BF_1) = R(BF_2)$ successfully protects $F_{12}$ in case (1.a), leads to $\chi(F'_{12}) \subseteq \chi(F_{12})$ according to Eq. (17), the infinite re-occurence of case (1.a) for some path implies that there is some round $r_0$ such that $\chi(F'_{12}) = \chi(F_{12}) = I \subset \Pi$ for all $r \geq r_0$. If this holds true, then case (1.a) must also re-occur perpetually in the lifted paths obtained by using the same $BF_1$ and $BF_2$ with $\chi(BF_{12}) = I$ in all rounds $r \geq r_0$.

For keeping track of possibly infinite re-occurrences of case (1.a), it is hence sufficient to determine, for every pair of facets $MF_1 \in \mathcal{P}_1$ and $MF_2 \in \mathcal{P}_1$, $MF_2 \neq MF_1$, intersecting in $MF_{12} \neq \emptyset$, the set of proper border facets $MF^1, \dots, MF^\ell \in \mathcal{P}_1$ with border root components $R^j = R(MF^j)$ satisfying $\chi(R^j) \subseteq \chi(MF_{12})$ for all $1 \leq j \leq \ell$. Clearly, every choice of $MF^j$, $MF^k$ is a possible candidate for the isomorphic re-occurring protecting facets $BF_1 = BF^j$ and $BF_2 = BF^k$ for case (1.a) in some $\mathcal{P}_{r-1}$, provided (i) $R(MF^j) = R(MF^k)$ and (ii) both $BF^j$ and $BF^k$ are in the connected component $\mathcal{C} \subseteq \mathcal{P}_{r-1}$ containing $F_1$ and $F_2$. If (ii) does not hold, one can safely drop $MF^j$, $MF^k$ from the set of candidates for infinitely re-occurring protecting facets in all subsequent rounds.

This can be operationalized in an elegant and efficient decision procedure by using an appropriately labeled and weighted version of the facets' *nerve graph* $\mathcal{N}$ of the 1-round uninterpreted complex $\mathcal{P}_1$. It is a topological version of the combinatorial decision procedure given in [39, Alg. 1] that works as follows: Every facet in $\mathcal{P}_1$ is a node $F$ in $\mathcal{N}$ and labeled by $w(F) = R(F)$, its root component in $\mathcal{P}_1$. Two nodes $F^1, F^2$ in $\mathcal{N}$ are joined by an (undirected) edge $(F^1, F^2)$, if they intersect in a simplex $\emptyset \neq F^{12} = F^1 \cap F^2$ in $\mathcal{P}_1$. The edge is labeled by $w((F^1, F^2)) = \{R^1, \dots, R^\ell\}$ (possibly empty), which is the maximal set of (necessarily: border) root components that satisfy the property $\chi(F^{12}) \supseteq \chi(R^i)$. Recall that the member sets of different border root components may satisfy $\chi(R^i) \cap \chi(R^j) \neq \emptyset$ and even $\chi(R^i) = \chi(R^j)$, albeit $R^i \cap R^j \neq \emptyset$ when taken as faces is impossible.

The procedure for deciding on consensus solvability proceeds in iterations, starting from $\mathcal{N}_0 = \mathcal{N}$, and defining $\mathcal{N}_{i+1}$ from $\mathcal{N}_i$ as follows. Let $V(\mathcal{N}_{i+1}) := V(\mathcal{N}_i)$ with the same node labels $w(F)$, initialize $E(\mathcal{N}_{i+1})$ to be the empty set, and add to it each edge $(F^1, F^2) \in E(\mathcal{N}_i)$ with a label $w_{i+1}((F^1, F^2))$ defined next, but only if this label is not empty. For a potential edge $(F^1, F^2) \in E(\mathcal{N}_i)$, set $R \in w_{i+1}((F^1, F^2))$ if the (unique) connected component $C_i^j$ of $\mathcal{N}_i$ with $(F^1, F^2) \in E(\mathcal{C}_i^j)$ contains some $F' \in V(\mathcal{C}_i^j)$ with $w(F') = R \in w_i((F^1, F^2))$. The construction stops when either (i) none of the connected components of $\mathcal{N}_i$ contains nodes representing facets with incompatible root components (consensus is solvable), or else (ii) if $\mathcal{N}_{i+1} = \mathcal{N}_i$ but there is at least one connected component containing nodes representing facets with incompatible root components (consensus is impossible).

For example, Fig. 12 shows the labeling of the facets with their root components for the RAS message adversary, where consensus can be solved. The sequence of nerve graphs $\mathcal{N}$, $\mathcal{N}_0$ and $\mathcal{N}_1$ is illustrated in Fig. 13. On the other hand, Fig. 14 and Fig. 15 show the same for the iRAS message adversary, where consensus cannot be solved.

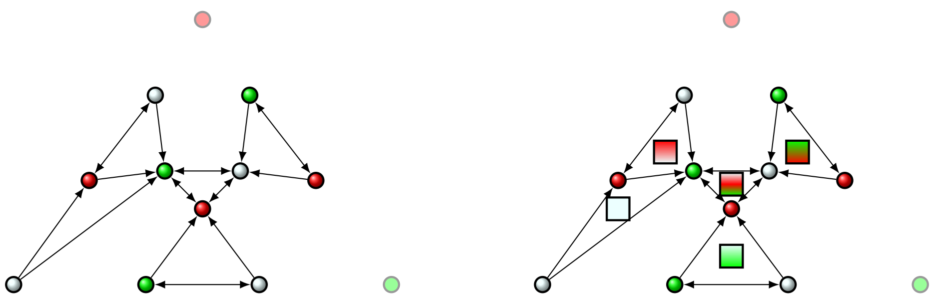

**Fig. 12.** Results of labeling the faces of the 1-layer protocol complex $\mathcal{P}_1$ of the RAS message adversary (left) by their root component (right).

