Stone Duality Proofs for Colorless Distributed Computability Theorems

Cameron Calk
Aix-Marseille University & CNRS UMR7020

Emmanuel Godard
Aix Marseille University & CNRS UMR7020

Abstract

We introduce a new topological encoding of executions of round-based, full-information distributed protocols via spectral spaces. Such protocols constitute a model of distributed computations which are functorially presented and englobe message adversaries. We give a characterization of the solvability of colorless tasks against compact adversaries. Message adversaries are distributed models that are known to be very expressive despite being round-based and crash-free. Colorless tasks are an important class of distributed tasks in which the specification does not depend upon the multiplicity of input or output values, examples thereof including the ubiquitous agreement tasks. Therefore, our result is a significant step toward unifying topological methods in distributed computing.

The main insight of this work is in considering global states obtained after finite executions of a distributed protocol not as abstract simplicial complexes as was previously done, but as spectral spaces, considering the Alexandrov topology on the associated face posets. Given an adversary $\mathcal{M}$ with a set of inputs $\mathcal{I}$, we define a limit object $\Pi^{\infty}_{\mathcal{M}}(\mathcal{I})$ by a projective limit in the category of spectral spaces. In some sense that is made explicit in this paper, this precisely encodes all distributed information about this adversary. Finally, this allows us to derive a new general distributed computability theorem using Stone duality: there exists an algorithm solving a colorless task $(\mathcal{I}, \mathcal{O}, \Delta)$ against the compact adversary $\mathcal{M}$ if, and only if, there exists a spectral map $f : \Pi^{\infty}_{\mathcal{M}}(\mathcal{I}) \to \mathcal{O}$ compatible with Δ.

From this general characterization, we then derive the known colorless computability theorems for the (colored or uncolored) Iterated Immediate Snapshot, while our method provides deep insights into the previously established connection between task-solvability and continuous maps between geometric realizations. Quite surprisingly, colored and uncolored models have the same distributed computability power, *i.e.* they solve the same tasks. Our new proofs give topological reasons for this equivalence, previously known through algorithmic reductions.

2012 ACM Subject Classification distributed algorithms, computability

Keywords and phrases simplicial complex, partial order, spectral spaces, Stone duality

1 Introduction

1.1 Topological methods for distributed computability

One of the most important results in distributed computing is the identification of finite combinatorial topology, in the form of simplicial complexes and their structure maps, as the natural setting of the subject along with the fact that, at least for a restricted computational protocol, continuous maps between the geometric realizations of associated simplicial complexes classify which tasks can be solved by the protocol. While many generalizations and applications have been developed, still to this day, the direct connection between distributed computing and continuous maps between geometric realizations was unclear. Here, we identify *spectral spaces*, arising as limits of the finite combinatorial topology models of distributed computing as the intrinsic phenomenon behind this remarkable connection. This additionally allows us to obtain a version of the original result which applies to a much wider class of protocols.

In a distributed system, several processes attempt to coordinate through some form of communication in order to solve a task. As opposed to sequential systems, in which computability can be characterized by universal objects such as Turing machines, task-solvability, *i.e.* computability, in the distributed case is limited by the ambiguity which arises from each process only having a local view of the global state of the system. Characterizing task-solvability for different models of communication has been a central question in distributed computing since its foundation. In ′99, Herlihy and Shavit [HS99], as well as Saks and Zaharoglou [SZ00], building on the work of Borowsky and Gafni [BG93a, BG97], were awarded the Gödel prize for showing that central problems in distributed computability can be understood and solved using topology.

Indeed, the possible epistemic states of the system fit nicely into combinatorial objects called simplicial complexes, which have a clear topological interpretation. Concretely, such an object consists of all possible global states of the system, as well as their subsets, thereby not only representing the combinatorics of possible configurations, but also encoding indistinguishability: a process cannot distinguish between two global states if it lies in their intersection. Such a structure, which can also be thought of as a space built by glueing together points, lines, triangles and their higher dimensional analogues, thus concisely encodes the epistemic ambiguity present in the system.

In sequential computing, a program specification consists of a relation between a set of inputs and a set of outputs. As described above, in a distributed system, inputs and outputs are encoded as simplicial complexes due to the multiplicity of processes. Taking this into account, a distributed task can be faithfully presented as a triple $(\mathcal{I}, \mathcal{O}, \Delta)$ where $\mathcal{I}$ and $\mathcal{O}$ are simplicial complexes, and Δ is a relation between global states in $\mathcal{I}$ and legal output states in $\mathcal{O}$ which respects the simplicial structure. On the other hand, a distributed protocol describes the modalities of communication, for example whether it be via message-passing or a shared-memory object, if it is synchronized or not, or whether processes can crash. This information can likewise be encoded as a relation between the input complex $\mathcal{I}$ and the so-called *protocol complex*, the latter consisting of all reachable global states from inputs in $\mathcal{I}$, glued along the indistinguishability relation, along with a relation specifying which global system states can be reached from which global input states. A protocol *solves* a task if each process can decide on an output value consistent with the specification Δ, which corresponds to the existence of a simplicial map, called the *decision map*, from the protocol complex to the output complex. This presentation of a distributed tasks and protocols provides a means of relating task-solvability, in the case of the Iterated Immediate Snapshot (IIS) model [BG97, HKR13], to topological information about the associated complexes.

In the IIS model, one round of computation corresponds topologically to a subdivision of the input complex, namely the standard chromatic subdivision Chr for the standard (colored) model of computation, *i.e.* in which processes have identifiers, and by the barycentric subdivision for the colorless version. The main result of Herlihy and Shavit [HS99] shows that a task $(\mathcal{I}, \mathcal{O}, \Delta)$ admits a solution in the IIS model if, and only if, there exists a simplicial map from $Chr^{\ell}\mathcal{I}$, the ℓ-fold chromatic subdivision, to $\mathcal{O}$, which respects the specification Δ. These results extend to the asynchronous wait-free read-write model, one of the fundamental models of distributed computing, since it was shown in [BG97], via algorithmic methods, that this model and IIS can solve the same class of tasks.

In this paper, we will focus on a specific subset of distributed tasks that is of particular interest for the distributed computing community: the colorless tasks. These include the *k-set agreement* tasks, in which output values must be amongst those that were initially present in the system, and such that there are at most k different output values. The case in which $k = 1$ is the *Consensus* task, while for $k + 1$ processes, this is called the *Set-Agreement* task, both being important primitives in distributed systems. Formally, a distributed task is said to be *colorless* if its specification does not depend of the multiplicity of values in the distributed system, meaning that we only have to encode the values present in the system, rather than also keeping track of process identifiers. For colorless tasks, the specification Δ is again a relation between simplicial complexes, but the computability theorems have simpler topological statements. For example, see [HKR13, Th. 4.3.1], a colorless task $(\mathcal{I}, \mathcal{O}, \Delta)$ is solvable in the IIS model if, and only if, there exists a continuous map $f : |\mathcal{I}| \to |\mathcal{O}|$ which respects Δ, where $|.|$ is the geometric realization. Using this topological characterization, the impossibility of $k-$set agreement is now a simple consequence of the No Retraction theorem from standard topology textbooks. This also explains why Saks and Zaharoglou [SZ00] were able to prove the impossibility of $k-$set agreement by a direct application of Sperner's Lemma, a known combinatorial topology presentation of the No Retraction theorem.

The main contribution of this paper is to extend this result to any colorless iterated protocol, rather than exclusively the IIS model. In the approach described above, this limitation arises as a result of using the geometric realization as the classifying object. Indeed, the topology of the geometric realization of the input complex is unchanged under successive subdivisions, and a decision map is obtained via simplicial approximation. Our approach avoids this limitation by using *spectral spaces* as classifying objects for colorless protocols. This method additionally provides a deep understanding of why geometric realizations provided a classification in the case of IIS.

1.2 Extending the topological framework to arbitrary models

Before detailing our contribution, we describe other lines of work extending the topological approach to distributed computing. As stated above, topological methods were pioneered by Herlihy and Shavit [HS99], Saks and Zaharoglou [SZ00], and Borowsky and Gafni [BG93b]. This yielded a first wave of applications that were collated in the book of Herlihy, Koslov and Rajsbaum [HKR13]. We refer the interested reader to this book for a complete introduction to topological methods in distributed computing, while a more concise account can be found in [Kei12].

