A Topology by Geometrization for Sub-Iterated Immediate Snapshot Message Adversaries and Applications to Set-Agreement

Yannis Coutouly
Laboratoire d'Informatique et des Systèmes - Université Aix-Marseille, France
CNRS, Marseille, France

Emmanuel Godard
Laboratoire d'Informatique et des Systèmes - Université Aix-Marseille, France
CNRS, Marseille, France

Abstract

The Iterated Immediate Snapshot model (IIS) is a central model in the message adversary setting. We consider general message adversaries whose executions are arbitrary subsets of the executions of the IIS message adversary. We present a complete and explicit characterization and lower bounds for solving *set-agreement* for general sub-IIS message adversaries.

In order to have this characterization, we introduce a new topological approach for such general adversaries, closely associating executions to *geometric* simplicial complexes. This way, it is possible to define and explicitly construct a topology directly on the considered sets of executions. We believe this topology by geometrization to be of independent interest and a good candidate to investigate distributed computability in general sub-IIS message adversaries, as this could provide both simpler and more powerful ways of using topology for distributed computability.

2012 ACM Subject Classification Theory of computation → Distributed algorithms

Keywords and phrases topological methods, geometric simplicial complex, set-agreement

Digital Object Identifier 10.4230/LIPIcs.DISC.2023.15

Funding This work has been funded by ANR project DUCAT – ANR-20-CE48-0006.

1 Introduction

The $k-$set-agreement problem is a standard problem in distributed computing and it is known to be a good benchmark for topological approaches. The $k-$set-agreement problem is a distributed task where processes have to agree on no more than k different initial values. The *set-agreement* problem is the $k-$set agreement problem with $k+1$ processes. A review by Raynal can be found in [23]. Since the seminal works of Herlihy-Shavit, Borowsky-Gafni and Saks Zaharoglou [14, 3, 25], using topological methods has proved very fruitful for distributed computing and for distributed computability in particular. In the shared memory model, the impossibility of wait-free $k-$set agreement for more than $k+1$ processes is one of the crowning achievements of topological methods in distributed computing.

Since those first results, the topological framework has been refined to be presented in a more effective way. In particular, the Iterated Immediate Snapshot model (IIS) is a special message adversary that has been proposed as a central model to investigate distributed computability. In this paper we consider the set-agreement problem in the context of message adversaries defined as subsets of executions of the IIS message adversary.

1.1 Main Contributions

The main contribution is the first complete and explicit characterization of sub-IIS message adversaries for which set-agreement is solvable. We introduce in Section 3 a geometrization mapping *geo* that associates a point in $\mathbb{R}^N$ (with a large enough N) to any execution of IIS_n, the set of IIS executions for $n+1$ processes. The characterization of Th. 26 states that set-agreement is solvable for $\mathcal{M} \subset IIS_n$ if and only if the geometrization of $\mathcal{M}$ has a "hole", *i.e.* $geo(\mathcal{M})$ is a strict subset of the convex hull of $\mathbb{S}^n$, the simplex of dimension n. In Section 4, we describe and prove important properties of the geometrization *geo*. In particular we give a combinatorial description of the sets $geo^{-1}(x)$, for $x \in \mathbb{R}^N$, in Th.25. Interestingly, we show that these sets can have only three possible size: 1, 2 and infinite size. Together with the previous theorem, this gives a explicit and complete characterization of the subsets of the executions of the IIS message adversary for which set-agreement is solvable. We also apply our technique to derive new lower bounds for general message adversaries solving set-agreement.

The *geo* mapping is central to our characterization. The second main contribution is to show that there is a natural topological interpretation of this mapping. Using *geo*, we present in Section 3 a new topology that is defined directly on the set of IIS executions. We believe this topology by geometrization to be of independent interest and a good candidate to investigate distributed computability in general sub-IIS message adversaries, as this could provide both simpler and more powerful ways of using topology for distributed computability of any task.

In order to handle general message adversaries, we consider here simplicial complexes primarily as *geometric* simplicial complexes. The standard chromatic subdivision is the combinatorial topology representation of one round of the Immediate Snapshot model. Its simple and regular structure makes topological reasoning attractive. In this paper, we introduce a new universal algorithm and show its relationship with the standard chromatic subdivision as exposed in the geometric simplicial complex setting. This new algorithm is called the Chromatic Averaging algorithm, it averages with specific weights vectors of $\mathbb{R}^N$ at each node. Running the Chromatic Average Algorithm in the IIS_n model yields a geometric counterpart in $\mathbb{R}^N$ to any given infinite execution of IIS_n. The geometrization mapping $geo(w)$ of an execution $w \in IIS_n$ is defined as the convergence value of running the Chromatic Average Algorithm under execution w.

The topology on the set of executions is then the topology induced from the standard topology in $\mathbb{R}^N$ by the mapping geo : the open sets are pre-images $geo^{-1}(\Omega)$ of the open sets Ω of R^N. The standard euclidean topology of $\mathbb{R}^N$ is simple and well understood, however, since *geo* is not injective, it is necessary to describe so-called "non-separable sets" in order to fully understand the new topology. In topology, two distinct elements x, y are said to be non-separable if for any two neighbourhoods Ω_x of x and Ω_y of y, we have $\Omega_x \cap \Omega_y \neq \emptyset$. In our setting, two executions are non-separable when they have the same image via the mapping *geo*, we call such pre-image sets $geo-classes$. Understanding those sets is central to the characterization of solvability of set-agreement. It is also central to precisely describe the properties of the geometrization topology. So we introduce first the geometrization topology and in Section 4, we investigate the $geo-$classes. In Section 5, we apply our framework to derive the characterization of computability of set-agreement and lower bounds. In the conclusion, we discuss the perspective of possible application of the geometrization topological framework to arbitrary tasks.

1.2 Related Works

In [9], the "two generals problem", that is the consensus problem for two processes is investigated for arbitrary sub-IIS models by Godard and Perdereau. Given that consensus for two processes is actually set-agreement, the characterization of solvability of set-agreement presented here is a generalization to any number of processes of the results of [9].

One of the most advanced results toward the investigation of general sub-ISS adversary are presented in the work of Kuznetsov and Rieutord [24, 17]. Their adversaries are iterated and are related to so-called affine task. Our work consider more general sub-IIS adversaries, including non-iterated adversaries, but the distributed computability is presented only for the set-agreement task.

In [8], Gafni, Kuznetsov and Manolescu investigate models that are more general subsets of the Iterated Immediate Snapshot model, where the execution sets are closed under a specific relation. We believe our tools can provide a simpler, and less error-prone (see [9, Section 5.1]), framework to investigate distributed computability of sub-IIS models. In particular, their closure relation is nicely interpreted here as exactly the non-separability relation of the geometrization topology.

In a series of works, averaging algorithms to solve relaxed versions of the Consensus problem, including approximate Consensus, have been investigated. In [6], Charron-Bost, Függer, and Nowak have used matrix oriented approaches to show the convergence of different averaging algorithms. We use a similar stochastic matrix technique here to prove the convergence of the Chromatic Average Algorithm. In [7], Függer, Nowak and Schwarz have shown tight bounds for solving approximate and asymptotic Consensus in quite general message adversaries.

In [20], Nowak, Schmid, and Winkler propose knowledge-based topologies for all message adversaries. It is then used to characterize message adversaries that can solve Consensus. The scope of [20] is larger than the scope of this paper, however, note that contrary to those topologies, that are *implicitly* defined by indistinguishability of local knowledge, the geometrization topology here is explicitly defined and fully described by Th. 25. Recently, in [2], Attiya, Castañeda and Nowak presented a corrected version of the general characterisation of [8] in this framework. They also give as application a characterisation for set-agreement based upon terminating subdivisions [2, Thm. 4.2]. We believe the characterisation given in Thm. 26 to be more precise. An interesting open question would be to compare the geometrization topology to the knowledge-based ones defined in [20, 2].

2 Models and Definitions

2.1 Message Adversaries

We introduce and present here our notations. Let $n \in \mathbb{N}$, we consider systems with $n+1$ processes. We denote $\Pi_n = [0, .., n]$ the set of processes. Since sending a message is an asymmetric operation, we will work with directed graphs. We recall the main standard definitions in the following.

We use standard directed graph (or digraph) notations: given G, $V(G)$ is the set of vertices, $A(G) \subset V(G) \times V(G)$ is the set of arcs.

▶ **Definition 1.** *We denote by $\mathcal{G}_n$ the set of directed graphs with vertices in Π_n.*

A dynamic graph $\mathbf{G}$ *is a sequence $G_1, G_2, \cdots, G_r, \cdots$ where G_r is a directed graph with vertices in Π_n. We also denote by $\mathbf{G}(r)$ the digraph G_r. A* message adversary *is a set of dynamic graphs.*

Since that n will be mostly fixed through the paper, we use Π for the set of processes and $\mathcal{G}$ for the set of graphs with vertices Π when there is no ambiguity.

Intuitively, the graph at position r of the sequence describes whether there will be, or not, transmission of some messages sent at round r. A formal definition of an execution under a scenario will be given in Section 2.3.

We will use the standard following notations in order to describe more easily our message adversaries [21]. A sequence is seen as a word over the alphabet $\mathcal{G}$.