Note that there is a small difference between the decision procedure [39, Alg. 1] and our topological version: Whereas the latter uses sets of border root components $w_i((F^1, F^2)) = \{R^1, \dots, R^\ell\}$ as the label of an edge $(F^1, F^2)$, the size of which may decrease during the iterations, the former uses the fixed set of processes that cannot distinguish $F^1$ and $F^2$ in $\mathcal{P}_1$ as its label $\ell((F^1, F^2))$. The latter does not change during the iterations, and can in fact be written as

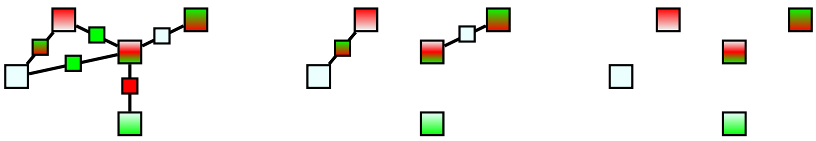

**Fig. 13.** Construction of the initial nerve graph $\mathcal{N}_0$ of the 1-layer protocol complex $\mathcal{P}_1$ of the RAS message adversary: After replacing the facets by their corresponding nodes (colored by their root component) and labeling all the edges (left), nerve graph $\mathcal{N}_0$ obtained by removing edges without a protecting root component (middle), nerve graph $\mathcal{N}_1$ (right). Note that $\mathcal{N}_1$ already reveals that consensus is solvable.

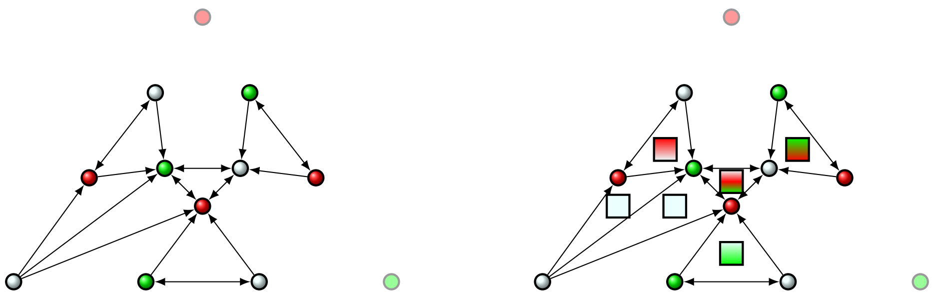

**Fig. 14.** Results of labeling the faces of the 1-layer protocol complex $\mathcal{P}_1$ of the iRAS message adversary (left) by their root component (right).

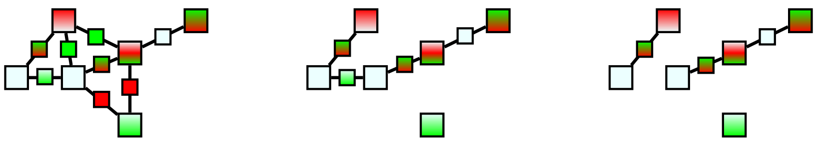

**Fig. 15.** Construction of the initial nerve graph $\mathcal{N}_0$ of the 1-layer protocol complex $\mathcal{P}_1$ of the iRAS message adversary: After replacing all facets by their corresponding nodes (colored by their root component) and labeling all the edges (left), nerve graph $\mathcal{N}_0$ obtained by removing edges without a protecting root component (middle), nerve graph $\mathcal{N}_1$ (right). Since $\mathcal{N}_2 = \mathcal{N}_1$, which still contains a component that connects incompatible root components, consensus is impossible.

$\ell((F^1, F^2)) = \chi(w_0((F^1, F^2))) = \chi(\{R^1, \ldots, R^\ell\})$. Whereas these different labeling schemes are equivalent in terms of correctly deciding consensus solvability/impossibility, ours might facilitate a more efficient data encoding and thus some advantages in computational complexity for certain message adversaries. On the other hand, we could of course also use the original labeling of [39, Alg. 1] in our decision procedure and detect successfully protecting border root components via proper inclusion of the member sets.

## 5 Consensus Termination Time

In this section, we will shift our attention from the principal question of whether consensus is solvable under a given message adversary $\mathbf{D}^\omega$ to the question of how long a distributed consensus algorithm may take to terminate.

Whereas it is immediately apparent that the number of iterations of the decision procedure in Section 4.5 is a lower bound for the consensus termination time, their exact relation is not clear: Case (2) of our classification for path lifting/breaking in Section 4 revealed an instance where the actual breaking of a path may happen up to $n-1$ rounds *after* detecting that it will eventually break. An interesting question is whether there are other effects that may even increase this gap between iteration complexity of the decision procedure and consensus termination time. And indeed, [39] provided an example that shows that this gap may even be exponential in $n$. In Section 5.1, we will provide an intuitive topological explanation of this gap, which is caused by the possibility of "bypassing". In Section 5.2, we finally propose a decision procedure, which allows to answer the question whether distributed consensus is solvable in $k$ rounds under a given message adversary $\mathbf{D}^\omega$.

### 5.1 Delayed path breaking due to bypassing

We mentioned already in Section 4.2 that in order for some incompatible border components to become disconnected, all paths connecting those must break. Consider the situation illustrated in Fig. 16, for the case of a system of $n = 5$ processes $(r, g, w, p, y)$, and a message adversary that comprises only 5 graphs, according to the uninterpreted 1-round protocol complex $\mathcal{P}_1$ illustrated in the top part of the figure. There are two paths $P_a \in \mathcal{C}_a$ and $P_b \in C_b$ in $\mathcal{P}_1$ that connect the same incompatible border facets (touching upon the $\{y, g\}$ resp. upon the $\{p, w\}$ border component carrier), lying in *different* connected components $\mathcal{C}_a$ and $\mathcal{C}_b$. The left path $P_a$ consists of facets $F_1$ (with $R(F_1) = \{y\}$) and $F_2$ (with $R(F_2) = \{w\}$), sharing the face $F_{12} = \{r\}$. The right path $P_b$ consists of $H_1$, $H_0$ and $H_2$ and involves the facet $H_0$ with border root component $R(H_0) = \{r\}$. According to case (1.b), both corresponding lifted paths in $\mathcal{P}_2$ break, since the shared faces $F_{12}$ between any two facets are (unsuccessfully) protected by the common root component of proper border facets in $\mathcal{P}_1$ lying in a *different* connected component only.