Our contribution fits into the second wave of applications, extending the results to many other models than IIS, the asynchronous wait-free read-write model, or the t-resilient model (at most t processes can fail) which are considered in [HKR13]. As explained above, although this approach has provided many results, the topological arguments are restricted to the IIS model, while extensions to more general models such as the asynchronous wait-free read-write model rely on ad-hoc simulations.

Recently, there have been two lines of work towards a generalization of the topological arguments to more general models. In [NSW19, ACN23], it was proved that for arbitrary round-based models $\mathcal{M}$ that there exists a general topological space $\mathcal{E}_{\mathcal{M}}$ such that a task $(\mathcal{I}, \mathcal{O}, \Delta)$ is solvable against $\mathcal{M}$ if and only if there is a chromatic simplicial map $f : \mathcal{E}_{\mathcal{M}} \to \mathcal{O}$ [ACN23, Thm.5.4]. This space was first introduced for the Consensus task in [NSW19] as an infinite simplicial complex. On the other hand, in [GP20, CG23, CG24], the generalization is furthered by considering arbitrary subsets of IIS. Many classical asynchronous models can be presented this way, so this was a step towards a general computability theorem.

The first line of work is more general (more adversaries, all tasks) but has somewhat abstract topological statements, while the second line of work has simpler geometric statements but only for specific adversaries and colorless tasks. To see why, despite the generality of the approach, it is not straightforward to derive the results from the second line of work from the first, one could compare the statement for computability of set-agreement in sub-IIS models in [ACN23, Thm. 4.2] and [CG23, CG24, Thm. 25]. By showing that the classifying topological space can naturally be interpreted as a spectral space, and leveraging Stone duality, our contribution unifies both approaches and gets the best from both lines of work : general topological statements for distributed computability that are amenable to simple geometric proofs for colorless tasks. See also the discussion comparing precisely the spaces introduced in these previous works on Section 5.4

1.3 Our Contributions

In complement to the introduction of a new topological encoding for round-based full-information distributed computations and illustrating examples, we present the following contributions, which show the interest of our new technique. First, given any input complex $\mathcal{I}$, we associate a spectral space $\Pi^\infty(\mathcal{I})$ to any round-based full information model Π. This is accomplished abstractly in terms of a projective limit, but also concretely as a space of sequences, see Proposition 6. These spaces characterize computability, as is shown in Theorem 10, which states that a protocol solves a task $(\mathcal{I}, \mathcal{O}, \Delta)$ if, and only if, there exists a spectral map from $\Pi^\infty(\mathcal{I})$ to the output complex $\mathcal{O}$ which respects the specification Δ. Finally, we show that this approach is consistent with the original approach using subdivisions, since in this case the geometric subdivision is a distinguished subspace of $\Pi^\infty(\mathcal{I})$. This allows us to apply our generalized theorem to the aforementioned standard models of distributed computing : IIS for colored and colorless protocols and colorless tasks.

Our approach relies on the interpretation of protocols as endofunctors on the category of simplicial complexes. Recent preprints [GHK$^+$25, FHG25b, FHG25a] have also adopted a categorical approach to distributed systems, using algebraic and sheaf theoretic methods, respectively, to describe distributed systems in colored models. While our functorial approach, although limited for the time being to the colorless case, is similar to that described in [GHK$^+$25], the algebraic point of view developed therein is complemented via Stone duality with an associated topological point of view in the form of the classifying spectral space.

2 Preliminaries

Let X and Y be sets. Given a map $f : X \to Y$ we denote by f^{-1} the inverse image map associated to f, *i.e.* $f^{-1} : \mathcal{P}(Y) \to \mathcal{P}(X)$, where $\mathcal{P}(X)$ is the powerset of X. The forward image map associated to f is denoted by $f_* : \mathcal{P}(X) \to \mathcal{P}(Y)$. Given an element $x \in X$ and subsets $A \subseteq X$ and $B \subseteq Y$, we write $f(x) \in Y$ to denote the image of x, $f[A]$ or $f_*[A]$ for the forward image of A, and $f^{-1}[B]$ for the inverse image of B. This helps to notationally distinguish at which level we are considering f: on the set or on the power-set.

We use standard notation and terminology for order and duality theory, see e.g [DST19, GvG24] for relevant definitions. Given a poset P, we denote by $\mathcal{U}(P)$ and $\mathcal{D}(P)$ the lattices of up-sets and down-sets of P, respectively.

2.1 Simplicial complexes

For the purposes of this paper, we will consider a *simplicial complex* $\mathcal{C}$ over a finite set of vertices V to be a down-set in the poset $(\mathcal{P}(V) \setminus \{\emptyset\}, \subseteq)$. This is a simple reformulation of the classical definition of simplicial complexes. Elements $\sigma \in \mathcal{C}$ are called *simplices*, minimal elements are called *vertices* and downsets $A \subseteq \mathcal{C}$ are called *subcomplexes*. We denote by $V_\mathcal{C}$ the set of vertices of $\mathcal{C}$, and by $\mathcal{D}(\mathcal{C})$ its lattice of subcomplexes. We denote by $\dim(\sigma) := |\sigma| - 1$ the *dimension* of σ, and by $\dim(\mathcal{C}) = \max\{\dim(\sigma) \mid \sigma \in \mathcal{C}\}$ the dimension of a complex $\mathcal{C}$.

A *simplicial map* $f : \mathcal{C} \to \mathcal{C}'$ is given by a map $g : V_{\mathcal{C}} \to V_{\mathcal{C}'}$ such that $f = g_*$. Note that the dimension of $f(\sigma)$ may be of smaller than that of σ. We say that f is *rigid* when it preserves dimension. The category of simplicial complexes with simplicial maps is denoted by **sComp**.

In the topological approach to distributed computing, another type of map is often also considered. Given two simplicial complexes $\mathcal{C}, \mathcal{C}'$, a *carrier map* from $\mathcal{C}$ to $\mathcal{C}'$ is a map $\phi : \mathcal{C} \to \mathcal{D}(\mathcal{C}')$, such that for $\sigma \subseteq \tau$ in $\mathcal{C}$, we have $\phi(\sigma) \subseteq \phi(\tau)$. A carrier map is *rigid* when it preserves dimension, *i.e.* for a simplex σ of dimension d, $\phi(\sigma)$ is a subcomplex of $\mathcal{C}'$ of dimension d. We say that it is *strict* if it preserves intersections, *i.e.* $\phi(\sigma \cap \tau) = \phi(\sigma) \cap \phi(\tau)$. However, note that when ϕ is not strict, we still have $\phi(\sigma \cap \tau) \subseteq \phi(\sigma) \cap \phi(\tau)$.

Finally, we endow these structures with a topology. It is well known [GvG24, Section 2.2] that finite posets can be endowed with a topology given by up- and down-sets. This is known as the Alexandrov topology. Since simplicial complexes are a special kind of poset, we will consider them as finite topological spaces endowed with this topology. Explicitly, we consider the topology on $\mathcal{C}$ in which open sets are up-sets $U \in \mathcal{U}(\mathcal{C})$ and closed sets are down-sets $D \in \mathcal{D}(\mathcal{C})$. Equipped with this topology, $\mathcal{C}$ is a finite T_0 space.

2.2 Spectral spaces

Finite T_0 spaces are examples of special topological spaces called *spectral* or *Stone spaces*, see [DST19] for a very complete presentation of these spaces and associated results. Recall that in a topological space X, a subset $A \subseteq X$ is said to be *compact* if any open cover of A can be refined to a finite open cover[1]. In non-Hausdorff spaces, a compact set is not necessarily closed. For this reason, in such spaces we may consider the set $\mathcal{KO}(X)$ of *compact-open* subsets of X, *i.e.* those which are compact *and* open. A T_0, compact and sober topological space X is a *spectral space* when $\mathcal{KO}(X)$ is closed under finite intersections and is in addition a basis for its topology, *i.e.* its open sets are generated under finite intersections and arbitrary unions of elements of $\mathcal{KO}(X)$. A *proper* or *spectral* map is a continuous function $f : X \to Y$ between spectral spaces, such that for every $U \in \mathcal{KO}(Y)$, $f^{-1}[U] \in \mathcal{KO}(X)$. The category of spectral spaces and spectral maps is denoted by **Spec**. In the case of a finite spectral space, which corresponds precisely to a poset P endowed with the Alexandrov topology, we have $\mathcal{KO}(P) = \mathcal{U}(P)$, and spectral maps are simply order-preserving maps [GvG24, Section 2.2]. In other words, the category $\mathbf{Pos}_f$ of finite posets and order-preserving maps is a full sub-category of **Spec**.