▶ **Definition 2.** *Given $A \subset \mathcal{G}$, A^* is the set of all finite sequences of elements of A, A^ω is the set of all infinite ones and $A^\infty = A^* \cup A^\omega$.*

Given $\mathbf{G} \in \mathcal{G}^\omega$, if $\mathbf{G} = \mathbf{HK}$, with $\mathbf{H} \in \mathcal{G}_n^*, \mathbf{K} \in \mathcal{G}_n^\omega$, we say that $\mathbf{H}$ is *a prefix* of $\mathbf{G}$, and $\mathbf{K}$ *a suffix*. $Pref(\mathbf{G})$ denotes the set of prefixes of $\mathbf{G}$. An adversary of the form A^ω is called an *oblivious adversary* or an *iterated adversary*. A word in $\mathcal{M} \subset \mathcal{G}^\omega$ is called a *communication scenario* (or *scenario* for short) of message adversary $\mathcal{M}$. Given a word $\mathbf{H} \in \mathcal{G}^*$, it is called a *partial scenario* and $len(\mathbf{H})$ is the length of this word. The prefix of $\mathbf{G}$ of length r is denoted $\mathbf{G}_{|r}$ (not to be confused with $\mathbf{G}(r)$ which is the r-th letter of $\mathbf{G}$, it the digraph at time r).

The following definitions provide a notion of causality when considering infinite word over digraphs.

▶ **Definition 3** ([5])**.** *Let $\mathbf{G}$ a sequence $G_1, G_2, \cdots, G_r, \cdots$. Let $p, q \in \Pi$. There is a* journey *in $\mathbf{G}$ at time r from p to q, if there exists a sequence $p_0, p_1, \ldots, p_t \in \Pi$, and a sequence $r \leq i_0 < i_1 < \cdots < i_t \in \mathbb{N}$ where we have*

- *$p_0 = p, p_t = q$,*
- *for each $0 < j \leq t$, $(p_{j-1}, p_j) \in A(G_{i_j})$*

This is denoted $p \underset{\mathbf{G}}{\overset{r}{\rightsquigarrow}} q$. We also say that p is causally influencing q from round r in $\mathbf{G}$.

2.2 Iterated Immediate Snapshot Message Adversary

We say that a graph G has the *Immediacy Property* if for all $a, b, c \in V(G)$, $(a, b), (b, c) \in A(G)$ implies that $(a, c) \in A(G)$. A graph G has the *containment Property* if for all $a, b \in V(G)$, $(a, b) \in A(G)$ or $(b, a) \in A(G)$.

▶ **Definition 4** ([12])**.** *We set $IS_n = \{G \in \mathcal{G}_n \mid G$ has the Immediacy and Containment properties$\}$. The Iterated Immediate Snapshot message adversary for $n+1$ processes is the message adversary $IIS_n = IS_n^\omega$.*

The Iterated Immediate Snapshot model was first introduced as a (shared) memory model and then has been shown to be equivalent to the above message adversary first as tournaments and iterated tournaments [4, 1], then as this message adversary [12, 13]. See also [22] for a survey of the reductions involved in these layered models.

We show how standard fault environments are conveniently described in our framework.

▶ **Example 5.** Consider a message passing system with $n+1$ processes where, at each round, all messages could be lost. The associated message adversary is $\mathcal{G}_n^\omega$.

▶ **Example 6.** Consider a system with two processes $\{\circ, \bullet\}$ where, at each round, only one message can be lost. The associated message adversary is $\{\circ\!\leftrightarrow\!\bullet, \circ\!\leftarrow\!\bullet, \circ\!\rightarrow\!\bullet\}^\omega$. This is IIS_1.

2.3 Execution of a Distributed Algorithm

Given a message adversary $\mathcal{M}$ and a set of initial configurations $\mathcal{I}$, we define what is an execution of a given algorithm $\mathcal{A}$ subject to $\mathcal{M}$ with initialization $\mathcal{I}$. An execution is constituted of an initialization step, and a (possibly infinite) sequence of rounds of messages exchanges and corresponding local state updates. When the initialization is clear from the context, we will use *scenario* and *execution* interchangeably.

An execution of an algorithm $\mathcal{A}$ under scenario $w \in \mathcal{M}$ and initialization $\iota \in \mathcal{I}$ is the following. This execution is denoted $\iota.w$. First, ι affects an initial state to all processes of Π.

A round is decomposed in 3 steps: sending, receiving, updating the local state. At round $r \in \mathbb{N}$, messages are sent by the processes using the `SendAll()` primitive. The fact that the corresponding receive actions, using the `Receive()` primitive, will be successful depends on $G = w(r)$, G is called the *instant graph* at round r.

Let $p, q \in \Pi$. The message sent by p is received by q on the condition that the arc $(p, q) \in A(G)$. Then, all processes update their state according to the received values and $\mathcal{A}$. Note that, it is usually assumed that p always receives its own value, that is $(p, p) \in A(G)$ for all p and G.

Let $w \in \mathcal{M}, \iota \in \mathcal{I}$. Given $u \in Pref(w)$, we denote by $\mathbf{s}_p(\iota.u)$ the state of process p at the $len(u)$-th round of the algorithm $\mathcal{A}$ under scenario w with initialization ι. This means that $\mathbf{s}_p(\iota.\varepsilon)$ represents the initial state of p in ι, where ε denotes the empty word.

A task is given by a set $\mathcal{I}$ of initial configurations, a set of output values Out and a relation Δ, the specification, between initial configurations and output configuration[1]. We say that a process *decides* when it outputs a value in Out. Finally and classically,

▶ **Definition 7.** *An algorithm $\mathcal{A}$ solves a Task $(\mathcal{I}, Out, \Delta)$ for the message adversary $\mathcal{M}$ if for any $\iota \in \mathcal{I}$, any scenario $w \in \mathcal{M}$, there exist u a prefix of w such that the states of the processes $out = (\mathbf{s}_0(\iota.u), \ldots, \mathbf{s}_n(\iota.u))$ satisfy the specification of the task, ie $\iota \Delta out$.*

3 A Topology by Geometrization

In this paper we present a new topological approach for investigating distributed computability. It extends the known simplicial complexes-based known method for finite executions to infinite executions without considering infinite additional complexes like in [8]. This enables to define directly a topology on the set of executions of the standard Iterated Immediate Snapshot model IIS_n.

3.1 Combinatorial Topology Definitions

3.1.1 Geometric Simplicial Complexes

Before giving the definition of the geometrization topology in Sect. 3.2.2, we state the definition of simplicial complexes, but not first as abstract complex, as is usually done in distributed computing, but primarily as geometrical objects in $\mathbb{R}^N$. This is the reason we call this definition the geometrization topology. Intuitively we will associatea point in $\mathbb{R}^N$ to any execution via a geometrization mapping *geo*. The geometrization topology is the topology induced by geo^{-1} from the standard topology in $\mathbb{R}^N$. This also makes *geo*

[1] Note that the standard definition in the topological setting involves carrier map that we do not consider here for we will consider only one specific task, the Set Agreement problem.

continuous by definition. In the standard approach, geometric simplices are also used but they are introduced as geometric realizations of the abstract simplicial complexes. As will be seen later, when dealing with infinite complexes, the standard topology of these simplices does not enable to handle the computability of distributed tasks since we will need to define an other topology. We show that the topology on infinite complexes, as defined in standard topology textbook, is different from the one we show here to be relevant for distributed computability. Note that to be correctly interpreted, the topology we construct is on the set of infinite executions, not on the complexes corresponding to finite executions.

The following definitions are standard definitions from algebraic topology [19]. We fix an integer $N \in \mathbb{N}$ for this part. We denote $||x||$ the euclidean norm in $\mathbb{R}^N$. For a bounded subset $X \subset \mathbb{R}^n$, we denote $diam(X)$ its diameter.

▶ **Definition 8.** *Let $n \in \mathbb{N}$. A finite set $\sigma = \{x_0, \ldots, x_n\} \subset \mathbb{R}^N$ is called a* simplex *of dimension n if the vectors $\{x_1 - x_0, \ldots, x_n - x_0\}$ are linearly independent. We denote by $|\sigma|$ the convex hull of σ.*

▶ **Definition 9** ([19])**.** *A* simplicial complex *is a collection C of* simplices *such that:*
(a) *If $\sigma \in C$ and $\sigma' \subseteq \sigma$, then $\sigma' \in C$,*
(b) *If $\sigma, \tau \in C$ and $|\sigma| \cap |\tau| \neq \emptyset$ then there exists $\sigma' \in C$ such that $|\sigma| \cap |\tau| = |\sigma'|$, $\sigma' \subset \sigma, \sigma' \subset \tau$.*

We denote $\wr C \wr = \bigcup_{S \in C} |S|$, this is the *geometrization of C*.

Note that these definitions do not require complexes to be a finite collection of simplices. The simplices of dimension 0 (singleton) of C are called vertices, we denote $V(C)$ the set of vertices of C. A complex is pure of dimension n if all maximal simplices are of dimension n. In this case, a simplex of dimension $n-1$ is called a facet. The *boundary* of a simplex $\sigma = \{x_0, \ldots, x_n\}$ is the pure complex $\bigcup_{i \in [0,n]} \{x_j \mid j \in [0,n], i \neq j\}$ of dimension $n-1$. It is denoted $\delta(\sigma)$, it is the union of the facets of σ.

Let A and B be simplicial complexes. A map $f \colon V(A) \to V(B)$ defines a *simplicial map* if it preserves the simplices, *i.e.* for each simplex σ of A, the image $f(\sigma)$ is a simplex of B. By linear combination of the barycentric coordinates, f extends to the linear simplicial map $f \colon \wr A \wr \to \wr B \wr$, which is continuous. See [19, Lemma 2.7].

We also have colored simplicial complexes. These are simplicial complexes C together with a function $\chi : V(C) \to \Pi$ such that the restriction of χ on any maximal simplex of C is a bijection. A simplicial map that preserves colors is called chromatic.

Finally, $\mathbb{S}.$ will denote "the" simplex of dimension n. Through this paper we assume a fixed embedding in $\mathbb{R}^N$ for $\mathbb{S}. = (x_0^*, \ldots, x_n^*)$. We will also assume that its diameter $diam(\mathbb{S}.)$ is 1.

3.1.2 The Standard Chromatic Subdivision

Here we present the standard chromatic subdivision, [12] and [15], as a geometric complex. We start with subdivisions and chromatic subdivisions.