Surprisingly, however, the $\{y, g\}$ and $\{p, w\}$ borders themselves are *not* separated in $\mathcal{P}_2$. Actually, it happens that the right path $P_b$ in $\mathcal{C}_b \subseteq \mathcal{P}_1$ gives rise to a *new* lifted path connecting facets with proper border components in the $\{y, g\}$ resp. $\{p, w\}$ borders in $\mathcal{P}_2$. This effect, called *bypassing*, is illustrated in the bottom part of Fig. 16: By applying $\mathcal{P}$ to both $F_1, F_2$, leading to $\mathcal{P}(F_1), \mathcal{P}(F_2) \subseteq \mathcal{P}_2$, one observes that $P_b$ now leads to a new lifted path connecting the incompatible border (root) components $R(H_1') = \{g\}$ (in the facet $H_1'$ corresponding to $H_1$ in $\mathcal{P}(F_1)$) and $R(H_2'') = \{p, w\}$ (in the facet $H_2''$ corresponding to $H_2$ in $\mathcal{P}(F_2)$) via the intersection of $H_0' \cap H_0'' = \{r\}$. In fact, the island created in $\mathcal{P}_2$ around the latter, due to case (1.b), which consists of $H_1'$, $H_2'$, $H_0'$, $H_0''$, $H_1''$ and $H_2''$ and nicely separates the connected components consisting of $F_1'$ and $F_2'$ from $F_1''$ and $F_2''$, is not an island, but rather connects two other incompatible border root components, namely $\{g\} \in H_1'$ and $\{w\} \in H_2''$. Whereas it can be inferred already in $\mathcal{P}_1$ that this new lifted path in $\mathcal{P}_2$ will eventually break as well, consensus cannot be solved in just two rounds here.

Even worse, for larger $n$, it is possible to iterate this construction: An additional path $P_c$ in a separate connected component $\mathcal{C}_c \subseteq \mathcal{P}_1$ could bypass both the shared face $\{w\}$ between $H_1$ and $H_0$ and $\{y\}$ between $H_2$ and $H_0$ in $P_b$, in the same way as the shared face $\{r\}$ in $P_a$ is bypassed. More specifically, if these shared faces in $P_b$ are (unsuccessfully) protected by the border root components of proper border facets in $P_c \in \mathcal{C}_c$, which connect proper border facets touching the incompatible