2.3 Stone duality

In this paper we will make extensive use of Stone duality, which establishes a useful connection between distributive lattices and spectral spaces. Given a lattice L, recall that a *filter* $F \subseteq L$ is a non-empty up-set which is closed under meets. A filter P is *proper* when $P \neq L$ and is *prime* if it is proper and if, for all $a, b \in L$, $a \vee b \in P$ implies that $a \in P$ or $b \in P$. We denote by *pFilt*(L) the set of prime filters of L. Endowing *pFilt*(L) with the topology generated by the sets $\hat{a} = \{F \in X \mid a \in F\}$, for $a \in L$, we obtain a spectral space $St(L)$. Denoting by **DL** the category of distributive lattices and lattice homomorphisms, and by **Spec** the category

[1] Note that we do not require X to be Hausdorff.

of spectral spaces and spectral maps, we have a pair of contravariant functors between **DL** and **Spec**:

$$St(-) : \mathbf{DL}^{op} \to \mathbf{Spec} \qquad \text{and} \qquad \mathcal{KO}(-) : \mathbf{Spec}^{op} \to \mathbf{DL}$$

establishing a *duality* between **DL** and **Spec**. Concretely, this means that, given a spectral map $f : X \to X'$ between spectral spaces, we obtain a lattice homomorphism $f^{-1} : \mathcal{KO}(X') \to \mathcal{KO}(X)$ given by inverse image. Conversely, a lattice homomorphism $h : L \to L'$ yields a spectral map $St(h) : St(L') \to St(L)$, and these assignments are mutually inverse. For a more in-depth treatment of Stone duality, we refer the reader to [GvG24].

As mentioned above, a finite spectral space X is isomorphic to a poset P endowed with the Alexandrov topology, and we have $\mathcal{KO}(X) \cong \mathcal{U}(P)$. Conversely, given a finite distributive lattice L, there is a bijection between $pFilt(L)$ and the set $\mathcal{J}(L)$ of join-prime elements of L. Summing this up, in the finite case, Stone duality specializes to a duality between $\mathbf{Pos}_f$ and $\mathbf{DL}_f$, the category of finite distributive lattices. This is called *Birkhoff duality*, see [GvG24] for more information. Note however that since $\mathcal{D}(P)$ is order-dual to $\mathcal{U}(P)$, this duality also applies to down-set lattices.

2.4 Protocol complexes

As a standard approach from distributed computing [HKR13], given the set of finite executions of length r of a round-based adversary $\mathcal{M}$, we construct a specific simplicial complex called the *protocol complex* $\boldsymbol{P}^r_{\mathcal{M}}(\mathcal{I})$. Its set of vertices is the set of states that can be reached after r rounds, and a subset σ of $State^r$ is a simplex if it appears as a tuple of local states for a given initialization $\iota \in \mathcal{I}$ and scenario $w \in \mathcal{M}$, that is, it can be obtained from execution $\iota.w_{|r}$, where $w_{|r}$ stands for the prefix of size r of w.

For example, it is well known that for colored IIS_d, this defines the chromatic subdivision $Chr(\mathcal{I})$ [HKR13], and the barycentric subdivision for the colorless IIS. See examples in Appendix A.3.

3 Functorial approach for full-information models

Due to the choice of models we are considering, namely round-based full-information adversaries, we can reformulate the operational definition of distributed tasks and protocols in terms of categorical and topological constructions. This line has also been taken in the recent pre-print [GHK+25]. In this section, we describe the mathematical objects which we use to model distributed protocols.

3.1 Distributed models, protocols and adversaries

We are interested in distributed computation models that can be encoded as the action of a functor on an input complex, along with a natural transformation relating input states to reachable states. Intuitively this corresponds to models where the computations are round-based and admit full-information algorithms. The reader should note that there are distributed models that are not *defined* as round-based (like asynchronous models) but are *computationnaly equivalent* to round-based ones.

A second assumption is that the models are "full information" models. This is not defined formally here at the operational level. What this means intuitively is that it is possible to devise any distributed algorithm on this model from a universal full-information algorithm. This algorithm [HKR13, Section 2.3] has the following simple structure : send the entire local state at each round. Broadly speaking, this means gathering all possible information (since it is distributed, this could be only partial information about the current global state of the system), and then trying to tell everybody else, gathering new information and repeating until there is enough information at hand to output a final (local) value. In particular, this requires that the local memory is unbounded, since the local (unprocessed) information gathered will strictly increase at each round. Another very important point here is that, despite local information being partial, it is possible, for an external actor, to infer the global state of the distributed system only from the collection of the local states. Even if the shared object that enables communication between processes is formally not part of any local state, it is sufficient to know all the local states to recover the internal state of the shared object. We give a precise example of this phenomenon, namely a specific message adversary called Iterated Immediate Snapshot (IIS). It is well-known that this synchronous crash-free model is computationally equivalent to the asynchronous shared memory wait-free (*i.e.* any process could crash at any point of the execution) model. The Iterated Immediate Snapshot model was first introduced as a (shared) memory model and was later shown to be equivalent to a message adversary, first as tournaments and iterated tournaments [BG97, AG13], and then as this message adversary [HKR13, HRR12]. See also [Raj10] for a survey of the reductions involved in these layered models.

Another point in such full information models is that the information can be collected in a way that is independent of the actual content of that information. An example is that generally, the set of possible actions for an omission adversary does not depend on the actual content of the messages (but could depend on the sender). We call these actions "abstract actions", to distinguish them from the corresponding concrete actions, in which the content of messages is specified. We emphasize that the adversary, being adaptive, could still choose, among its possible concrete actions, to omit a message based on its actual content. This aspect of models is called "oblivious" in [GHK+25].

Since the functorial protocols as defined below are not as classical as message adversaries, we will restrict applications to this family of models, but our contribution can be applied to many other models. See Appendix A.4 for a full justification of this approach.

3.2 Protocols as endofunctors

For round-based full-information models (called FI models), the protocol complex is simply defined and appears as a functorial construction describing how the simplices (*i.e.* global states) of the input complex are modified by the action of the protocol in one round. This is proposed in [GHK+25, Section 4], without an associated natural transformation.

A *protocol*, is a pair (Π, π) where $\Pi : \mathbf{sComp} \to \mathbf{sComp}$ is a functor and π is a natural transformation between $U : \mathbf{sComp} \to \mathbf{Pos}_f$, the forgetful functor, and $\mathcal{D} \circ \Pi : \mathbf{sComp} \to \mathbf{Pos}_f$. Given a simplicial complex $\mathcal{C}$, this means that the component $\pi_{\mathcal{C}} : \mathcal{C} \to \mathcal{D}(\Pi(\mathcal{C}))$ is a strict carrier map, and is such that

$$\bigvee_{\sigma \in \mathcal{I}} \pi_{\mathcal{I}}(\sigma) = \Pi(\mathcal{I}), \tag{1}$$

these being the conditions on carrier maps seen in [HKR13]. We also require these maps to be order embeddings, as is justified below in Remark 4. Thus, similarly to [FHG25a, GHK+25], we consider the action of a protocol as an endofunctor, but additionally encode the associated carrier maps as a natural transformation. Recall that given any order preserving map $f : P \to \mathcal{D}(Q)$, where P and Q are posets, we obtain a join-preserving map, called the *join-lift* of f.

$$\begin{aligned} f^\vee :& \mathcal{D}(P) \longrightarrow \mathcal{D}(Q) \\ & A \longmapsto \bigvee\{f(a) \mid a \in A\}, \end{aligned}$$

Since $\pi_\mathcal{C}$ is strict and the equality (1) holds, it is routine to show that $\pi_\mathcal{C}^\vee$ is a lattice homomorphism. Since this holds for all $\mathcal{C}$, we can consider the components $\pi_{\Pi^n(\mathcal{C})}$ for each $n \in \mathbb{N}$, which to simplify notation below we denote by π_n, obtaining a sequence of maps

$$\mathcal{D}(\mathcal{C}) \xrightarrow{\pi_0^\vee} \mathcal{D}(\Pi(\mathcal{C})) \xrightarrow{\pi_1^\vee} \cdots \xrightarrow{\pi_2^\vee} \mathcal{D}(\Pi^n(\mathcal{C})) \xrightarrow{\pi_n^\vee} \mathcal{D}(\Pi^{n+1}(\mathcal{C})) \xrightarrow{\pi_{n+1}^\vee} \cdots, \tag{2}$$

in the category **DL**. For $n \leq n'$, we denote by $\pi_{n,n'}^\vee : \mathcal{D}(\Pi^n(\mathcal{C})) \to \mathcal{D}(\Pi^{n'}(\mathcal{C}))$ the composition $\pi_n^\vee \circ \cdots \circ \pi_{n'-1}^\vee$.