▶ **Definition 10** (Subdivision)**.** *A subdivision of a simplex S is a simplicial complex C with $\wr C \wr = |S|$.*

▶ **Definition 11** (Chromatic Subdivision)**.** *Given $(S, \mathcal{P})$ a chromatic simplex, a chromatic subdivision of S is a chromatic simplicial complex $(C, \mathcal{P}_C)$ such that*
- *C is a subdivision of S*(i.e.$\wr C \wr = |S|$),
- *$\forall x \in V(S), \mathcal{P}_C(x) = \mathcal{P}(x)$.*

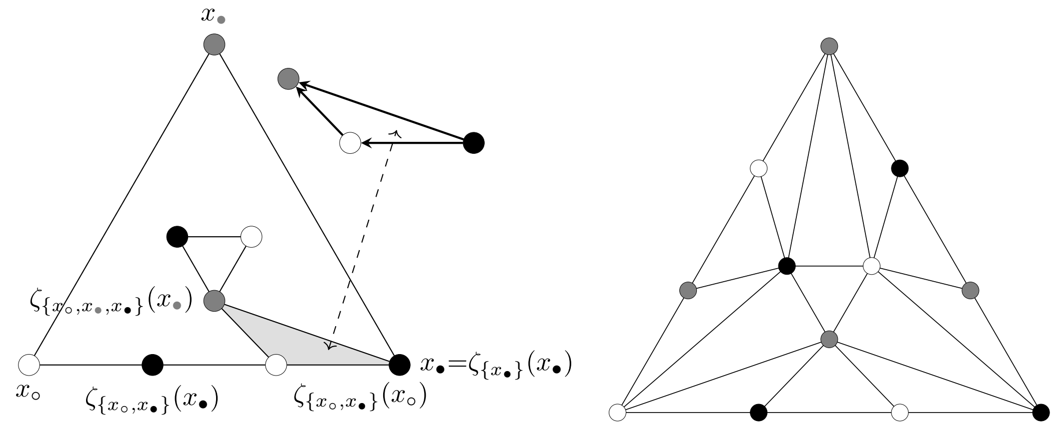

(a) Images by ζ_V for various $V \subset \{\circ, {\color{gray}\bullet}, \bullet\}$.

(b) $Chr(\mathbb{S}^2)$.

Figure 1 Standard chromatic subdivision construction for dimension 2. On the left, the association between an instant graph of IS_2 (top) and a simplex of $Chr(\mathbb{S}^2)$ (grey area) is illustrated.

Note that it is not necessary to assume $V(S) \subset V(C)$ here, since the vertices of the simplex S being extremal points, they are necessarily in $V(C)$.

We start by defining some geometric transformations of simplices (here seen as sets of points). The choice of the coefficients will be justified later.

▶ **Definition 12.** *Consider a set $V = (y_0, \ldots, y_d)$ of size $d+1$ in $\mathbb{R}^N$. We define the function $\zeta_V : V \longrightarrow \mathbb{R}^N$ by, for all $j \in [0, d]$*

$$\zeta_V(y_j) = \frac{1 - \frac{d}{2d+1}}{d+1} y_j + \sum_{i \neq j} \frac{1 + \frac{1}{2d+1}}{d+1} y_i$$

We now define in a geometric way the *standard chromatic subdivision* of colored simplex $(S, \mathcal{P})$, where $S = \{x_0, x_1, \ldots, x_n\}$ and $\mathcal{P}(x_i) = i$.

The chromatic subdivision $\mathrm{C}hr(S)$ for the colored simplex $S = \{x_0, \ldots, x_n\}$ is a simplicial complex defined by the set of vertices $V(Chr(S)) = \{\zeta_V(x_i) \mid i \in [0, n], V \subset V(S), x_i \in V\}$.

For each pair (i, V), $i \in [0, n]$ and $V \subset V(S)$, there is an associated vertex y of $\mathrm{C}hr(S)$, and conversely each vertex has an associated pair. The *color* of (i, V) is i. The set V is called the *view*. We define Φ the following *presentation* of a vertex y, $\Phi(y) = (\mathcal{P}(y), V_y)$ where $\mathcal{P}(y) = i$ and $V_y = V$.

The simplices of $\mathrm{C}hr(S)$ are the set of $d+1$ points $\{\zeta_{V_0}(x_{i_0}), \cdots, \zeta_{V_d}(x_{i_d})\}$ where

- there exists a permutation π on $[0, d]$ such that $V_{\pi(0)} \subseteq \cdots \subseteq V_{\pi(d)}$,
- If $i_j \in \mathcal{P}(V_\ell)$ then $V_j \subset V_\ell$.

In Fig. 1, we present the construction for $Chr(\mathbb{S}^2)$. For convenience, we associate $\circ, {\color{gray}\bullet}, \bullet$ to the processes $0, 1, 2$ respectively. In Fig. 1a, we consider the triangle $x_\circ, x_{\color{gray}\bullet}, x_\bullet$ in $\mathbb{R}^2$, with $x_\circ = (0,0)$, $x_\bullet = (1,0)$, $x_{\color{gray}\bullet} = (\frac{1}{2}, \frac{\sqrt{3}}{2})$. We have that $\zeta_{\{x_\circ, x_\bullet\}}(x_\bullet) = (\frac{1}{3}, 0)$, $\zeta_{\{x_\circ, x_\bullet\}}(x_\circ) = (\frac{2}{3}, 0)$ and $\zeta_{\{x_\circ, x_{\color{gray}\bullet}, x_\bullet\}}(x_{\color{gray}\bullet}) = (\frac{1}{2}, \frac{\sqrt{3}}{10})$. The relation between instant graph (top) and simplex $\left\{(\frac{2}{3}, 0), (1, 0), (\frac{1}{2}, \frac{\sqrt{3}}{5})\right\}$ (grey area) is detailed in the following section.

In the following, we will be interested in iterations of $Chr(\mathbb{S}^n, \mathcal{P})$. The last property of the definition of chromatic subdivision means with we can drop the C index in the coloring of complex C and use $\mathcal{P}$ to denote the coloring at all steps. From its special role, it is called the *process color* and we drop $\mathcal{P}$ in $Chr(S, \mathcal{P})$ using in the following $Chr(S)$ for all simplices S of iterations of $Chr(\mathbb{S}^n)$.

In [16], Kozlov showed how the standard chromatic subdivision complex relates to Schlegel diagrams (special projections of cross-polytopes), and used this relation to prove the standard chromatic subdivision was actually a subdivision.

In [12, section 3.6.3], a general embedding in R^n parameterized by $\epsilon \in \mathbb{R}$ is given for the standard chromatic subdivision. The geometrization here is done choosing $\epsilon = \frac{d}{2d+1}$ in order to have "well balanced" drawings.

3.2 Encoding Iterated Immediate Snapshots Configurations

3.2.1 Algorithms in the Iterated Immediate Snapshots Model

It is well known, see *e.g.* [12, Chap. 3&4, Def. 3.6.3], that each maximal simplex $S = \{\zeta_{V_0}(x_{i_0}), \cdots, \zeta_{V_n}(x_{i_n})\}$ from the chromatic subdivision of $\mathbb{S}^n$ can be associated with a graph of IS_n denoted $\Theta(S)$. We have $V(\Theta(S)) = \Pi_n = [0,n]$ and set $\Theta(\zeta_{V_j}(x_{i_j})) = \mathcal{P}(x_{i_j})$.The arcs are defined using the representation Φ of points, $A(\Theta(S)) = \{(i,j) \mid i \neq j, V_i \subseteq V_j\}$. The mapping θ will denote Θ^{-1}. We can transpose this presentation to an averaging algorithm called the *Chromatic Average* Algorithm presented in Algorithm 1.

Algorithm 1 The Chromatic Average Algorithm for process i.

$x \leftarrow x_i^*$;
Loop *forever*
  `SendAll`$((i,x))$;
  $V \leftarrow$`Receive()` // set of all received messages;
  $d \leftarrow sizeof(V) - 1$ // i received $d+1$ messages including its own ;
  $x = \frac{1-\frac{d}{2d+1}}{d+1} x + \sum_{(j,x_j)\in V, j\neq i} \frac{1+\frac{1}{2d+1}}{d+1} x_j$;

Executing one round of the loop in Chromatic Average for instant graph G, the state of process i is $x'_i = \zeta_{V_i}(x_i^*)$, where V_i is the view of i on this round, that is the set of (j, x_j) it has received; with $\Theta(\{\zeta_{V_0}(x_0^*), \cdots, \zeta_{V_n}(x_n^*)\}) = G$. See eg. in Fig. 1a, the simplex of the grey area corresponds to the ordered sequence of views $\{x_\bullet\} \subset \{x_\bullet, x_\circ\} \subset \{x_\bullet, x_\circ, x_{\bullet}\}$, associated to the directed graph depicted at the top right. Adjacency for a given i corresponds to the smallest subset containing x_i. By iterating, the chromatic subdivisions $Chr^r(\mathbb{S}^n)$ are given by the global state under all possible r rounds of the Chromatic Average Algorithm. Finite rounds give the Iterated Chromatic Subdivision (hence the name). This is an algorithm that is not meant to terminate (like the full information protocol). The infinite runs are used below to define a topology on IIS_n.

The Chromatic Average algorithm is therefore the geometric counterpart to the Full Information Protocol that is associated with Chr [12]. In particular, any algorithm can be presented as the Chromatic Average together with a terminating condition and an output function of x.

This one round transformation for the canonical $\mathbb{S}^n$ can actually be done for any simplex S of dimension n of $\mathbb{R}^N$. For $G \in IS_n$, we denote $\mu_G(S)$ the geometric simplex that is the image of S by one round of the Chromatic average algorithm under instant graph G.