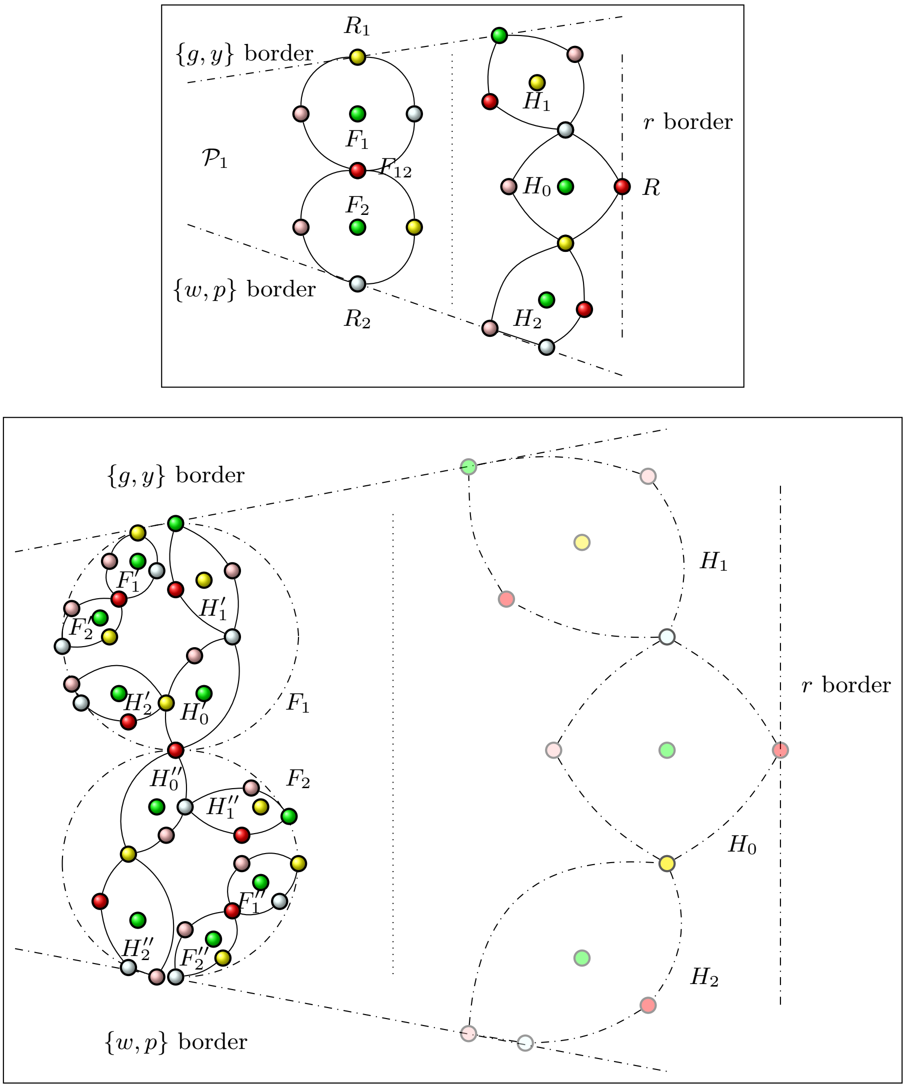


**Fig. 16.** Illustration of delayed path creation in the evolution of a protocol complex, for $n = 5$. The top part shows $\mathcal{P}_1$, which consists of two paths $P_a$ and $P_b$ connecting incompatible border (root) components (on the $\{y, g\}$ resp. $\{p, w\}$ border), lying in different connected components $\mathcal{C}_a$ and $\mathcal{C}_b$. The left path $P_a$ consists of facets $F_1$ and $F_2$, sharing $F_{12} = \{r\}$. The right path $P_b$ consists of $H_0$–$H_2$ and involves also a facet $H_0$ with border (root) component $R(H_0) = \{r\}$. Note that our restriction to $n = 5$ implies that the pink vertices in $F_1$ and in $F_2$ are actually the same, and so are the red vertices in $H_1$ and $H_2$. The bottom part shows $\mathcal{P}_2$: Whereas the corresponding paths for both $P_a$ and $P_b$ break in round 2 according to case (1.b), it also happens that $P_b$ creates a lifted path in $\mathcal{P}_2$ (running within $\mathcal{P}(F_1)$ and $\mathcal{P}(F_2)$) that connects "new" incompatible proper border components there. This lifted path will break only in $\mathcal{P}_3$.

$\{g, y\}$ border to the $\{w\}$ border, to the $\{r\}$ border, to the $\{y\}$ border, and finally to the $\{w, p\}$ border, one gets a path connecting the $\{g, y\}$ and the $\{w, p\}$ borders in $\mathcal{P}_2$, carried by $P_b$ in $\mathcal{P}_1$, in exactly the same way as we got the path carried by $P_a$ described above. Note carefully that the border root component of the proper border facet touching the $\{r\}$ border in $P_c$ must be different from the one touching the $\{r\}$ border in $P_b$, since $\mathcal{C}_b \neq \mathcal{C}_c$. Since $\mathcal{P}_3 = \mathcal{P}(\mathcal{P}_2)$, this finally causes the creation of a new path in $\mathcal{P}_3$ that also connects the incompatible $\{g, y\}$ and $\{w, p\}$ borders, carried by $P_a$ in $\mathcal{P}_1$.

Whereas successive bypassing cannot go on forever, it stops only if there are no "new" connected components in $\mathcal{P}_1$ that allow to bypass shared faces. Indeed, there are natural limits of the number of such bypassing connected components:

(1) The bypassing connected components $\mathcal{C}_1, \mathcal{C}_2, \dots$ in $\mathcal{P}_1$ must connect incompatible borders, but must be disjoint. The root components of the facets that touch some specific border in different components must hence all be different (when taken as faces) as well. However, the example worked out in [39] demonstrates that there can be exponentially (in $n$) many of those.
(2) A connected component $\mathcal{C}_{x+1}$ that contains a border root component $R$ that unsuccessfully protects a shared face $F_{12}$ in the connected component $\mathcal{C}_x$ to accomplish bypassing must be such that it connects both the incompatible borders of $\mathcal{C}_x$ and some border containing $R$. Since each such $F_{12}$ must be protected by some proper border facet in $\mathcal{C}_{x+1}$, the length of the paths connecting two particular borders must be strictly increasing.

This ultimately provides a very intuitive "geometric" explanation of the quite unexpected exponential blowup of the gap between the iteration complexity of the decision procedure and the termination time of distributed consensus. In particular, (1) explains the surprising fact that the *number* of connected components in $\mathcal{P}_1$ plays a major role here.

### 5.2 A decision procedure for $k$-round distributed consensus

Reviewing the decision procedure of Section 4.5 in the light of bypassing as described in Section 5.1, it is apparent that the nerve graph based approach removes edges/labels *eagerly*. Regarding decision time, this is of course most advantageous: In the example of Fig. 16, it would terminate already after one iteration, telling that consensus is solvable.

There is a less eager alternative decision procedure, which builds a sequence of *(border) root reachability graphs* $\mathcal{RRG}_i$, $i \geq 0$, that tell which proper border facets are reachable from each other in $\mathcal{P}_{i+1}$. First, it builds the initial root reachability graph $\mathcal{RRG}_0$, the vertices of which (represented as square nodes in our illustrating figures) are the border root components (which are the same as the border components for all proper border facets) of the 1-round uninterpreted complex $\mathcal{P}_1$, see Fig. 3 (top right), and where two such vertices are connected by an undirected edge if they are connected via a path in $\mathcal{P}_1$ (irrespectively of the type of edges in $\mathcal{P}_1$), see Fig. 17 (top left). We obtain $\mathcal{RRG}_1$ by replacing every facet $F$ in $\mathcal{P}_1$ by an instance of $\mathcal{RRG}_0$, in such a way that the replacements of two facets $F_1$, $F_2$ that intersect in a simplex $F_{12}$ (case (1.a) in Section 4.2) that is protected by the common root $R(MF_1) = R(MF_2)$ of the proper border facets $MF_1$ and $MF_2$ in $\mathcal{P}_1$, i.e., $\chi(R(MF_1)) = \chi(R(MF_2)) \subseteq \chi(F_{12})$, share a node labeled $\chi(R(MF_1)) = \chi(R(MF_2))$. Note that an actual root component is represented by a fat square node in our figures, whereas the node representing an intersection is displayed by a non-fat square node.

Fig. 17 shows $\mathcal{RRG}_0$ (top left part), obtained directly from the top-right part of Fig. 3, and $\mathcal{RRG}_1$ (top right part), which consists of several connected components. It is apparent, however, that it no longer connects incompatible border components. In particular, the bottom left border root component consisting of the white fat square node is no longer connected by a path to the red-green fat square node on the right side of the outer triangle in $\mathcal{RRG}_1$. That is, the connection between these two border root components (present in $\mathcal{RRG}_0$) has disappeared!