3.3 The example of message adversaries

As a specific FI model, we consider the classical general message adversary models. We sketch here the main definitions, full operational definitions and other examples are detailed in Appendix A.

For $d \in \mathbb{N}$, we denote by $[d] = \{0, .., d\}$ the set of processes. We use standard directed graph (or digraph) notations : given G, $V(G)$ is the set of vertices, $A(G) \subseteq V(G) \times V(G)$ is the set of arcs. 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 $[d]$. We also denote by $\mathbf{G}(r)$ the digraph G_r. A *message adversary* is a set of dynamic graphs.

Intuitively, the arcs (u, v) of $\mathbf{G}(r)$ describe the transmission of messages from u to v at round r. If there is no arc from u to v, no information is transmitted. Note that u is not a priori aware whether a given broadcast message has been received. A formal definition of an execution under a dynamic graph is given in Section A.2.

We will use the following standard notations in order to more easily describe our message adversaries [PP04]. A dynamic graph is seen as an infinite word over the alphabet $\mathcal{G}_d$, set of all graphs with vertices $[d]$.Given $U \subseteq \mathcal{G}_d$, U^* is the set of all finite sequences of elements of U, U^ω is the set of all infinite sequences, and $U^\infty = U^* \cup U^\omega$.

An adversary $\mathcal{M}$ is said to be *compact* (in the sense of the profinite topology on infinite words [PP04]) if for any word $\mathbf{G} \in \mathcal{G}_d^\omega$, $\mathbf{G}_{|r} \in Pref_r(\mathcal{M})$ for all $r \in \mathbb{N}$ implies $\mathbf{G} \in \mathcal{M}$, where $Pref_r(\mathcal{M})$ is the set of prefixes of length r of $\mathcal{M}$.

We show how standard fault environments are conveniently described in our framework. We fix an integer d.

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

Given a graph G, we denote by $In_G(a) = \{b \in V(G) \mid (b, a) \in A(G)\}$ the set of incoming vertices of a in G. A graph G has the *containment property* if for all $a, b \in V(G)$, $In_G(a) \subseteq In_G(b)$ or $In_G(b) \subseteq In_G(a)$.

▶ **Definition 2** (Chap. 14.2 in [HKR13]). *We set $S_d = \{G \in \mathcal{G}_d \mid G$ has the Containment property$\}$. The Iterated Snapshot message adversary for $d+1$ processes is the message adversary $IS_d = S_d^\omega$.*

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 $(a, c) \in A(G)$.

▶ **Definition 3** (Chap. 4.3 and 14 in [HKR13]). *We set* $ImS_d = \{G \in \mathcal{G}_d \mid G$ *has the Immediacy and Containment properties* $\}$. *The Iterated Immediate Snapshot message adversary for* $d+1$ *processes is the message adversary* $IIS_d = ImS_d^{\omega}$.

There exists variant for these models, the colored and colorless variants. When, with the `Receive` primitive, a process is able to distinguish messages with the same content arriving from different processes, this is the (standard) colored version. This amounts to processes having unique identities. When the value returned by the `Receive` primitive is actually a set of the broadcast values, which means that a process is not able to distinguish messages with the same content arriving from different processes – including itself –, this is the colorless variant. Note that any homonymous model (processes have labels that may not be unique) can also be expressed in that setting.

The one round complex of the colored IIS (resp. colorless) protocol is the standard chromatic subdivision (resp. barycentric subdivision). The protocol complex for the colored IS model is describe in Appendix A.3. Finally, the corresponding functors can be recovered by a Yoneda extension argument, see Appendix A.4 below, or [GHK+25, Section 4.1] for a similar argument.

3.4 From endofunctors to inverse limit systems

The main insight of this paper is to view a simplicial complex as a spectral space, that is, in contrast to the usual approach, not as a combinatorial object nor a Hausdorff topological space built from higher-dimensional triangles. Rather, we view a simplicial complex as a poset endowed with the Alexandrov topology. This establishes the appropriate relationship between simplicial complexes and their lattices of sub-complexes, thereby also simplifying the relationship between carrier maps and simplicial maps. This construction can be taken to the (co-)limit, producing a general spectral space and a dual distributive lattice associated to a given protocol.

We fix a simplicial complex $\mathcal{I}$ for the remainder of this section.

Applying Birkhoff duality to the sequence (2), we deduce the existence of order preserving maps $\rho_n : \Pi^{n+1}(\mathcal{I}) \to \Pi^n(\mathcal{I})$ on the underlying simplicial complexes. This gives rise to an inverse limit system

$$\mathcal{I} \xleftarrow{\rho_0} \Pi(\mathcal{I}) \xleftarrow{\rho_1} \Pi^2(\mathcal{I}) \xleftarrow{\rho_2} \cdots \xleftarrow{\rho_{n-2}} \Pi^{n-1}(\mathcal{I}) \xleftarrow{\rho_{n-1}} \Pi^n(\mathcal{I}) \xleftarrow{\rho_n} \cdots \tag{3}$$

in the category of spectral spaces, since each $\Pi^n(\mathcal{I})$ is a spectral space when endowed with the Alexandrov topology. Results from the theory of spectral spaces [Hoc69] tell us that the limit of this diagram, which we denote by $\Pi^\infty(\mathcal{I})$, is a spectral space. We denote by $\rho_{n',n} : \Pi^{n'} \to \Pi^n$ the composition $\rho_{n'} \circ \cdots \circ \rho_{n-1}$, *i.e.* the dual of $\pi^\vee_{n,n'}$, and by $p_n : \Pi^\infty(\mathcal{I}) \to \Pi^n(\mathcal{I})$ the canonical projections associated with the limit construction.

▶ Remark 4. In terms of message adversaries, a word $G_1, G_2, \ldots, G_n, \ldots$ corresponds to a sequence $\sigma_1, \sigma_2, \ldots, \sigma_n, \ldots$ of maximal simplices, with $\sigma_i \in \Pi^i(\mathcal{I})$. The maps ρ_i encode which sequences of simplices are accepted by the message adversary, in the sense that the associated sequence $G_1, G_2, \ldots, G_n, \ldots$ is permitted by the message adversary precisely when $\rho_i(\sigma_{i+1}) = \sigma_i$. For this reason, it is reasonable to consider these as surjective maps. This justifies the specification that the maps $\pi^\vee_n$ are injective for a protocol (Π, π).

Moreover, since we consider up-sets as open sets, the colimit of the dual diagram (2) which we denote by $\mathcal{D}^\infty(\mathcal{I})$, is the lattice of closed sets of $\Pi^\infty(\mathcal{I})$ whose complements are compact, which generate its closed sets under arbitrary intersections and finite unions. We denote by $i_n : \mathcal{D}(\Pi^n(\mathcal{I})) \to \mathcal{D}^\infty(\mathcal{I})$ the canonical maps associated with the colimit.

Summing this up, we obtain our first theorem.

▶ **Theorem 5.** *Given a protocol* (Π, π)*, we have functors*

$$\Pi^\infty : \mathbf{sComp} \to \mathbf{Spec} \qquad \textit{and} \qquad \mathcal{D}^\infty : \mathbf{sComp}^{op} \to \mathbf{DL},$$

such that for any simplicial complex $\mathcal{I}$*,* $\mathcal{D}^\infty(\mathcal{I})$ *generates the closed sets of* $\Pi^\infty(\mathcal{I})$*.*

Proof. The assignments $\mathcal{I} \mapsto \mathcal{D}^\infty(I), \Pi^\infty(I)$ are well defined by co-completeness of $\mathbf{DL}$ and since directed limits of finite T_0 spaces are spectral (see [Hoc69], Proposition 10), respectively. Functoriality is a consequence of the universal property of (co-)limits. Explicitly, for a simplicial map $\phi : \mathcal{I} \to \mathcal{I}'$, we deduce morphisms $\Pi^n\phi : \Pi^n(\mathcal{I}) \to \Pi^n(\mathcal{I}')$ meaning that $\Pi^\infty(\mathcal{I})$ is a cone under $(\mathcal{I}' \leftarrow \Pi(\mathcal{I}') \leftarrow \cdots)$. By the universal property of limits, we obtain a unique arrow $\Pi^\infty\phi : \Pi^\infty(\mathcal{I}) \to \Pi^\infty(\mathcal{I}')$. For $\mathcal{D}^\infty$ the proof is similar except for that the dual universal property for colimits implies contravariance.