The definitions of the previous section can be considered as mostly textbook (as in [12]), or folklore. To the best of our knowledge, the Chromatic Average Algorithm, as such, is new, and there is no previous complete proof of the link between the Chromatic Average Algorithm and iterated standard chromatic subdivisions. However, one shall remark that people are, usually, actually *drawing* standard chromatic subdivisions using the Chromatic Average Algorithm.

3.2.2 A Topology for IIS_n

Let $w \in IIS_n$, $w = G_1 G_2 \cdots$. For the prefix of w of size r, S a simplex of dimension n, we define $geo(w_{|r})(S) = \mu_{G_r} \circ \mu_{G_{r-1}} \circ \cdots \circ \mu_{G_1}(S)$. Finally, we set $geo(w) = \lim_{r \longrightarrow \infty} geo(w_{|r})(\mathbb{S}^n)$. We prove in Section A.1 that this actually converges.

We define the *geometrization topology* on the space IIS_n by considering as open sets the sets $geo^{-1}(\Omega)$ where Ω is an open set of $\mathbb{R}^N$. A collection of sets can define a topology when any union of sets of the collection is in the collection, and when any finite intersection of sets of the collection is in the collection. This is straightforward for a collection of inverse images of a collection that satisfies these properties.

A neighbourhood for point x is an open set containing x. In topological spaces, a pair of distinct points x, y is called *non-separable* if there does not exist two disjoint neighbourhoods of these points. The pre-images $geo^{-1}(x)$ that are not singletons are *non-separable sets*. We will see that we always have non-separable sets and that they play an important role for task solvability.

Subset of IIS_n will get the subset topology, that is , for $\mathcal{M} \subseteq IIS_n$, open sets are the sets $geo^{-1}(\Omega) \cap \mathcal{M}$ where Ω is an open set of $\mathbb{R}^N$. We set $\wr\mathcal{M}\wr = geo(\mathcal{M})$ the geometrization of $\mathcal{M}$.

Note that the geometrization should not be confused with the standard *geometric realization*. They are the same at the set level but not at the topological level, see in Section A.2. At times, in order to emphasize this difference, for a simplex $S \subset \mathbb{R}^N$, we will also use $\wr S\wr$ instead of $|S|$. The *geometrization* of C, denoted $\wr C\wr$, that is the union of the convex hulls $|\sigma|$ of the simplices σ of C, is endowed with the standard topology from $\mathbb{R}^N$. We also note this topological space as $\wr C\wr$.

4 Geometrization Equivalence

As will be be shown later, the geometrization has a crucial role in order to understand the relationship between sets of possible executions and solvability of distributed tasks. In this section, we describe more precisely the pre-images sets, that is subsets of IIS_n of the form $geo^{-1}(x)$ for $x \in |\mathbb{S}^n|$. In particular, we will get a description of the non-separable sets of execution.

4.1 Definitions

We say that two executions $w, w' \in IIS_n$ are *geo-equivalent* if $geo(w) = geo(w')$. The set of all w' such that $geo(w) = geo(w')$ is called the equivalence class of w. Since the topology we are interested in for $\wr IIS_n\wr$ is the one induced by the standard separable space $\mathbb{R}^N$ via the geo^{-1} mapping, it is straightforward to see that non-separable sets are exactly the geo-equivalence classes that are not singletons. In this section, we describe all equivalence classes and show that there is a finite number of possible size for these sets.

We define the sets $Solo(P)$, that correspond to subsets of instant graphs where the processes in $P \subset \Pi$ have no incoming message from processes outside of P. We have $Solo(\Pi) = IS_n$.

▶ **Definition 13.** *In the complex* $\mathrm{Chr}(\mathbb{S}^n)$*, with* $P \subset \Pi$*,* $Solo(P)$ *is the set of simplices* $T \in \mathrm{Chr}(\mathbb{S}^n)$ *such that* $\forall (p,q) \in A(\Theta(T)), q \in P \Rightarrow p \in P$.

We denote by $\mathcal{K}_\Pi$ the instant graph that is complete on Π. An execution w is said *fair* for P, $w \in Fair(P)$, if $w \in Solo(P)^\omega$ and for all $p, q \in P$, $\forall r \in \mathbb{N}$, we have $p \overset{r}{\underset{w}{\rightsquigarrow}} q$. Fairness for P means that processes in P are only influenced by processes in P, and that any process always influence other processes infinitely many times. We have the equivalent, and constructive definition:

▶ **Proposition 14.** *Let $w \in Solo(P)^\omega$. An execution w is $Fair$ for P if and only if w has no suffix in $\bigcup_{Q \neq \emptyset, Q \subsetneq P} Solo(Q)^\omega$.*

Proof. Assume we have a suffix s for w in $Solo(Q)^\omega$ with $Q \subsetneq P$.Let $p \in P \setminus Q$ and r the starting index of the suffix. Then $\forall q \in Q$, we must have $p \overset{r}{\underset{w}{\rightsquigarrow}} q$ by definition of fairness for P. Denote q_0 the first element of Q to be causally influenced by p at some time $t \geq r$. So q_0 receive a message from some $p' \in \Pi$, $p' \neq q_0$ at time t. Since $s \in Solo(Q)^\omega$, this means that q_0 can only receive message from processes in Q. Hence $p' \in Q$ and p' was influenced by p at time $t-1$. A contradiction with the minimality of q_0. So w is not in $Fair(P)$.

Conversely, assume that $w \notin Fair(P)$. Then $\exists p, q \in P, \exists s, \forall r \geq s, \neg p \overset{r}{\underset{w}{\rightsquigarrow}} q$. We set Q as the set of processes that causally influence q for all $r \geq s$. We have $p \notin Q$ so $Q \subsetneq P$. We denote s_0, a time at which no process of $\Pi \setminus Q$ influence a process in Q. By construction, the suffix at step s_0 is in $Solo(Q)^\omega$. ◀

4.2 First Results on Geometrization

We start by presenting a series of results about geometrization. Lemma 35 gives the following immediate corollaries.

▶ **Corollary 15.** *Let w a run in IIS_n, then $\forall r \in \mathbb{N}$, $geo(w) \in |geo(w_{|r})(\mathbb{S}^n)|$.*

▶ **Proposition 16.** *Let w, w' two geo-equivalent runs in IIS_n, then $\forall r \in \mathbb{N}$, $geo(w_{|r})(\mathbb{S}^n) \cap geo(w'_{|r})(\mathbb{S}^n) \neq \emptyset$.*

Proof. The intersection of the geometrizations $|geo(w_{|r})(\mathbb{S}^n)|$ and $|geo(w'_{|r})(\mathbb{S}^n)|$ contains at least $geo(w)$ by previous corollary. Since the simplices $geo(w_{|r})(\mathbb{S}^n)$ and $geo(w'_{|r})(\mathbb{S}^n)$ belong to the complex $\mathrm{C}hr^r(\mathbb{S}^n)$, they also intersect as simplices. ◀

▶ **Proposition 17.** *Let S a maximal simplex of the chromatic subdivision $\mathrm{C}hr\mathbb{S}^n$ that is not $\theta(\mathcal{K}(\Pi))$. Then there is $P \subsetneq \Pi$ such that $\Theta(S) \in Solo(P)$.*

Conversely we can describe $Solo(P)$ more precisely. We denote by $\delta(\mathbb{S}^n, P)$ the sub-simplex of $\mathbb{S}^n$ corresponding to $P \subset \Pi$. This is the boundary relative to P in $\mathbb{S}^n$, and we have that $\bigcup_{P \subsetneq \Pi} \delta(\mathbb{S}^n, P) = \bigcup_{p \in \Pi} \delta(\mathbb{S}^n, \pi \setminus p) = \delta(\mathbb{S}^n)$.

▶ **Proposition 18** (Boundaries of $\mathrm{C}hr$ are $Solo$)**.** *Let P a subset of Π. Then $Solo(P) = \{S \mid S$ a maximal simplex of $\mathrm{C}hr(\mathbb{S}^n), |S| \cap |\delta(\mathbb{S}^n, P)| \neq \emptyset\}$.*

Proof. Denote q such that $\Pi = P \cup \{q\}$. Then by construction, $Solo(P)$ corresponds exactly to the simplex where the processes in P do not receive any message from q, ie the simplex intersecting the boundary $\delta(\mathbb{S}^n, P)$. ◀

For a given size s of P, the $Solo$ sets are disjoint, however this does not form a partition of $\mathrm{C}hr$. Finally, by iterating the previous proposition, the boundaries of $\mathbb{S}^n$ are described by $Solo(P)^\omega$.

▶ **Proposition 19.** *Let P a subset of Π. We have $\wr Solo(P)^\omega \wr = |\delta(\mathbb{S}^n, P)|$.*

We can now state the main result that links geometrically fair executions and corresponding simplices: in a fair execution, the corresponding simplices, that are included by convexity, have to eventually be strictly included in the interior.

▶ **Proposition 20** (Geometric interpretation of $Fair$). *Let w an execution that is Fair for Π, then $\forall s \in \mathbb{N}, \exists r > s \in \mathbb{N}$, such that $\delta(geo(w_{|s})(S)) \cap geo(w_{|r})(S) = \emptyset$.*

Proof. Let $s \in \mathbb{N}$, and an execution w. We denote $T = geo(w_{|s}$. Consider a process $p \in \Pi$, for all process $q \neq p$ we have $p \underset{w}{\overset{s}{\rightsquigarrow}} q$. Since w is fair in Π, we can consider $r > s$ the time at which p is influencing all q from round s. At his step, for all $q \neq p$, the barycentric coordinate of the vertex of $geo(w_{|r}(S)$ of colour q relative to the vertex of $geo(w_{|r}(S)$ of colour p is strictly positive. This means that $geo(w_{|r}(S)$ does not intersect $\delta(T, \Pi \setminus p)$.