This immediately gives us a recursive procedure for deciding consensus solvability: Rather than starting from the initial $\mathcal{RRG}_0 = \mathcal{RRG}_0^{(0)}$, we start inductively from the previously constructed $\mathcal{RRG}_0^{(i)}$, $i \geq 0$, and plug it into $\mathcal{P}_1$ exactly as before to construct $\mathcal{RRG}_1^{(i+1)}$. Note that, for $i \geq 1$, $\mathcal{RRG}_1^{(i+1)}$ has at most the same number of edges than $\mathcal{RRG}_1^{(i)}$. This process can be repeated until $\mathcal{RRG}_1^{(m+1)} = \mathcal{RRG}_1^{(m)}$ for some $m \geq 0$. Consensus is possible if and only if $\mathcal{RRG}_1^{(m)}$ contains no component that connects incompatible fat square nodes. In the example of Fig. 17, already $\mathcal{RRG}_1^{(1)}$ does not connect incompatible fat square nodes, so consensus is solvable under RAS.

This recursive RRG construction is equivalent to the the following iterative procedure, which operates directly on the root reachability graphs: Starting out from $\mathcal{RRG}_i$, initially $\mathcal{RRG}_1$, construct $\mathcal{RRG}_{i+1}$ by removing every edge incident to a node (= a non-fat square node) where the common border root component (= the fat square node $R(BF_1) = R(BF_2)$) of the protecting matching border facets is not in the same connected component. The procedure stops if either the resulting $\mathcal{RRG}_i$ contains no component that connects incompatible fat square nodes (in which case consensus is solvable), or else if $\mathcal{RRG}_{i+1} = \mathcal{RRG}_i$ (in which case consensus is not solvable if incompatible fat square nodes are still connected).

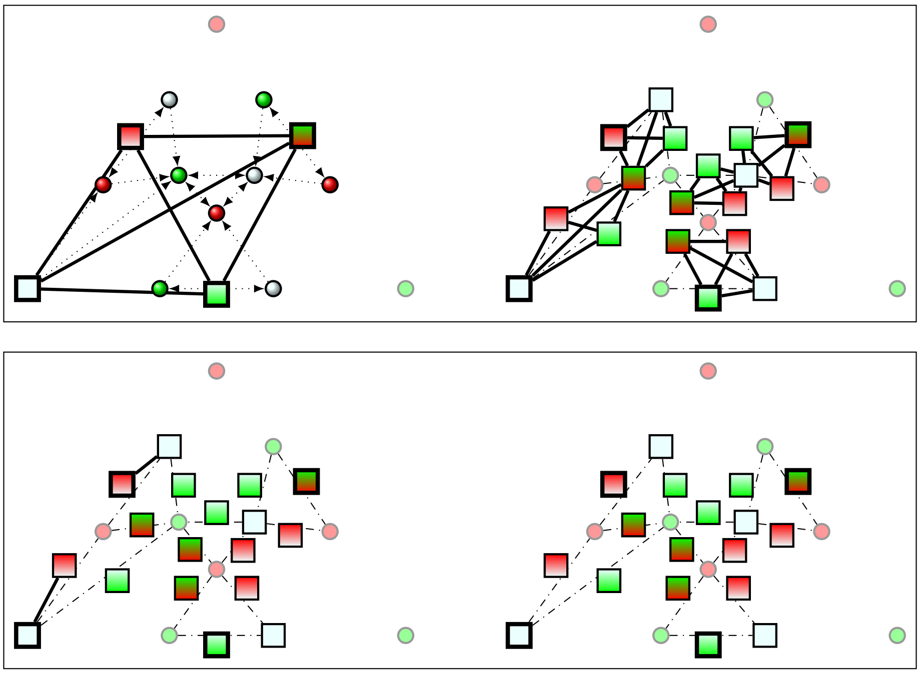

**Fig. 17.** Construction of the root components reachability graphs $\mathcal{RRG}_0$–$\mathcal{RRG}_3$ from the 1-layer protocol complex for RAS: Initial root reachability graph $\mathcal{RRG}_0$ (top left) and $\mathcal{RRG}_1$ (top right). Since $\mathcal{RRG}_1$ already partitions into several connected components that no longer contain incompatible border root components, one can already decide here that consensus is solvable. For completeness, we also show $\mathcal{RRG}_2$ (bottom left) and $\mathcal{RRG}_3$ (bottom right), where all edges have finally been removed.

Fig. 18 and Fig. 14 show the RRG construction for the iRAS message adversary, where consensus is impossible, as introduced in Fig. 10. The case of the 2-chain message adversary 2C is illustrated in in Fig. 19.

Returning to the example of Fig. 16, it is apparent that the root reachability graph based decision would not terminate as early as the nerve graph based procedure, since it explicitly keeps track of all paths between border root components. More specifically, whereas the path $P_a$ between the root components $R_1 = \{y\}$ and $R_2 = \{w\}$ in $\mathcal{P}_1$ in the top part of Fig. 16 has vanished in $\mathcal{P}_2$, and hence also in $\mathcal{RRG}_1$, the path $P_b$ connecting the border root components $\{g\}$ and $\{p, w\}$ is lifted to $\mathcal{P}_2$ and hence still present in $\mathcal{RRG}_1$. Consequently, the decision procedure would proceed to $\mathcal{RRG}_2$ before it can decide that consensus is solvable. In general, it would faithfully track paths/connected components that bypass each other until they have been exhausted.