Under Stone duality, the dual lattice to $\Pi^\infty(\mathcal{I})$ is its set of compact-opens. In terms of colimits, this means we are considering the spectral space dual to the colimit $\mathcal{U}^\infty(I)$ of the sequence

$$\mathcal{U}(I) \longrightarrow \mathcal{U}(\Pi(I)) \longrightarrow \cdots \longrightarrow \mathcal{U}(\Pi^n(I)) \longrightarrow \mathcal{U}(\Pi^{n+1}(I)) \longrightarrow \cdots$$

where $\mathcal{U}(\Pi^n(I))$is the up-set lattice of $\Pi^n(I)$ and the maps are those induced from (3) by Birkhoff duality. Since each $\mathcal{D}(\Pi^n(I))$ consists of the complements of $\mathcal{U}(\Pi^n(I))$, we see that this means that $\mathcal{D}^\infty(\mathcal{I})$ is the set of closed sets in $\Pi^\infty(\mathcal{I})$ whose complements are compact-open, and thus generates its closed sets. ◀

3.5 The spectral space of a protocol

Here we describe the spectral space $\Pi^\infty(\mathcal{I})$ in terms of executions of the protocol. Indeed, we show that points in $\Pi^\infty(\mathcal{I})$ correspond to sequences of simplices associated with executions, and that the topology on $\Pi^\infty(\mathcal{I})$ characterizes task solvability, *cf.* Theorem 10.

To this end, we consider sequences $(\sigma_n)_n$ of simplices such that for all n,

- $\sigma_n \in \Pi^n(\mathcal{I})$,
- $\sigma_{n+1} \in \pi_n(\downarrow \sigma_n)$, or equivalently, $\rho_n(\sigma_{n+1}) = \sigma_n$.

We call these *protocol sequences*. Since $\Pi^\infty(\mathcal{I})$ is the limit of the sequence (3), we have [DST19, Section 2.3.9] that the underlying set of $\Pi^\infty(\mathcal{I})$ is

$$\{(\sigma) \in \prod \Pi^n(\mathcal{I}) \mid \rho_n(\sigma_{n+1}) = \sigma_n\},$$

that is, the set of protocol sequences. By [DST19], the specialization order $\preceq$ on $\Pi^\infty(\mathcal{I})$ is determined point-wise. More explicitly, given protocol sequences (σ) and (τ), we have $(\sigma) \preceq (\tau)$ if, and only if, for all n $\sigma_n \preceq_n \tau_n$, where $\preceq_n$ is the specialization order in the finite T_0 space Π^n. The latter is given by $\sigma_n \preceq_n \tau_n \iff \uparrow\sigma_n \subseteq \uparrow\tau_n$, whereby we deduce $(\sigma) \preceq (\tau)$ if, and only if, $\sigma_n \geq \tau_n$ for all n. Finally, we have the following characterization of $\Pi^\infty(\mathcal{I})$:

▶ **Proposition 6.** *Given a protocol* (Π, π)*, we obtain a functor* $\Pi^\infty : \mathbf{sComp} \to \mathbf{Spec}$ *sending a simplicial complex* $\mathcal{I}$ *to the spectral space* $\Pi^\infty(\mathcal{I})$*, such that*

- *The points of* $\Pi^\infty(\mathcal{I})$ *are in one-to-one correspondence with protocol sequences* (σ)*,*
- *The open topology is generated by the sets* $U_{n,v} := \{(\tau) \in \Pi^\infty(\mathcal{I}) \mid \tau_n \geq \sigma\} = p_n^{-1}(\uparrow v)$*.*
- *The specialization order* $\leq$ *on* $\Pi^\infty(I)$ *is defined by* $x_\sigma \leq x_\tau \iff \forall n, \sigma_n \geq \tau_n$, i.e. *the opposite order of the point-wise order on protocol sequences.*

4 Colorless Computability

4.1 Generalized colorless computability

We consider a fixed number $d+1$ of processes. A colorless task is given by a finite simplicial complex $\mathcal{I}$ of initial configurations, a simplicial complex of output values $\mathcal{O}$ and a relation Δ, the specification, between initial configurations and output configurations. It is required, see [HKR13, Section 4.1.4], that Δ be a order-preserving mapping from $\mathcal{I}$ (seen as a poset of faces) to the set of subcomplexes of $\mathcal{O}$, $\mathcal{D}(\mathcal{O})$.

Formally, a *colorless distributed task* is a triple $(\mathcal{I}, \mathcal{O}, \Delta)$, where $\mathcal{I}$ and $\mathcal{O}$ are simplicial complexes, and $\Delta : \mathcal{I} \to \mathcal{D}(\mathcal{O})$ is a carrier map. This is equivalent to [HKR13], since for a colorless task, given any simplex in the output complex, a sub-simplex is also legal. A protocol *solves* a task $(\mathcal{I}, \mathcal{O}, \Delta)$ if there exists an $n \in \mathbb{N}$ and a spectral map $\delta : \Pi^n(\mathcal{I}) \to \mathcal{O}$, such that for all $\sigma \in \Pi^n(\mathcal{I})$, $\delta(\sigma) \subset \Delta(\rho_{0,n}(\sigma))$. This is equivalent to the classical definition (see [HKR13]) with the following relaxation : a process can output a simplex of $\mathcal{O}$. Classically, it is said that a process *decides* when it outputs a value in $V(\mathcal{O})$. However, as is shown in Lemma 8, we can anyway get an output that is a vertex of $\mathcal{O}$. So this is a flexibility of our framework that actually does not change the computability. Finally, in our setting, solving is to say that the diagram on the left sub-commutes, *i.e.* for all $\sigma \in \mathcal{I}$, $\delta_*[\pi_{0,n}(\sigma)] \leq \Delta(\sigma)$.

$$\begin{tikzcd} \mathcal{I} \arrow[rr, "\Delta"] \arrow[dr, "\pi_{0,n}"'] & & \mathcal{D}(\mathcal{O}) \\ & \mathcal{D}(\Pi^n(I)) \arrow[ur, "\delta_*"'] & \end{tikzcd} \qquad \begin{tikzcd} \mathcal{D}(\mathcal{I}) \arrow[rr, "\Delta^\vee"] \arrow[dr, "\pi^\vee_{0,n}"'] & & \mathcal{D}(\mathcal{O}) \arrow[dl, "\delta^{-1}"] \\ & \mathcal{D}(\Pi^n(I)) & \end{tikzcd} \tag{4}$$

Unwinding the definitions of the join-lift of a map and using Birkhoff duality, one easily proves that the diagram on the left sub-commutes if, and only if, the diagram on the right sub-commutes, *i.e.* for $A \in \mathcal{D}(\mathcal{I})$, $\pi^\vee_{0,n}[A] \leq \delta^{-1}[\Delta^\vee[A]$. Note that the right-hand diagram is not in **DL** since $\Delta^\vee$ only preserves joins in general.

▶ Remark 7. Note that, contrary to the presentation in [HKR13], it is not required here that δ be simplicial. It is sufficient to have δ a spectral map since, as underlined above, any sub-simplex of a legal output simplex is also legal for a colorless task.

Actually, if Δ is rigid (it associates only vertices to vertex of $\mathcal{I}$), this will force δ to be simplicial. However, it is possible to have general Δ, that is task definitions where a vertex is associated to simplices of higher dimensions.