Since w is fair in Π, we can repeat this argument for any $p \in \Pi$. We denote the r^* the maximal such r and since $\bigcup_{p \in \Pi} \delta(T, \Pi \setminus p) = \delta(T)$, we have that $\delta(geo(w_{|s})(S)) \cap geo(w_{|r^*})(S) = \emptyset$. ◀

4.3 A Characterization of Geo-Equivalence

We start by simple, but useful, sufficient conditions about the size of geo-classes.

▶ **Proposition 21** ($Fair(\Pi)$ is separable). *Let $w \in IIS_n$, denote Σ the geo-class of w. If w is $Fair$ on Π, then $\#\Sigma = 1$.*

Proof. Let $w' \in \Sigma$. We will show that w' shares all prefixes of w. Let $r \in \mathbb{N}$. From Prop. 20 and Lemma 35, we get that $geo(w)$ does not belong to the boundary of $geo(w_{|r})(S)$, nor to the boundary of $geo(w'_{|r})(S)$. Assume that w and w' have not the same prefix of size r, that is $geo(w_{|r})(S) \neq geo(w'_{|r})(S)$. From Prop. 16 $geo(w_{|r})(S)$, $geo(w'_{|r})(S)$ have to intersect (as simplices), and since they are different, they can intersect only on their boundary. This means that $geo(w)$ would belong to the boundary, a contradiction.

So they have the same prefixes and $w' = w$. ◀

▶ **Proposition 22** (Infinite Cardinal). *Let $n \geq 2$. Let w, w' two distinct executions such that $geo(w) = geo(w')$ and there exist $s \in \mathbb{N}$ such that $\forall r > s \exists T geo(w_{|r})(S) \cap geo(w'_{|r})(S) = T$ with T a simplex of dimension $k \leq n-2$. Then, the geo-equivalence class of w is of infinite size.*

Proof. Let w, w' two executions with $geo(w) = geo(w')$ and $\forall r > s$, $geo(w_{|r})(S) \cap geo(w'_{|r})(S) = T$ with T of dimension $k \leq n-2$. Denote P the colors of T. Since $k \leq n-2$, we have at least $p_1 \neq p_2 \in \Pi \setminus P$. The suffix at length s of w is in $Solo(P)$.

Hence, for the processes in P, when running in w or w', it is not possible to distinguish these 3 cases about the induced subgraph by $\{p_1, p_2\}$ in the instant graphs: $p_1 \leftarrow p_2$, $p_1 \leftrightarrow p_2$ and $p_1 \rightarrow p_2$.

So $\forall r > s$, we have 3 possible ways of completing what is happening on the induced subgraph by processes in P in $G \in Solo(P)$. So we have infinitely many different executions, the cardinality of the geo-class of w is infinite. ◀

Let's consider the remaining cases. Let $w \in IIS_n$, denote Σ the geo-class of w.

▶ **Proposition 23** (Boundaries of $\mathbb{S}^n$ are separable). *If w is $Fair$ on $\Pi \setminus \{p\}$ for some $p \in \Pi$ then $\#\Sigma = 1$.*

Proof. We denote $Q = \Pi \setminus \{p\}$. We apply Prop. 21 for $n-1$ to w' the restriction of w to the set of processes Q (this is possible by definition of $Fair$: no process of Q receives message from outside of Q). Since w' satisfies the condition for Prop. 21 (by definition of Fair), which means that the geo-class of w' is of size 1.

Since there is only one unique way of completing an execution restricted to Q to one in $Solo(Q)$ (adding $(q,p), \forall q \in Q$), we get that there is only w in the equivalence class. ◀

A suffix of a word w is *strict* if it is not equal to w.

▶ **Proposition 24.** *If w has only a strict suffix that is $Fair$ on $\Pi \setminus \{p\}$ for some $p \in \Pi$ then $\#\Sigma = 2$.*

Proof. We denote $Q = \Pi \setminus \{p\}$. We can write $w = uav$ where $u \in IS_n^*$, $a \in IS_n$ and v is Fair on Q but av is not. We can choose u such that u has the shortest length.

We consider w' such that $geo(w') = geo(w)$. Let r be the length of ua. We denote by T the facet of $geo(ua)(\mathbb{S}^n)$ with colors Q. Since v is Fair for Q, we can apply to $v_{|Q}$ the restriction of v to Q Prop. 20. So $geo(w)$ is not on the boundaries of T which means, from Prop. 16, that either $geo(w'_{|r})(\mathbb{S}^n) = geo(w_{|r})(\mathbb{S}^n)$ either $geo(w'_{|r})(\mathbb{S}^n) \cap geo(w_{|r})(\mathbb{S}^n) = T$.

In both cases, we can apply Prop. 21 to v' the restriction of w to Q. Which means that there is only one restricted execution in Q. Since there is only one way to complete to p, there are as many elements in the class that simplices at round r that include T. Since we have a subdivision, we have exactly two simplices sharing the facet T.

In the first case, this means that $w'_{|r} = ua$ and $w = w'$.

In the second case, we have that $w'_{|r} = ub$ for some $b \neq a$. We remark that if $w'_{|r-1} \neq u$ this would contradict the minimality of u. Indeed, the prefixes of length $r-1$ are different, this means that av is Fair for Q. ◀

Using these previous propositions, and remarking that for any w, there exists P such that w has a suffix in $Fair(P)$, we can now present our main result regarding the complete classification of geo-equivalence classes. Let $n \in \mathbb{N}$ and Σ a geo-equivalence class on $\mathbb{S}^n$. Then there are exactly 3 cardinals that are possible for Σ (only 2 when $n = 1$, the case of [9]):

▶ **Theorem 25.** *Let $w \in IIS_n$, denote Σ the geo-class of w.*
$\mathcal{C}_1$: If w is $Fair$ on Π or on $\Pi \setminus \{p\}$ for some $p \in \Pi$, then $\#\Sigma = 1$;
$\mathcal{C}_2$: w has only a strict suffix that is $Fair$ on $\Pi \setminus \{p\}$ for some $p \in \Pi$ then $\#\Sigma = 2$;
$\mathcal{C}_\infty$: otherwise Σ is infinite.

5 The Set-Agreement Problem

For all n, the *set-agreement problem* is defined by the following properties [18]. Given initial *init* values in $[0, n]$, each process outputs a value such that

Agreement the size of the set of output values is at most n,
Validity the output values are initial values of some processes,
Termination All processes terminates.

We will consider in this part sub-IIS message adversaries $\mathcal{M}$, that is $\mathcal{M} \subseteq IIS_n$. It is well known that set-agreement is impossible to solve on IIS_n, we prove the following characterization.

▶ **Theorem 26.** *Let $\mathcal{M} \subset IIS_n$. It is possible to solve Set-Agreement on $\mathcal{M}$ if and only if $\wr\mathcal{M}\wr \neq |\mathbb{S}^n|$.*

5.1 Impossibility Result

On the impossibility side, we will prove a stronger version with non-silent algorithms. An algorithm is said to be *non-silent* if it sends message forever. Here, this means that a process could have decided a value while still participating in the algorithm.

▶ **Theorem 27.** *Let $\mathcal{M} \subset IIS_n$. If $\wr\mathcal{M}\wr = \wr IIS_n\wr = |\mathbb{S}^n|$ then it is not possible to solve Set-Agreement on $\mathcal{M}$, even with a non-silent algorithm.*

We will need the following definition from combinatorial topology.

▶ **Definition 28** (Sperner Labelling)**.** *Consider a simplicial complex C that is a subdivision of a chromatic simplex (S, χ). A labelling $\lambda : V(C) \longrightarrow \Pi$ is a* Sperner labelling *if for all $x \in V(C)$, for all $\sigma \subset S$, we have that $x \in |\sigma| \Rightarrow \lambda(x) \in \chi(\sigma)$.*

▶ **Lemma 29** (Sperner Lemma [26])**.** *Let a simplicial complex C that is a subdivision of a chromatic simplex (S, χ) with Sperner labelling λ. Then there exists $\sigma \in C$, such that* $\lambda(\sigma) = \Pi$.

A simplex σ with labelling using all Π colors is called *panchromatic*.

Proof of Theorem 27. By absurd, we assume there is a non-silent algorithm $\mathcal{A}$ (in full information protocol form) solving set-agreement on $\mathcal{M}$. We run the algorithm on initial inputs $init(i) = i$. We translate the full information protocol to the chromatic average, non-silent form: the initial value of i is x_i^*; when the decision value is given, we still compute and send the chromatic average forever. We can also assume a "normalized" version of the algorithm: when a process receives a decision value from a neighbour, it will decide instantly on this value. Such a normalization does not impact the correctness of the algorithm since set-agreement is a colorless task.

The proof will use the Sperner Lemma with labels obtained from the eventual decision value of the algorithm. However it is not possible to use directly the Sperner Lemma for the "full subdivision under $\mathcal{M}$" (which we won't define), since this subdivision could be infinite. The following proof will use König Lemma to get an equivalent statement.

Given $t \geq 0$, we consider $\mathrm{C}hr^t(\mathbb{S}^n)$ under our algorithm with initial values $init(i) = i$. For any vertex, we define the following labelling λ_t: if the process i has not terminated at time t with state $x \in V(\mathrm{C}hr^t(\mathbb{S}^n))$, then the Sperner label $\lambda_t(x) = i$, otherwise it is the decided value. Since the decided value depends only on the local state, the label of a vertex at time t is independent of the execution leading to it. The goal of the following is to show that there is an entire geo-equivalence class that does not belongs to $\mathcal{M}$.

By Integrity property, we have that the value decided on a face of $\mathbb{S}^n$ of processes $i_1, \ldots, i_n$, ie for $Solo(i_1, \ldots, i_n)^\omega$ are taken in $i_1, \ldots, i_n$. From Prop. 19, at any t, this labelling defines therefore a Sperner labelling of a (chromatic) subdivision of S.