It follows that the $\mathcal{RRG}$-based decision procedure would be a natural candidate for developing a decision procedure that can tell whether distributed consensus is solvable within $k$ rounds. However, like the nerve graph based procedure, it does not cover delayed path breaking due to case (2). Whereas a simple way to also accommodate this would be to scale the number of rounds required for termination by a factor of $n-1$, i.e., to infer from a number of iterations $k$ of the decision

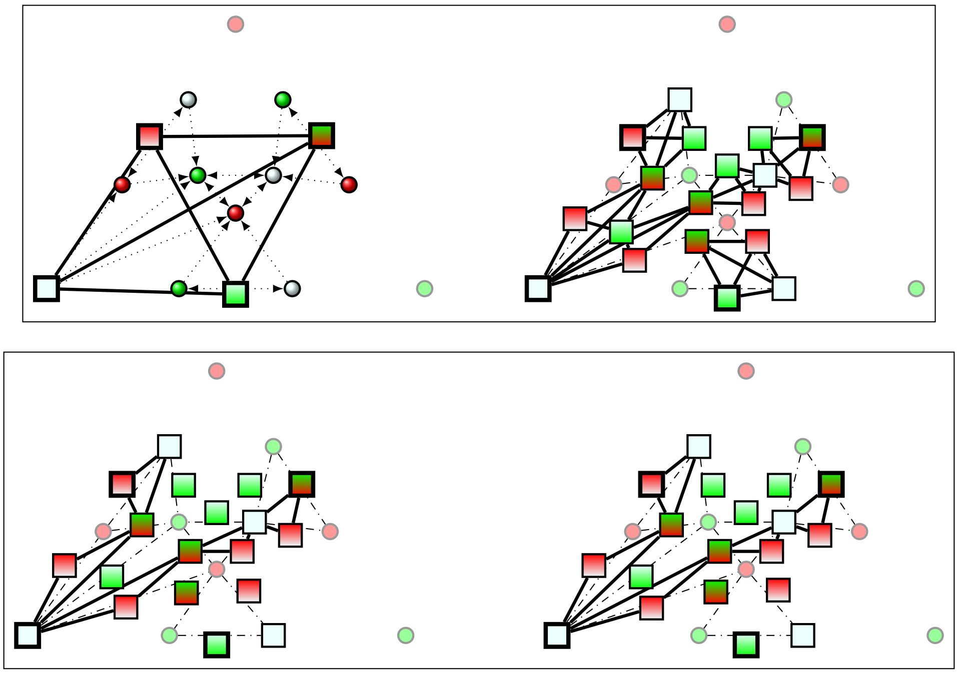

**Fig. 18.** Construction of the root components reachability graphs $\mathcal{RRG}_0$–$\mathcal{RRG}_3$ from the 1-layer protocol complex for the iRAS message adversary: Initial root reachability graph $\mathcal{RRG}_0$ (top left) and $\mathcal{RRG}_1$ (top right). Since $\mathcal{RRG}_1$ partitions into several connected components containing incompatible border root components, one has to construct $\mathcal{RRG}_2$ (bottom left). As incompatible border root compenents are still connected $\mathcal{RRG}_2$, another iteration finally provides $\mathcal{RRG}_3 = \mathcal{RRG}_2$, so consensus is not solvable here.

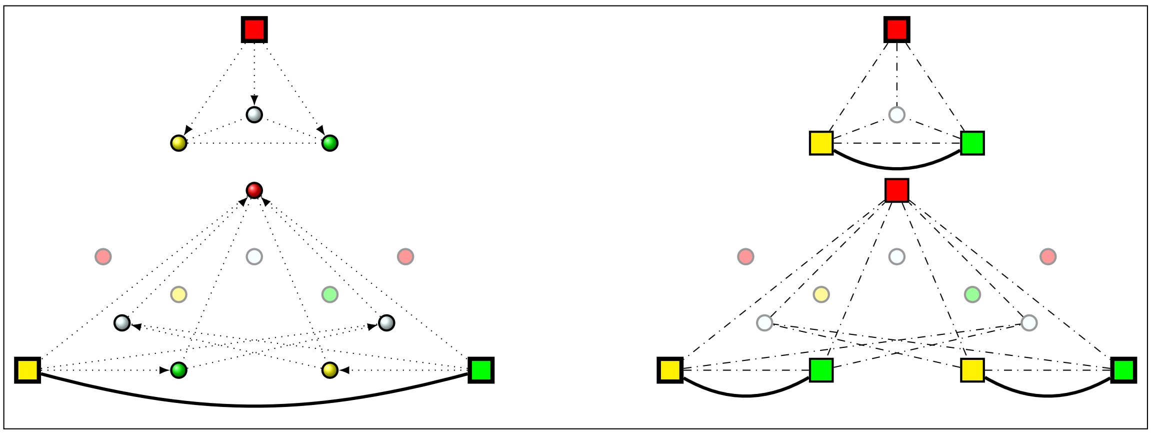

**Fig. 19.** Construction of the root components reachability graphs $\mathcal{RRG}_0$ and $\mathcal{RRG}_1$ from the 1-layer protocol complex for the 2-chain message adversary 2C. Since $\mathcal{RRG}_1$ already partitions into connected components containing only compatible border root components, consensus is solvable here. Note that one additional iteration even removes all edges in $\mathcal{RRG}_2$.

procedure a consensus termination time bound of $k(n-1)$, this is quite conservative. The major disadvantage of the $\mathcal{RRG}$-based decision procedure is its computational complexity, however: After all, the number of different border root components is exponential in $n$ and thus makes the initial root reachability graph $\mathcal{RRG}_1$ exponentially larger than the initial nerve graph $\mathcal{N}_0$.

## 6 Conclusion

We presented a topological view on deciding consensus solvability in dynamic graphs controlled by oblivious message adversaries. Compared to the purely combinatorial approach [39], it not only provides additional insights into the roots of the possible exponential blowup of both the iteration complexity of the decision procedure and the termination time of distributed consensus, but also results in a decision procedure for consensus termination within $k$ rounds. Thanks to our novel concept of a communication pseudosphere, which can be viewed as the message passing analogon of the chromatic subdivision, it is also a promising basis for further generalizations, e.g., for other decision problems and other message adversaries.