Given a task $(\mathcal{I}, \mathcal{O}, \Delta)$ and a protocol (Π, π), we say that a spectral map $f : \Pi^\infty(\mathcal{I}) \to \mathcal{O}$ is *carried* by Δ provided that the following diagram sub-commutes, that is, $i_0[A] \leq f^{-1}[\Delta^\vee[A]]$.

$$\begin{tikzcd} \mathcal{D}(\mathcal{I}) \arrow[rr, "\Delta^\vee"] \arrow[dr, "i_0"'] & & \mathcal{D}(\mathcal{O}) \arrow[dl, "f^{-1}"] \\ & \mathcal{D}^\infty(\mathcal{I}), & \end{tikzcd}$$

4.2 Colorless Computability theorem

We can now state our general computability theorem. First, we prove the following lemma:

▶ **Lemma 8.** *Given a a spectral map $f : \Pi^\infty(\mathcal{I}) \to \mathcal{O}$, there exists n such that f factorizes as $\Pi^\infty(\mathcal{I}) \xrightarrow{p_n} \Pi^n(\mathcal{I}) \xrightarrow{\delta} \mathcal{O}$, where δ is a simplicial map, and p_n is the canonical projection.*

Proof. This essentially follows from compactness of $\Pi^\infty(\mathcal{I})$. Indeed, note that the family $(\uparrow w)_{w\in V_\mathcal{O}}$ is a covering of $\mathcal{O}$, so $(f^{-1}(\uparrow w))_{w\in V_\mathcal{O}}$ is a covering of $\Pi^\infty(\mathcal{I})$. Each of these sets being open, and since by Propostion 6 the $(U_{n,v})_{v\in\Pi^n(\mathcal{I})}$ are a basis for the open sets of $\Pi^\infty(\mathcal{I})$, we have $f^{-1}(\uparrow w) = \bigcup_{(n,v)\in P_w} U_{n,v}$, where $P_w = \{(n,v) \mid p_n^{-1}(\uparrow v) \subseteq f^{-1}(\uparrow w)\}$. Writing $P = \bigcup_{w\in V_\mathcal{O}} P_w$, we thus have $\Pi^\infty(\mathcal{I}) = \bigcup_{(n,v)\in P} U_{n,v}$. By compactness, we have a finite sub-cover. Furthermore, we can choose a uniform n to consider since $\mathcal{D}(\Pi^n(\mathcal{I}))$ is the union of the $\mathcal{D}(\Pi^{n'}(\mathcal{I}))$ for $n' \leq n$. Explicitly, these means that we have an n and $v_1,\ldots,v_k \in V_{\Pi^n(\mathcal{I})}$, such that $p_n^{-1}(\uparrow v_i) \subseteq f^{-1}(\uparrow w)$ for some $w \in V_\mathcal{O}$, and $\Pi^\infty(\mathcal{I}) = \bigcup_{i=1}^k p_n^{-1}(\uparrow v_i)$. However, notice that for this to hold, $v_1,\ldots,v_k$ must be all of the vertices of $\Pi^n(\mathcal{I})$. Summing up, for all $v \in V_{\Pi^n(\mathcal{I})}$, there exists some $w_v \in V_\mathcal{O}$ with $p_n^{-1}(\uparrow v) \subseteq f^{-1}(\uparrow w_v)$. The assignment $v \mapsto w_v$ defines a simplicial map δ, and by construction, f factorizes as desired. ◀

▶ Remark 9. The above proof is similar to the proof using simplicial approximation found in [HKR13, Proposition 3.7.3], generalized to the case of spectral maps. Indeed, the conclusion is reached by showing that f satisfies the so-called *star condition*, namely that there exists n such that for all $v \in V_{\Pi^n(\mathcal{I})}$, there exists $w \in V_\mathcal{O}$ such that $f(p_n^{-1}(\uparrow v)) \subseteq \uparrow w$.

▶ **Theorem 10.** *A protocol (Π,π) solves a colorless distributed task $\Delta : \mathcal{I} \to \mathcal{D}(\mathcal{O})$ if, and only if, there exists a spectral map $f : \Pi^\infty(\mathcal{I}) \to \mathcal{O}$ carried by Δ.*

Proof. We take as definition of solving a task the sub-commutation of the right-hand diagram from (4). The diagram below proves that if (Π,π) solves the protocol, then there exists a spectral map $f : \Pi^\infty(\mathcal{I}) \to \mathcal{O}$ carried by Δ.

$$\begin{tikzcd}
\mathcal{D}(\mathcal{I}) \arrow[rr, "\Delta^\vee"] \arrow[drr, "\pi^\vee_{0,n}"] \arrow[ddr, "i_0"'] & & \mathcal{D}(\mathcal{O}) \arrow[d, "\delta^{-1}"] \\
& & \mathcal{D}(\Pi^n(I)) \arrow[dl, "i_n"] \\
& \mathcal{D}^\infty(\mathcal{I}) &
\end{tikzcd}$$

Indeed, writing $f^{-1} = i_n \circ \delta^{-1}$, we see that the outer diagram sub-commutes since the upper triangle sub-commutes and the lower triangle commutes on the nose as a result of the limit construction.

Conversely, given a spectral map $f : \Pi^\infty(\mathcal{I}) \to \mathcal{O}$ and applying Lemma 8, we obtain the same diagram as above. Since the lower triangle commutes on the nose, the upper triangle sub-commutes under the hypothesis that f is carried by Δ. ◀

5 Applications to IIS models

Here we describe the applications of spectral topology and duality theory to distributed computing. This culminates in proving, via novel methods, the following classical results:

▶ **Theorem 11** (Colorless Distributed Computability Theorem for colorless IIS)**.** *[HKR13, Thm. 4.2.3] Let $(\mathcal{I},\mathcal{O},\Delta)$ be a colorless task. It is solvable by a colorless immediate iterated snapshot protocol for $n+1$ processes if and only if there exists a continuous map $f : |\mathcal{I}| \to |\mathcal{O}|$ carried by Δ.*

▶ **Theorem 12** (Colorless Distributed Computability Theorem for (colored) IIS)**.** *[HKR13, Thm. 4.3.1] Let $(\mathcal{I},\mathcal{O},\Delta)$ be a colorless task. It is solvable by a (colored) immediate iterated* *snapshot protocol for $n+1$ processes if and only if there exists a continuous map $f : |\mathcal{I}| \to |\mathcal{O}|$ carried by Δ.*

▶ **Theorem 13** (Strong Colorless Distributed Computability Theorem for IIS)**.** *Let $(\mathcal{I}, \mathcal{O}, \Delta)$ be a colorless task. It is solvable by a colorless immediate iterated snapshot protocol if and only if it is solvable by a colored immediate iterated snapshot protocol.*

The third theorem is a consequence of the previous two. It is demonstrated in a direct, algorithmic way in [HKR13, Chap. 4].

5.1 Mesh-shrinking subdivision operators

Here we refine our characterization of spectral spaces obtained from protocols in the case of mesh-shrinking protocols, *i.e.* those for which the associated endofunctor on simplicial complexes Π is a mesh-shrinking subdivision operator. Indeed, in this case, the spectral space contains the geometric realization of the considered input complex $\mathcal{I}$ as its subspace of maximal points.

A *mesh-shrinking subdivision operator* is a functor $\Pi : \mathbf{sComp} \to \mathbf{sComp}$ such that

- $|\Pi(\mathcal{C})|$ is homeomorphic to $|\mathcal{C}|$ (*subdivision*),
- The diameter of realizations of protocol sequences vanish (*mesh-shrinking*), that is, $diam(|\sigma_n|) \overset{n\to\infty}{\longrightarrow} 0$ where, for $A \subseteq |\mathcal{C}|$, $diam(A) = \max\{d(x, y) \mid x, y \in A\}$.

Now, consider a protocol (Π, π) where Π is a mesh-shrinking subdivision operator. We write $R_n = \{|A| \mid A \in \mathcal{D}(\Pi^n(\mathcal{I}))\}$ for the lattice of realizations of elements of $\mathcal{D}(\Pi^n(\mathcal{I}))$. Since $|A \cup B| = |A| \cup |B|$ and $|A \cap B| = |A| \cap |B|$, we have $R_n \cong \mathcal{D}(\Pi^n(\mathcal{I}))$. Abusing notation, we will write $\pi_n : R_n \to R_{n+1}$. The colimit of these maps gives a basis for the closed sets of $|\mathcal{I}|$. Before showing this, we need some definitions. Given a point $x \in |\mathcal{I}|$, let

- σ_n^x the unique simplex of minimal dimension in $\Pi^n(\mathcal{I})$ such that $x \in |\sigma_n^x|$,
- $C_n^x :=\downarrow(\uparrow\sigma_n^x)$ the closed star of σ_n^x.
- $\Gamma_n^x := \{\sigma \in \Pi^n(\mathcal{I}) \mid \sigma_n^x \not\leq \sigma\}$,

By construction, $|\sigma_n^x|$ is the smallest closed set containing x in R_n, and we have $|\sigma_{n+1}^x| \subseteq |\sigma_n^x|$, so $\bigcap |\sigma_n^x| = \{x\}$. Moreover, observing that the set Γ_n^x is a downset in $\Pi^n(\mathcal{I})$, and the complement of its realization is $(|C_n^x|)^o$, the interior of $|C_n^x|$, we conclude that the latter is the smallest open containing x which is obtainable from complementing elements of R_n. Finally, since Π is mesh-shrinking, we know that the diameters of the sequences $(|\sigma_n^x|)_n$ and $(|C_n^x|)_n$ go to zero as n goes to infinity.