We consider the set $\mathcal{S}$ of all simplices S of dimension n of $\mathrm{C}hr^t(\mathbb{S}^n)$, for all t. For a given t, from Sperner Lemma, there is at least a simplex of $\mathrm{C}hr^t(\mathbb{S}^n)$, that is panchromatic. There is therefore an infinite number of simplices S that are panchromatic in $\mathcal{S}$. We consider now $\mathcal{T} \subset \mathcal{S}$, the set of simplices $T \in \mathcal{S}$ such that there is an infinite number of panchromatic simplex S such that $|S| \subset |T|$. Note that T needs not be panchromatic. Since the number of simplices of $\mathrm{C}hr^t(\mathbb{S}^n)$ is finite for a given t, there is at least one simplex of $\mathrm{C}hr^t(\mathbb{S}^n)$ that is in $\mathcal{T}$. Therefore the set $\mathcal{T}$ is infinite.

We build a rooted-tree structure over $\mathcal{T}$: the root is $\mathbb{S}^n$ (indeed it is in $\mathcal{T}$), the parent-child relationship between T and T' is defined when $T' \in \mathrm{C}hr(T)$. We have an infinite tree with finite branching. By König Lemma, we have an infinite simple path from the root. We denote

T_t the vertex at level t of this path. We have $|T_{t+1}| \subset |T_t|$ and $(T_t)_{t\in\mathbb{N}}$ converges (same argument as the end of Section A.1) to some $y \in |\mathbb{S}^n|$. The increasing prefixes corresponding to T_t define an execution w of IIS_n.

We will now consider two different cases, not on the fact whether or not, $w \in \mathcal{M}$, but on the result of $\mathcal{A}$ on execution w.

For first case, assume that algorithm $\mathcal{A}$ has eventually decided on all processes on run w at some time t_0. Since it could be that $w \notin \mathcal{M}$, we cannot conclude yet. But since all processes have decided, they do not change their label in subsequent steps. By definition, $T_{t_0} = geo(w_{|t_0}(\mathbb{S}^n)$ contains an infinite number of panchromatic simplices, *i.e.* at least one. So the simplex $geo(w_{|t_0}(\mathbb{S}^n)$ is panchromatic. Hence any run with prefix $w_{|t_0}$ cannot be in $\mathcal{M}$, since $\mathcal{A}$ solves set-agreement on $\mathcal{M}$. Therefore $w' = w_{|t_0}\mathcal{K}_\Pi^\omega$ (where $\mathcal{K}_\Pi$ is the complete graph), is a fair execution that does not belong to $\mathcal{M}$. Its entire geo-equivalence class, which is a singleton, is not in $\mathcal{M}$.

The second case is when algorithm $\mathcal{A}$ does not eventually decide on all processes on run w. Therefore $w \notin \mathcal{M}$. Now we show that all elements w' of the geo-class of w are also not in $\mathcal{M}$. Assume otherwise, then $\mathcal{A}$ halts on w'. By Prop. 16, at any t, the simplex corresponding to $w'_{|t}$ intersects T_t on a simplex of smaller dimension whose geometrization contains y. Consider t_0 such that the execution has decided at this round for w'. Consider now T_{t_0+1}, it intersects the decided simplex of $Chr^{t_0}(\mathbb{S}^n)$ corresponding to w', which means that the processes corresponding to the intersection were solo in $w'(t_0+1)$ and in $w(t_0+1)$. When a process does not belong to a set of solo processes of the round, it receives all their values. So by normalization property of algorithm $\mathcal{A}$, this means that in T_{t_0+1}, all processes have decided. A contradiction with the fact that $\mathcal{A}$ does not decide on all processes on run w. ◀

This impossibility result means that there are many strict subsets $\mathcal{M}$ of IIS_n where it is impossible to solve set-agreement, including cases where $IIS_n \setminus \mathcal{M}$ is of infinite size.

5.2 Algorithms for Set-Agreement

In this section, we consider message adversaries $\mathcal{M}$ that are of the form $IIS_n \setminus geo^{-1}(y)$ for a given $y \in |\mathbb{S}^n|$. We note $w \in IIS_n$, such that $geo(w) = y$. We have $w \notin \mathcal{M}$. In other words, $\mathcal{M} = IIS_n \setminus \mathcal{C}$, where $\mathcal{C} = geo^{-1}(geo(w))$ is the equivalence class of w. We also denote $\sigma_y(r)$ the simplex $geo(w_{|r})(\mathbb{S}^n)$.

5.2.1 From Sperner Lemma to Set-Agreement Algorithm

Remark that the protocol complex at time r is exactly $Chr^r(\mathbb{S}^n)$, there is no hole “appearing” in finite time for such $\mathcal{M}$. From Sperner Lemma, any Sperner labelling of a subdivision of $\mathbb{S}^n$ admits at least one simplex that is panchromatic. In order to solve set-agreement, the idea of Algorithm 2 is to try to confine the panchromatic, problematic but unavoidable, simplex of $Chr^t(\mathbb{S}^n)$ to $\sigma_y(r)$. Since the geo-class of w is not in $\mathcal{M}$, any execution will eventually diverge from $\sigma_y(r)$ and end in a non panchromatic simplex. We now define a special case of Sperner labelling of the Standard Chromatic Subdivision that admits exactly one given simplex that is panchromatic.

We consider the generic colored simplex (S,χ) where $S = (x_0, \ldots, x_n)$ and coloring function χ, that could be different from $\mathcal{P}$. We consider labellings of subdivisions C of S.

▶ **Definition 30.** *Let $\tau \in C$ a subdivision of S. $f : V(T) \longrightarrow \Pi$ is a* Sperner $\tau-$panlabelling *if: f is a Sperner labelling of C; for all simplex $\sigma \in C$, $f(\sigma) = \Pi$ if and only if $\sigma = \tau$.*

▶ **Proposition 31.** *Let τ be a face of $Chr(S,\chi)$, there exists a $\tau-$panlabelling λ of $Chr(S,\chi)$.*

Algorithm 2 Algorithm $\mathcal{A}_w$ for process i.

```
1 x ← x*_i; r ← 0;
2 Loop while x ∈ V(geo(w_{|r})(S^n))
3     r ← r + 1;
4     Send((i, x));
5     V ←Receive() // set of all received messages;
6     d ← sizeof(V) − 1 // i receives d + 1 messages ;
7     x = (1 − d/(2d+1))/(d+1) x + Σ_{(j,x_j)∈V, j≠i} (1 + 1/(2d+1))/(d+1) x_j;
8 Output: Ψ_w(r)(x);
```

This technical proposition is proved in the appendix. Denote $\lambda_\tau(S,\chi)$ such a Sperner $\tau-$panlabelling of $Chr(S,\chi)$.

Before stating the algorithm, we show how to construct a sequence of panlabellings for $Chr^r(\mathbb{S}^n)$. Let $r \in \mathbb{N}$, we denote $\Psi_w(r)$ the following labelling defined by recurrence.

Intuitively, it is the following labelling. In $Chr^r(\mathbb{S}^n)$, we have $\sigma_y(r)$ that is panchromatic, all other simplices using at most n colors. In $Chr^{r+1}(\mathbb{S}^n)$, we label vertices that do not belongs to the subdivision of $\sigma_y(r)$ by the labels used at step r. In vertices from $Chr\sigma_y(r)$, we use $\lambda_{\theta(w(r+1))}$ the Sperner $\tau-$panlabelling associated with $\theta(w(r+1))$ to complete the labelling that uses at most n colors on a given simplex, except at $\sigma_y(r)$. In order to simplify notation, we also note λ_G the labelling $\lambda_{\theta(G)}$. Of course, we apply $\lambda_{w(r+1)}$ using as input (corner) colors, the colors from $\Psi_w(r)$. This way, on the neighbours of $\sigma_y(r)$ the labelling is compatible.

We denote $\gamma_r(x)$ the precursor of level r of $x \in V(Chr^{r+1}(\mathbb{S}^n))$, that is the vertex of $V(Chr^r(\mathbb{S}^n))$ from which x is originating.

▶ **Definition 32.** *We set $\Psi_w(1)(x) = \lambda_{w(1)}(\mathbb{S}^n, \mathcal{P})(x)$ for all $x \in V(Chr^r(\mathbb{S}^n))$, and for $r \in \mathbb{N}^*$*

$$\begin{aligned}\Psi_w(r+1)(p) \quad = \quad & \Psi_w(r)(\gamma_r(x)) \text{ if } x \notin |geo(w_{|r})(\mathbb{S}^n)| \\ & \lambda_{w(r+1)}(\Psi_w(r)(\sigma_y(r))(x) \text{ if } x \in |geo(w_{|r})(\mathbb{S}^n)|\end{aligned}$$

▶ **Proposition 33.** *For all r, $\Psi_w(r)$ is a Sperner $\sigma_y(r)-$panlabelling of $Chr^r(\mathbb{S}^n)$.*

Proof. The proof is done by recurrence. The case $r = 1$ is Prop 31. Assume that $\Psi_w(r)$ is a Sperner $\sigma_y(r)-$panlabelling of $Chr^r(\mathbb{S}^n)$.

Consider now $\Psi_w(r+1)$ for $Chr^{r+1}(\mathbb{S}^n)$. By construction and recurrence assumption, panchromatic simplices can only lay in $|\sigma_y(r)|$. Since $\lambda_{w(r+1)}$ is a Sperner panlabelling and that the corner colors for σ_y are taken from $\Psi_w(r)$, we have that $\sigma_y(r)$ is the only panchromatic simplex of $Chr^{r+1}(\mathbb{S}^n)$. ◀

We now prove the correctness of $\mathcal{A}_w$ presented in Algorithm 2. Consider an execution $v \in \mathcal{M}$. For Termination: since elements of the geo-class of w are not in $\mathcal{M}$, there exists a round r at which $v_{|r} \neq w'_{|r}$ for all $w' \in geo^{-1}(geo(w)$, *i.e.* the conditional at line 2 is false for all processes and the algorithm is terminating. For Agreement: when terminating at round r, i is not in $\sigma_y(r)$, by loop conditional, so since $\Psi_w(r)$ is only panchromatic on $\sigma_y(r)$, the number of decided values is less than n. Integrity comes from the fact that $\Psi_w(r)$ is a Sperner labelling.