## References

1. Abraham, I., Malkhi, D., et al.: The blockchain consensus layer and BFT. Bulletin of EATCS **3**(123) (2017)
2. Afek, Y., Gafni, E.: Asynchrony from synchrony. In: Frey, D., Raynal, M., Sarkar, S., Shyamasundar, R., Sinha, P. (eds.) Distributed Computing and Networking, Lecture Notes in Computer Science, vol. 7730, pp. 225–239. Springer Berlin Heidelberg (2013). https://doi.org/10.1007/978-3-642-35668-1_16, `http://dx.doi.org/10.1007/978-3-642-35668-1_16`
3. Attiya, H., Castañeda, A.: A non-topological proof for the impossibility of k-set agreement. Theor. Comput. Sci. **512**, 41–48 (2013)
4. Attiya, H., Castañeda, A., Herlihy, M., Paz, A.: Bounds on the step and namespace complexity of renaming. SIAM J. Comput. **48**(1), 1–32 (2019). https://doi.org/10.1137/16M1081439, `https://doi.org/10.1137/16M1081439`
5. Biely, M., Robinson, P., Schmid, U.: Agreement in directed dynamic networks. In: Proceedings 19th International Colloquium on Structural Information and Communication Complexity (SIROCCO'12). pp. 73–84. LNCS 7355, Springer-Verlag (2012)
6. Biely, M., Robinson, P., Schmid, U., Schwarz, M., Winkler, K.: Gracefully degrading consensus and k-set agreement in directed dynamic networks. Theoretical Computer Science **726**, 41–77 (2018). https://doi.org/https://doi.org/10.1016/j.tcs.2018.02.019, `http://www.sciencedirect.com/science/article/pii/S0304397518301166`
7. Biely, M., Schmid, U., Weiss, B.: Synchronous consensus under hybrid process and link failures. Theoretical Computer Science **412**(40), 5602 – 5630 (2011), `http://www.sciencedirect.com/science/article/pii/S0304397510005359`, http://dx.doi.org/10.1016/j.tcs.2010.09.032
8. Castañeda, A., Fraigniaud, P., Paz, A., Rajsbaum, S., Roy, M., Travers, C.: Synchronous t-resilient consensus in arbitrary graphs. In: SSS. Lecture Notes in Computer Science, vol. 11914, pp. 53–68. Springer (2019)
9. Castañeda, A., Fraigniaud, P., Paz, A., Rajsbaum, S., Roy, M., Travers, C.: A topological perspective on distributed network algorithms. Theor. Comput. Sci. **849**, 121–137 (2021)
10. Charron-Bost, B., Függer, M., Nowak, T.: Approximate consensus in highly dynamic networks: The role of averaging algorithms. In: Automata, Languages, and Programming, Lecture Notes in Computer Science, vol. 9135, pp. 528–539. Springer Berlin Heidelberg (2015)
11. Charron-Bost, B., Schiper, A.: The Heard-Of model: computing in distributed systems with benign faults. Distributed Computing **22**(1), 49–71 (Apr 2009). https://doi.org/10.1007/s00446-009-0084-6
12. Coulouma, É., Godard, E., Peters, J.G.: A characterization of oblivious message adversaries for which consensus is solvable. Theor. Comput. Sci. **584**, 80–90 (2015). https://doi.org/10.1016/j.tcs.2015.01.024, `http://dx.doi.org/10.1016/j.tcs.2015.01.024`
13. El-Hayek, A., Henzinger, M., Schmid, S.: Asymptotically tight bounds on the time complexity of broadcast and its variants in dynamic networks. In: 14th Innovations in Theoretical Computer Science (ITCS) (2023)
14. Fevat, T., Godard, E.: 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. pp. 1001–1011 (2011). https://doi.org/10.1109/IPDPS.2011.96, `https://doi.org/10.1109/IPDPS.2011.96`