▶ **Lemma 14.** *The sets $((|C_n^x|)^o)$ for $x \in |\mathcal{I}|$ and $n \in \mathbb{N}$ generate the open topology of $|\mathcal{I}|$. The colimit R of the diagram of maps $(\pi_n : R_n \to R_{n+1})$ is a basis for its closed sets.*

Proof. Let $x \in |\mathcal{I}|$ and X be a closed set in $|\mathcal{I}|$ not containing x. By the above remarks, we know that there exists n such that $(|C_n^x|)^o \subseteq X^c$. Thus $X \subseteq |\Gamma_n^x|$, and by construction $x \notin |\Gamma_n^x|$. ◀

This result states that the co-frame of closed sets for $\Pi^\infty(\mathcal{I})$ and $|\mathcal{I}|$ are isomorphic. This is sufficient to conclude that $|\mathcal{I}|$ is homeomorphic to the subspace of maximal points of $\Pi^\infty(\mathcal{I})$ by the works pioneered by Wallman and Frink, see Cornish [Cor74] for more details. For self-containment, we provide a proof for this specific case in Appendix B.1.

▶ **Proposition 15.** *Given a mesh-shrinking protocol (Π, π) and a simplicial complex $\mathcal{I}$, the geometric realization $|\mathcal{I}|$ is isomorphic to the subspace of $\Pi^\infty(\mathcal{I})$ consisting of maximal points with respect to the specialization order, and the map $\max : \Pi^\infty(\mathcal{I}) \to |\mathcal{I}|$ is continuous.*

Moreover, each point $x \in |\mathcal{I}|$ corresponds to the protocol sequence (σ_n^x), and its principal down-set in the specialization order is order-dual to the set of protocol sequences (σ) such that $x \in \bigcap |\sigma_n|$, ordered point-wise.

5.2 Final proof

Theorems 11, 12, and 13 are direct consequences of Theorem 10 and the following lemma. First we need the following definition: given a task $(\mathcal{I}, \mathcal{O}, \Delta)$ and a continuous map $|\mathcal{I}| \to |\mathcal{O}|$, we say that f is *carried* by Δ provided that for every $\sigma \in \mathcal{I}$, $f(|\sigma|) \subseteq |\Delta(\sigma)|$.

▶ **Lemma 16.** *For a mesh-shrinking protocol (Π, π), and a task $(\mathcal{I}, \mathcal{O}, \Delta)$, there exists a continuous map $|\mathcal{I}| \to |\mathcal{O}|$ carried by Δ if, and only if, there exists a spectral map $\Pi^\infty(\mathcal{I}) \to \mathcal{O}$ carried by Δ.*

Proof. Suppose there exists a continuous map $f : |\mathcal{I}| \to |\mathcal{O}|$ carried by Δ. Note that $|\Pi^n(\mathcal{I})| \cong |\mathcal{I}|$ for all n, so we can consider this to be the domain of f. Now, take a vertex $w \in \mathcal{O}$ and consider $|\uparrow w|$. Since $\uparrow w$ is the complement of a down-set, its realization is open. Similarly, for any vertex $v \in \Pi^n(\mathcal{I})$, $|\uparrow v|$ is open, and since Π is mesh-shrinking, these opens become arbitrarily small for large enough n. Thus, for large enough n, for every vertex $v \in \Pi^n(\mathcal{I})$, there exists some vertex $w \in \mathcal{O}$ with $|\uparrow v| \subseteq f^{-1}[|\uparrow w|]$. We write $d(v) = w$, and show that $\delta = d_*$ is simplicial. Let $\sigma = \{v_0, \ldots, v_k\} \in \Pi^n(\mathcal{I})$ a k-simplex. We have $(|\sigma|)^o \subseteq |\uparrow v_i|$ for all i, meaning in particular that $f((|\sigma|)^o \subseteq \bigcap |\uparrow d(v_i)|$. This intersection is non-empty if, and only if, there is a simplex τ in $\mathcal{O}$ containing the vertices $d(v_i)$, meaning that $\delta(\sigma)$ is indeed a simplex in $\mathcal{O}$. Thus we have a simplicial map $\delta : \Pi^n(\mathcal{I}) \to \mathcal{O}$ which is easily verified to be carried by Δ. Precomposing by the projection $p_n : \Pi^\infty(\mathcal{I}) \to \Pi^n(\mathcal{I})$ gives the desired map.

Conversely, suppose we have a spectral map $g : \Pi^\infty(\mathcal{I}) \to \mathcal{O}$ carried by Δ. By Lemma 8, we obtain an order-preserving map $\delta : \Pi^n(\mathcal{I}) \to \mathcal{O}$ carried by Δ, where we can take n to be arbitrarily large. Since Π is mesh-shrinking, for large enough n and v a vertex of $\Pi^n(\mathcal{I})$, we have $|\delta_*[\uparrow v]| \subseteq |\uparrow w|$ for some vertex w of $\mathcal{O}$. This is enough to conclude that the map δ is simplicial, and therefore that its realization is continuous. Since $|\Pi^n(\mathcal{I})| \cong |\mathcal{I}|$, we obtain the desired continuous map $|\mathcal{I}| \to |\mathcal{O}|$. ◀

5.3 Spectral space of the colorless IIS model

It is well-known [HKR13, Chap. 4.2] that the immediate iterated snapshot protocol, in the context of colorless distributed computing, corresponds to the barycentric subdivision functor, which is mesh-shrinking. Below, we explicit the construction of the limit space for this protocol for an arbitrary number of initial values, and illustrate them for the special cases of two and three initial values.

Consider the immediate iterated snapshot protocol $(\mathbf{Ch}, b)$, where $\mathbf{Ch} : \mathbf{sComp} \to \mathbf{sComp}$ sends a simplicial complex C to the poset $\mathbf{Ch}(C)$ of its chains and their inclusions, which is also a simplicial complex, and where $b_{\mathcal{I}} : \mathcal{I} \longrightarrow \mathcal{D}(\mathbf{Ch}(\mathcal{I}))$ sends a simplex σ to $\downarrow \{\gamma \in \mathbf{Ch}(C) \mid \bigvee \gamma = \sigma\}$. The maps $b_{\mathbf{Ch}^n(\mathcal{I})}$ send a chain $\gamma \in \mathbf{Ch}^n(\mathcal{I})$ to $\downarrow \{\Gamma \in \mathbf{Ch}^{n+1}(C) \mid \bigvee \Gamma = \gamma\}$. Now we define maps for every n as follows:

$$\begin{aligned} m_n : \mathbf{Ch}^{n+1}(\mathcal{I}) &\longrightarrow \mathbf{Ch}^n(\mathcal{I}) \\ \Gamma &\longmapsto \bigvee \Gamma. \end{aligned}$$

We claim that m_n is the dual map associated to $b^\vee_{\mathbf{Ch}^n(\mathcal{I})}$. Indeed, m_n is order-preserving map because if $\Gamma \leq_{\mathbf{Ch}^{n+1}(\mathcal{I})} \Gamma'$ we know that $\bigvee \Gamma \leq_{\mathbf{Ch}^n(\mathcal{I})} \bigvee \Gamma'$ since the former inequality is an inclusion of chains in $\mathbf{Ch}^n(\mathcal{I})$. Moreover, given a down-set $A \in \mathcal{D}(\mathbf{Ch}^n(\mathcal{I}))$, its inverse image under m_n is

$$\begin{aligned} m_n^{-1}[A] &= \{\Gamma \in \mathbf{Ch}^{n+1}(\mathcal{I}) \mid \exists \gamma \in A, \bigvee \Gamma = \gamma\} \\ &= \bigcup_{\gamma \in A} \{\Gamma \in \mathbf{Ch}^{n+1}(\mathcal{I}) \mid \bigvee \Gamma = \gamma\} \\ &= \bigcup_{\gamma \in A} \downarrow \{\Gamma \in \mathbf{Ch}^{n+1}(\mathcal{I}) \mid \bigvee \Gamma = \gamma\} = b^{\vee}_{\mathbf{Ch}^n(\mathcal{I})}(A). \end{aligned}$$

It is clear that $m_n^{-1}[A] \subseteq b^{\vee}_{\mathbf{Ch}^n(\mathcal{I})}(A)$. Conversely, given $\Gamma \in b^{\vee}_{\mathbf{Ch}^n(\mathcal{I})}(A)$, there exist $\gamma' \in A$ and $\Gamma' \in \mathbf{Ch}^{n+1}(\mathcal{I})$ such that $\Gamma \leq \Gamma'$ and $\bigvee \Gamma' = \gamma'$. Writing $\gamma = \bigvee \Gamma$, we deduce $\gamma \leq \gamma'$. Since A is a down-set, we have $\gamma \in A$, proving that $\Gamma \in m_n^{-1}[A]$.