5.2.2 Lower Bounds

It is possible to use the impossibility result to prove the following lower bound. Algorithm 2 is therefore optimal for fair w.

▶ **Theorem 34.** *Let $\mathcal{A}$ be an algorithm that solves set-agreement on $\mathcal{M} = IIS_n \backslash geo^{-1}(geo(w))$ with $w \in IIS_n$. Then, for any execution $v \in \mathcal{M}$, $t \in \mathbb{N}$, such that $v_{|t} = w'_{|t}$ for some $w' \in geo^{-1}(geo(w))$, $\mathcal{A}$ has not terminated at t.*

Proof. Suppose $\mathcal{A}$ has decided on all process at t, with $v_{|t} = w'_{|t}$ for some $w' \in geo^{-1}(geo(w))$. So $\mathcal{A}$ solve set-agreement on w'. A contradiction with Th. 27 since $\wr \mathcal{M} \cup \{w'\} \wr = |\mathbb{S}^n|$. ◀

6 Conclusion and Implications for Topological Methods

In this note, we have presented how to construct a topology directly on the set of executions of IIS_n the Iterated Immediate Snapshot message adversary. Though this is not a simple textbook topology as usual, since there are non-separable points, the properties we presented enables to fully understand it. As a important application on using the underlying geometrization mapping *geo*, we were able to characterize precisely for the first time general subsets of IIS_n where set-agreement is solvable and give a topological interpretation of this result.

We also believe this new approach could be successfully applied to other distributed tasks and distributed models. When considering the input complex embedded in $\mathbb{R}^N$, the geometrization topology could be applied on all simplices, in effect providing a new topological framework for so called protocol complex. This could be done by applying the Chromatic Average algorithm. This was not detailed here as we did not need it to investigate set-agreement. Moreover, note that this construction works also for any model of computation that corresponds to a mesh-shrinking subdivision.

This geometrization topology provides also a nice topological interpretation for the characterization theorem. In particular the topology as defined here is the coarsest topology such that the mapping *geo* is continuous. Therefore the impossibility theorem could also be stated, using the No-Retraction Theorem of standard topology [11, Cor. 2.15]: set-agreement is solvable on message adversary $\mathcal{M}$ only if there exists a continuous function $f : \mathcal{M} \longrightarrow \partial\mathbb{S}^n$, where $\mathcal{M}$ has the geometrization topology. This is interesting since $\partial\mathbb{S}^n$, the boundary of $\mathbb{S}^n$, is exactly the output complex of the set-agreement task.

It should also be noted that, since we do have non-separable sets in our setting, it shows that the standard abstract simplicial complexes approach is actually not always directly usable, since abstract simplicial complexes are known to have separable topology. It means that, for the first time, we have to primarily use the geometric version of simplicial complexes to fully investigate general distributed computability. We call the topology defined here the *geometrization topology* to emphasize this change of paradigm.

References

1 Yehuda Afek and Eli Gafni. *Asynchrony from Synchrony*, pages 225–239. Number 7730 in Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2013.

2 Hagit Attiya, Armando Castañeda, and Thomas Nowak. Topological characterization of task solvability in general models of computation. In Rotem Oshman, editor, *Proceedings of the 37th International Symposium on Distributed Computing (DISC'23)*, volume 281 of *LIPICS*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2023.

3 Elizabeth Borowsky and Eli Gafni. Generalized flp impossibility result for t-resilient asynchronous computations. In *STOC '93: Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 91–100, New York, NY, USA, 1993. ACM Press. doi:10.1145/167088.167119.

4 Elizabeth Borowsky and Eli Gafni. A simple algorithmically reasoned characterization of wait-free computation (extended abstract). In *Proceedings of the Sixteenth Annual ACM Symposium on Principles of Distributed Computing*, PODC '97, pages 189–198. ACM, 1997. doi:10.1145/259380.259439.

5 Arnaud Casteigts, Paola Flocchini, Walter Quattrociocchi, and Nicola Santoro. Time-varying graphs and dynamic networks. *Int. J. Parallel Emergent Distributed Syst.*, 27(5):387–408, 2012.

6 Bernadette Charron-Bost, Matthias Függer, and Thomas Nowak. Approximate consensus in highly dynamic networks: The role of averaging algorithms. In *ICALP (2)*, volume 9135 of *Lecture Notes in Computer Science*, pages 528–539. Springer, 2015.

7 Matthias Függer, Thomas Nowak, and Manfred Schwarz. Tight bounds for asymptotic and approximate consensus. *J. ACM*, 68(6):46:1–46:35, 2021.

8 Eli Gafni, Petr Kuznetsov, and Ciprian Manolescu. A generalized asynchronous computability theorem. In Magnús M. Halldórsson and Shlomi Dolev, editors, *ACM Symposium on Principles of Distributed Computing, PODC '14, Paris, France, July 15-18, 2014*, pages 222–231. ACM, 2014.

9 Emmanuel Godard and Eloi Perdereau. Back to the coordinated attack problem. *Math. Struct. Comput. Sci.*, 30(10):1089–1113, 2020.

10 Darald J. Hartfiel. Behavior in Markov set-chains. In Darald J. Hartfiel, editor, *Markov Set-Chains*, Lecture Notes in Mathematics, pages 91–113. Springer, Berlin, Heidelberg, 1998. doi:10.1007/BFb0094591.

11 Allen Hatcher. *Algebraic Topology.* Cambridge University Press, 2002.

12 Maurice Herlihy, Dmitry N. Kozlov, and Sergio Rajsbaum. *Distributed Computing Through Combinatorial Topology.* Morgan Kaufmann, 2013.

13 Maurice Herlihy, Sergio Rajsbaum, and Michel Raynal. Computability in distributed computing: A tutorial. *SIGACT News*, 43(3):88–110, 2012.

14 Maurice Herlihy and Nir Shavit. The topological structure of asynchronous computability. *J. ACM*, 46(6):858–923, 1999.

15 Dmitry N. Kozlov. *Combinatorial Algebraic Topology*, volume 21 of *Algorithms and computation in mathematics.* Springer, 2008.

16 Dmitry N. Kozlov. Chromatic subdivision of a simplicial complex. *Homology Homotopy Appl.*, 14(2):197–209, 2012. URL: http://projecteuclid.org/euclid.hha/1355321488.

17 Petr Kuznetsov, Thibault Rieutord, and Yuan He. An asynchronous computability theorem for fair adversaries. In Calvin Newport and Idit Keidar, editors, *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing, PODC 2018, Egham, United Kingdom, July 23-27, 2018*, pages 387–396. ACM, 2018. URL: https://dl.acm.org/citation.cfm?id=3212765.

18 Nancy A. Lynch. *Distributed Algorithms.* Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1996.

19 James R. Munkres. *Elements Of Algebraic Topology.* Addison Wesley Publishing Company, 1984.

20 Thomas Nowak, Ulrich Schmid, and Kyrill Winkler. Topological characterization of consensus under general message adversaries. In *PODC*, pages 218–227. ACM, 2019.

21 J.E. Pin and D. Perrin. *Infinite Words*, volume 141 of *Pure and Applied Mathematics.* Elsevier, 2004.

22 Sergio Rajsbaum. Iterated shared memory models. In Alejandro López-Ortiz, editor, *LATIN 2010: Theoretical Informatics*, pages 407–416, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.

23 Michel Raynal. *Set Agreement*, pages 1956–1959. Springer New York, New York, NY, 2016. doi:10.1007/978-1-4939-2864-4_367.

24 Thibault Rieutord. *Combinatorial characterization of asynchronous distributed computability. (Caractérisation combinatoire de la calculabilité distribuée asynchrone)*. PhD thesis, University of Paris-Saclay, France, 2018. URL: https://tel.archives-ouvertes.fr/tel-02938080.

25 M. Saks and F. Zaharoglou. "wait-free k-set agreement is impossible: The topology of public knowledge. *SIAM J. on Computing*, 29:1449–1483, 2000.

26 Emanuel Sperner. Neuer beweis für die invarianz der dimensionszahl und des gebietes. *Math. Sem. Univ. Hamburg*, 6:265–272, 1928.

A Geometrization Topology

A.1 Convexity and Metric Results

We present some metric results relating vertices of the iterated chromatic subdivision. In particular we prove that the sequences $geo(w_{|r})(S)$ converge to a point. This is related to the known fact that the standard chromatic subdivision is *mesh-shrinking* [12].

The following lemma comes from the convexity of the μ_G transforms.

▶ **Lemma 35.** *Let w a run, let $r, r' \in \mathbb{N}, r < r'$ then $|geo(w_{|r'})(\mathbb{S}^n)| \subset |geo(w_{|r})(\mathbb{S}^n)|$.*

Proof. Consider only one step. We have that $\frac{1-\frac{d}{2d+1}}{d+1} + d \times \frac{1+\frac{1}{2d+1}}{d+1} = \frac{1-\frac{d}{2d+1}+d+\frac{d}{2d+1}}{d+1} = 1$. So one step of the Chromatic Average gives, on each process, a linear combination with non-negative coefficients that sums to 1, it is therefore a barycentric combination on the points of the simplex at the beginning of the round. It is therefore a convex mapping of this simplex. Since composing convex mapping is also convex, and that $\mathbb{S}^n$ is a convex set, we get the result by recurrence. ◀

▶ **Lemma 36.** *There exists reals $0 < K' < K < 1$, such that for all G of IS_n, all $p, q \in V(S)$, $p', q' \in V(\mu_G(S))$, such that $\mathcal{P}(p) = \mathcal{P}(p')$ and $\mathcal{P}(q) = \mathcal{P}(q')$, we have*

$$K'||p - q|| \leq ||p' - q'|| \leq K||p - q||.$$

Proof. This is a consequence of μ_G transforms being convex when $G \in IS_n$. It corresponds to a stochastic matrix (non-negative coefficients and all lines coefficient sums to one) that is scrambling (there is a line without null coefficients) hence contractive. See *e.g.* [10, Chap. 1] for definitions and a proof for any given G of IS_n.