15. Fraigniaud, P., Paz, A.: The topology of local computing in networks. In: 47th International Colloquium on Automata, Languages, and Programming, ICALP. LIPIcs, vol. 168, pp. 128:1–128:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020). https://doi.org/10.4230/LIPIcs.ICALP.2020.128, `https://doi.org/10.4230/LIPIcs.ICALP.2020.128`
16. Függer, M., Nowak, T., Schwarz, M.: Tight bounds for asymptotic and approximate consensus. In: Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing. pp. 325–334. PODC '18, ACM, New York, NY, USA (2018). https://doi.org/10.1145/3212734.3212762, `http://doi.acm.org/10.1145/3212734.3212762`
17. Gafni, E.: Round-by-round fault detectors (extended abstract): unifying synchrony and asynchrony. In: Proceedings of the Seventeenth Annual ACM Symposium on Principles of Distributed Computing. pp. 143–152. ACM Press, Puerto Vallarta, Mexico (1998). https://doi.org/10.1145/277697.277724
18. Galeana, H.R., Rajsbaum, S., Schmid, U.: Continuous tasks and the asynchronous computability theorem. In: Braverman, M. (ed.) 13th Innovations in Theoretical Computer Science Conference, ITC S 2022, January 31 - February 3, 2022, Berkeley, CA, USA. LIPIcs, vol. 215, pp. 73:1–73:27. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2022). https://doi.org/10.4230/LIPIcs.ITCS.2022.73, `https://doi.org/10.4230/LIPIcs.ITCS.2022.73`
19. Godard, E., Perdereau, E.: k-set agreement in communication networks with omission faults. In: Fatourou, P., Jiménez, E., Pedone, F. (eds.) 20th International Conference on Principles of Distributed Systems, OPODIS 2016, December 13-16, 2016, Madrid, Spain. LIPIcs, vol. 70, pp. 8:1–8:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2016). https://doi.org/10.4230/LIPIcs.OPODIS.2016.8, `https://doi.org/10.4230/LIPIcs.OPODIS.2016.8`
20. Herlihy, M., Kozlov, D.N., Rajsbaum, S.: Distributed Computing Through Combinatorial Topology. Morgan Kaufmann (2013), `https://store.elsevier.com/product.jsp?isbn=9780124045781`
21. Herlihy, M., Rajsbaum, S.: Concurrent computing and shellable complexes. In: Lynch, N.A., Shvartsman, A. (eds.) Distributed Computing, Lecture Notes in Computer Science, vol. 6343, pp. 109–123. Springer Berlin Heidelberg (2010). https://doi.org/10.1007/978-3-642-15763-9_10, `http://dx.doi.org/10.1007/978-3-642-15763-9_10`
22. Herlihy, M., Rajsbaum, S., Tuttle, M.R.: An overview of synchronous message-passing and topology. Electron. Notes Theor. Comput. Sci. **39**(2), 1–17 (2001). https://doi.org/10.1016/S1571-0661(05)01148-5, `https://doi.org/10.1016/S1571-0661(05)01148-5`
23. 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`
24. Keidar, I., Shraer, A.: Timeliness, failure detectors, and consensus performance. In: Proceedings of the twenty-fifth annual ACM SIGACT-SIGOPS symposium on Principles of Distributed Computing (PODC'06). pp. 169–178. ACM Press, New York, NY, USA (2006)
25. Kozlov, D.N.: Chromatic subdivision of a simplicial complex. Homology, Homotopy and Applications **14**(1), 197–209 (2012). https://doi.org/https://dx.doi.org/10.4310/HHA.2012.v14.n2.a12
26. Kozlov, D.N.: Structure theory of flip graphs with applications to weak symmetry breaking. CoRR **abs/1511.00457** (2015), `http://arxiv.org/abs/1511.00457`
27. Kozlov, D.N.: Combinatorial Topology of the Standard Chromatic Subdivision and Weak Symmetry Breaking for Six Processes, pp. 155–194. Springer International Publishing (2016)
28. Kuhn, F., Oshman, R.: Dynamic networks: Models and algorithms. SIGACT News **42(1)**, 82–96 (2011)
29. Kuhn, F., Lynch, N.A., Oshman, R.: Distributed computation in dynamic networks. In: STOC. pp. 513–522 (2010)
30. Kuhn, F., Oshman, R., Moses, Y.: Coordinated consensus in dynamic networks. In: Proceedings of the 30th annual ACM SIGACT-SIGOPS symposium on Principles of distributed computing. PODC '11, ACM (2011)
31. Newport, C., Kotz, D., Yuan, Y., Gray, R.S., Liu, J., Elliott, C.: Experimental Evaluation of Wireless Simulation Assumptions. SIMULATION: Transactions of The Society for Modeling and Simulation International **83**(9), 643–661 (Sep 2007). https://doi.org/10.1177/0037549707085632, `http://dx.doi.org/10.1177/0037549707085632`
32. Nowak, T., Schmid, U., Winkler, K.: Topological characterization of consensus under general message adversaries. In: Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019. pp. 218–227 (2019). https://doi.org/10.1145/3293611.3331624, `https://doi.org/10.1145/3293611.3331624`, (full version: http://arxiv.org/abs/1905.09590)
33. Ongaro, D., Ousterhout, J.: In search of an understandable consensus algorithm. In: Proc. USENIX Annual Technical Conference (ATC). pp. 305–319 (2014)
34. Santoro, N., Widmayer, P.: Time is not a healer. In: Proc. 6th Annual Symposium on Theor. Aspects of Computer Science (STACS'89). pp. 304–313. LNCS 349, Springer-Verlag, Paderborn, Germany (Feb 1989)

35. Santoro, N., Widmayer, P.: Agreement in synchronous networks with ubiquitous faults. Theoretical Computer Science **384**(2–3), 232–249 (October 2007)
36. Schmid, U., Weiss, B., Keidar, I.: Impossibility results and lower bounds for consensus under link failures. SIAM Journal on Computing **38**(5), 1912–1951 (2009). https://doi.org/10.1137/S009753970443999X, https://epubs.siam.org/doi/pdf/10.1137/S009753970443999X
37. Schwarz, M., Winkler, K., Schmid, U.: Fast consensus under eventually stabilizing message adversaries. In: Proceedings of the 17th International Conference on Distributed Computing and Networking. pp. 7:1–7:10. ICDCN '16, ACM, New York, NY, USA (2016). https://doi.org/10.1145/2833312.2833323, http://doi.acm.org/10.1145/2833312.2833323
38. Shimi, A., Castañeda, A.: K-set agreement bounds in round-based models through combinatorial topology. In: Proceedings of the 39th Symposium on Principles of Distributed Computing. p. 395–404. PODC '20, Association for Computing Machinery, New York, NY, USA (2020). https://doi.org/10.1145/3382734.3405752, https://doi.org/10.1145/3382734.3405752
39. Winkler, K., Paz, A., Rincon Galeana, H., Schmid, S., Schmid, U.: The Time Complexity of Consensus Under Oblivious Message Adversaries. In: Tauman Kalai, Y. (ed.) 14th Innovations in Theoretical Computer Science Conference (ITCS 2023). Leibniz International Proceedings in Informatics (LIPIcs), vol. 251, pp. 100:1–100:28. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2023). https://doi.org/10.4230/LIPIcs.ITCS.2023.100, https://drops.dagstuhl.de/opus/volltexte/2023/17603
40. Winkler, K., Schmid, U.: An overview of recent results for consensus in directed dynamic networks. Bulletin of the EATCS **128** (2019), http://bulletin.eatcs.org/index.php/beatcs/article/view/581
41. Winkler, K., Schmid, U., Moses, Y.: A characterization of consensus solvability for closed message adversaries. In: 23rd International Conference on Principles of Distributed Systems, OPODIS 2019, December 17-19, 2019, Neuchâtel, Switzerland. LIPIcs, vol. 153, pp. 17:1–17:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2019). https://doi.org/10.4230/LIPIcs.OPODIS.2019.17, https://doi.org/10.4230/LIPIcs.OPODIS.2019.17
42. Winkler, K., Schwarz, M., Schmid, U.: Consensus in directed dynamic networks with short-lived stability. Distributed Computing **32**(5), 443–458 (2019). https://doi.org/https://doi.org/10.1007/s00446-019-00348-0, https://doi.org/10.1007/s00446-019-00348-0