Then K (resp. K') is the largest (resp. smallest) such bounds over all $G \in IS_n$. ◀

While iterating the chromatic subdivision, we remark that the diameter of the corresponding simplices is contracting. From Lemma 36, we have

▶ **Lemma 37.** *Let S a simplex of $\mathbb{R}^{\mathbb{N}}$, then $diam(\mu_{G_r} \circ \mu_{G_{r-1}} \circ \cdots \circ \mu_{G_1}(S)) \leq K^r diam(S)$, where K is the constant from the previous lemma.*

Since the simplices are contracted by the μ_G functions, the sequence of isobarycenters of $(geo(w_{|r}(S))_{r \in \mathbb{N}^*}$ has the Cauchy property and this sequence is therefore convergent to some point $x \in \mathbb{R}^N$. Since the diameter of the simplices converges to 0, it makes senses to say that the limit of the simplices is the point x. Note that it would also be possible to formally define a metric on the convex subsets of $\mathbb{R}^N$ and consider the convergence of the simplices in this space.

A.2 Geometrization Topology vs Geometric Realization Topology

In this section, we provide an example of a simplicial complex whose topology as a geometric realization is different from the topology it has in the ambient R^N space, here with $N = 1$ but that can be generalized to any N. This is actually quite well known, see e.g. [15].

We consider $C = \{0\} \cup \{[\frac{1}{r+1}, \frac{1}{r}] \mid r \in \mathbb{N}^*\}$.

We denote $|C|$ the topological space of C defined as a geometric realization. The closed sets of $|C|$ are the sets F such that $F \cap S$ is closed (in $\mathbb{R}$) for all $S \in C$, see [19]. Therefore $|C|$ has two connected components. We have $F =]0,1]$ is closed in $|C|$ since $F \cap [\frac{1}{r+1}, \frac{1}{r}] = [\frac{1}{r+1}, \frac{1}{r}]$, hence is closed for all r. Moreover, $F \cap \{0\} = \emptyset$ which is also closed in $\mathbb{R}$. We also have that $\{0\}$ is closed in $|C|$, so C can be covered by two disjoint closed sets, it is not connected.

On the other end, at the set level, $\wr C \wr$ is exactly $[0, 1]$. So within the standard ambient topology of $\mathbb{R}$, $\wr C \wr$ is connected.

Since they do not have the same number of connected components, the two spaces C as a geometric realization and with the subset topology cannot be homeomorphic.

A full discussion of these differences could be very interesting. Given that the topologies are the same when the complex is finite, the question at stake seems to be the passage to the limit.

B Sperner Panlabellings of the Standard Chromatic Subdivision

In this section, n is fixed. We show how to construct a Sperner panlabelling of the standard chromatic subdivision. We consider the generic colored simplex (S, χ) where $S = (x_0, \ldots, x_n)$ and coloring function χ, that could be different from $\mathcal{P}$. We consider labellings of the colored complex $Chr(S, \chi)$.

We show the following combinatorial result about Sperner labellings.

▶ **Theorem 38.** *Let τ be a maximal simplex of $Chr(S, \chi)$, then there exists a $\tau-$panlabelling λ of $Chr(S, \chi)$.*

We start by some definitions related to proving the above theorem. It is possible to associate to any simplex σ of $Chr(S)$ a pre-order $\succ$ on Π that corresponds to the associate graph $\Theta(\sigma)$: $i \succ j$ when $(i, j) \in A(\Theta(\sigma))$. We call equivalence classes for $\Theta(\sigma)$, the classes of the equivalence relation defined by $i \succ j \wedge j \succ i$. It corresponds actually to the strongly connected components of the directed graph $\Theta(\sigma)$.

We define the *process view* of a point. This is the color of points in the view V of vertex (i, V) of the standard chromatic subdivision.

▶ **Definition 39** (Process View). *The* process view *of point $x = (\chi(x), V) \in V(\mathrm{C}hr(S, \chi))$ is defined by : $V_x = \{\chi(y) | y \in V\}$.*

For $\tau \in \mathrm{C}hr(S, \chi)$, we also define the *process view relative to τ* of a process p, denoted V_p^τ. It is the process view of the point of τ whose color is p. It is linked to pre-order $\succ$: we have $V_p^\tau = \{q \mid q \succ p\}$.

Let τ be a fixed maximal simplex of $Chr(S)$. We show how to construct a $\tau-$panlabelling. We choose a permutation φ on Π such that it defines circular permutations on the equivalent classes of $\Theta(\tau)$. Let $p \in \Pi$, given $W \subset V_p^\tau$ such that $p \in W$, we denote by $min^*(p, W) = \min\{i \in \mathbb{N}^* \mid \varphi^i(p) \in W\}$. Note that since φ is a permutation, there exists $j > 0$ such that $\varphi^j(p) = p$, and since $p \in W$, the minimum is taken over a non-empty set. Finally we set $\varphi^*(p, W) = \varphi^{min^*(p,W)}(p)$. This is the first point of W that is in the orbit of p in φ.

▶ **Definition 40.** *We define* $\lambda_\tau : V(Chr(S)) \to V(S)$*, for* $x \in V(Chr(S))$*, we set*
$$\lambda_\tau(x) = \begin{cases} q & \text{if } \exists q \in V_x \text{ and } q \notin V^\tau_{\chi(x)} \\ \varphi^*(\chi(x), V_x \cap V^\tau_{\chi(x)}) & \text{otherwise.} \end{cases}$$

Intuitively, for a given vertex of $Chr(S)$ with view V, if the process sees an other process q than in τ, then it is labelled by this q, otherwise it will choose the first process in the circular orbit of φ that is in its view.

▶ **Proposition 41.** *The labelling* λ_τ *is a* $\tau-$*panlabelling.*

Proof. First we show that it is indeed a Sperner labelling. In both cases of the definition, $\lambda_\tau(x)$ belongs to V_x. For $x \in V(Chr(S))$, for $\sigma \subset S$, $x \in |\sigma|$, with σ minimum for this property, means that the presentation of x is $\Phi(x) = (i, \sigma)$ for some x_i such that $x_i \in V(\sigma)$ and $\chi(x_i) = i$.

Now we show that the only panchromatic simplex is τ. By construction, with $x \in V(\tau)$, $\varphi^*(\chi(x), V_x) = \varphi(\chi(x))$ since in this case $V_x = V^\tau_{\chi(x)}$. So τ is panchromatic through λ_τ.

Now we consider $\sigma \neq \tau$. We have two possible cases:

1. $\exists x \in V(\sigma), q \in V_x, q \notin V^\tau_{\chi(x)}$,
2. $\forall x \in V(\sigma), V_x \subseteq V^\tau_{\chi(x)}$.

We start with the first case, we denote by C the highest, for $\succ$ in σ, class such there is x in C satisfying the clause (1). We show that $\#\lambda_\tau(C) \cap C < \#C$, where $\#$ denotes the cardinal of a set. By definition of C, $\lambda_\tau^{-1}(C) \subseteq C$. Since $\lambda_\tau(x) \notin C$, this means that $\#\lambda_\tau(C) \cap C \neq \#C$. By assumption all classes C' that are higher than C choose colors in C', so σ is not panchromatic under λ_τ.

Now, we assume we do not have case (1), this means that $\forall x \in V(\sigma), \lambda_\tau(x) = \varphi^*(\chi(x), V_x)$. Since $\sigma \neq \tau$, there exists $x \in V(\sigma), V_x \subsetneq V^\tau_{\chi(x)}$. We choose the lowest such x for $\succ$ in τ. We consider C_x the class of x in σ. We show that $\#\lambda_\tau(C_x) < \#C_x$.

We denote C^τ_x the class of color $\chi(x)$ in τ. First we show that $C_x \subseteq C^\tau_x$. Indeed, assume there is $y \in C_x$ such that $y \notin C^\tau_x$. Since the view of elements of the same class are the same, this means that $\chi(x) \in V_y$ and y would satisfy property 1. A contradiction to the case we are considering. And this is true for all $y \in C_x$.

Now we show $C_x \subsetneq C^\tau_x$. We have $V_x \subsetneq V^\tau_{\chi(x)}$. Let $y \in V^\tau_{\chi(x)} \setminus V_x$. If $y \notin C^\tau_x$, by the same previous argument, we get a contradiction. Hence $y \in C^\tau_x$ and therefore $C_x \subsetneq C^\tau_x$.

We denote $p = \varphi^*(\chi(x), C^\tau_x \setminus C_x)$. We note $p' = \varphi^{-1}(p)$. We have by definition of $\varphi^*(., C^\tau_x \setminus C_x)$, that $p' \in C_x$, therefore $C_{\chi^{-1}_{|\sigma}(p')} = C_x$. Now we set $p'' = \varphi^*(p', C_x)$. The color p'' has at least two predecessors in the labelling: p' by construction (since x was chosen the lowest for $\succ$ then $V_{\chi^{-1}_{|\sigma}(p')} = V^\tau_{p'}$) and $p''' = \varphi^{-1}(p'')$ which is not p' since $\varphi(p') = p \notin C_x$.

So $\#\lambda_\tau(\chi^{-1}_{|\sigma}(V^\tau_x)) < \#V^\tau_x$, and $\lambda_\tau(\sigma) \neq \Pi$. ◀