The maps m_n thus give us the projective limit system defining the spectral space $\mathbf{Ch}^\infty(\mathcal{I})$. Using Propositions 6 and 15 and the specificity of this protocol, we will now describe the limit space $\mathbf{Ch}^\infty(\mathcal{I})$. In particular, we will characterize the downsets of maximal points in terms of the terminal co-dimension of their associated minimal sequences when $\mathcal{I}$ is the standard d-simplex for some $d \geq 0$.

First, by Proposition 6, we know that the points of $\mathbf{Ch}^\infty(\mathcal{I})$ correspond to sequences (Γ) such that $\Gamma_n \in \mathbf{Ch}^n(\mathcal{I})$ and $\Gamma_n = \bigvee \Gamma_{n+1}$ for all $n \geq 0$. Now we divide the points $x \in |\mathcal{I}|$ into three classes, depending on the terminal co-dimension of the associated minimal sequences (Γ^x). More precisely, given $x \in |\mathcal{I}|$,

- $x \in C_1$ if $\operatorname{codim}(\Gamma^x_n) = 0$ for all n,
- $x \in C_3$ if $\exists N$ such that $\forall n \geq N$, $\operatorname{codim}(\Gamma^x_n) = 1$,
- $x \in C_\infty$ if $\exists N$ such that $\forall n \geq N$, $\operatorname{codim}(\Gamma^x_n) \geq 2$.

Before proving Proposition 18, we need the following lemma, proved in Appendix B.2.

▶ **Lemma 17.** *Let $\mathcal{I}$ be the standard d-simplex. For $\sigma \in \mathbf{Ch}^n(\mathcal{I})$ with $\operatorname{codim}(\sigma) = 1$, then*

- *If $|\sigma| \subseteq (|\mathcal{I}|)^o$, there exist two unique simplices $\tau_1, \tau_2 \in \mathbf{Ch}^n(\mathcal{I})$ such that $\sigma \leq \tau_1, \tau_2$.*
- *If $|\sigma| \subseteq \partial|\mathcal{I}|$, there exists a unique $\tau \in \mathbf{Ch}^n(\mathcal{I})$, $\sigma \leq \tau$.*

Now we are ready to characterize the number of elements in each down-set under the specialization order. See proof in Appendix B.3

▶ **Proposition 18.** *For $x \in (|\mathcal{I}|)^o$, we have $x \in C_i \iff \#(\downarrow x) = i$ where the downset is taken in $\mathbf{Ch}^\infty(I)$ and $i \in \{1, 3, \infty\}$. For $x \in \partial|\mathcal{I}|$, the same holds excepting that $x \in C_3$ if, and only if, $\#(\downarrow x) = 2$.*

Below, we illustrate the spectral spaces $\mathbf{Ch}^\infty(\mathcal{I})$ when $\mathcal{I}$ is the d-simplex for $d = 1, 2$. For the 1-simplex, $\mathbf{Ch}^\infty(\mathcal{I})$ is the space pictured below in blue. We have included the second barycentric subdivision of the standard 1-simplex above it for reference. Firstly, we observe that the terminal co-dimension of every sequence is either 0 or 1, so each point will either have 0 or 2 points below it, respectively. The subspace of maximal points is isomorphic to the real interval $[0, 1]$, this being the geometric realization of $\mathcal{I}$. At each point x of the form $\frac{k}{2^n}$, for $n \geq 1$ and $0 < k < 2^n$, *i.e.* those that correspond to a vertex in one of the subdivisions, we have two points below x in the specialization order, as is the case for x_3. Extremal points like x_2 have terminal co-dimension 1, but only have one element below them since they are in the border of the realization. All other points in the interval have no elements below them. This is case for $x_1 = \frac{1}{3}$, for example.

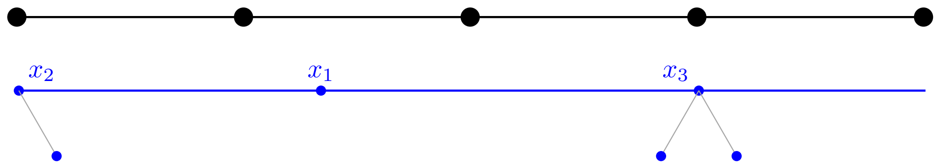

For $d = 2$, we have pictured the first barycentric subdivision on the left, and the limit space $\mathbf{Ch}^\infty(\mathcal{I})$ on the right. The point x_1 is an example of a point such that $\sigma_n^{x_1}$ is of dimension 2 for all n, thus there are no elements below it. Points whose associated minimal sequences are of co-dimension 1 have exactly two elements below them if they are interior, like x_2, or one element below them if not, as is the case for x_3. Finally, points like x_4 and x_5 have a poset of points below them of infinite width and depth of two.

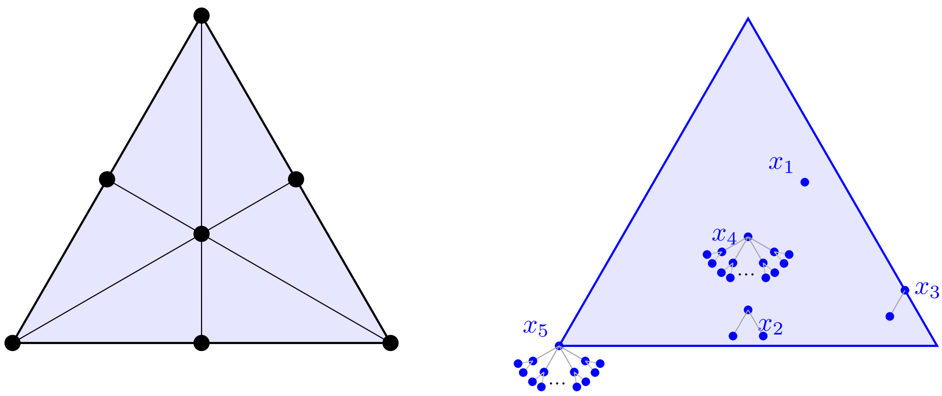

5.4 Comparison with previously defined spaces

As mentioned in the introduction, this contribution aims at unifying the two lines of work [ACN23, CG23] generalizing computability theorems with topological methods. Here, we will compare our approach with theirs via the examples in the previous section.

We underline that our spectral space $\Pi_{\mathcal{M}}^\infty(\mathcal{I})$ is not the space from [ACN23, Theorem 5.4], which is Hausdorff; nor from [NSW19], which contains only executions. Our approach using spectral spaces allows us to separate distinct executions while retaining a strong topological relation when they are indistinguishable in the sense of distributed computation. This is now apparent within the specialization order of the T_0 topology, since executions are indistinguishable precisely when they are below the same maximal point (in the specialization order).

In [CG23], the colored IIS model is investigated, that is, the chromatic subdivision protocol. The set of executions is projected onto a geometric subset of $\mathbb{R}^N$ by a *geo* mapping. In [CG23, Th. 25], it is shown that it is possible to classify the geometric points by the number of pre-images in the set of executions : 1, 2 or infinity. This would have provided a similar classification if the colorless IIS model had been considered. This agrees with our Prop. 18 classifying down-sets in the specialization order. As noted previously, execution sequences, *i.e.* protocol sequences consisting of maximal simplices, are in one-to-one correspondence with executions of the message adversary. So it seems that the spectral space contains both the geometric points and the associated executions, with the following structure : whenever a point has only one pre-image, it corresponds to an execution sequence, whereas, when the point has at least two pre-images, we have both the geometric point, as a maximal point in the spectral space, and the associated execution sequences as minimal points in its down-set under the specialization order. So the spectral space we construct could be interpreted as containing the merge (without unnecessary duplication) of the set of execution sequences and the set of "limit geometrization" points, as well as all intermediary protocol sequences, all organized into a nice T_0 space.

6 Conclusion and Future Work

In this paper we have introduced a new topological encoding for round-based full-information distributed models by spectral spaces and shown that this leads to a colorless computability characterization generalizing the classical colorless computability theorems. These preliminary results demonstrate the effectiveness of this approach. However, the main interest of this method would actually be to get new results for arbitrary adversaries (like the $t-$resilient adversary). Future extensions include determining equivalence of models by studying the topological relationship between the associated spectral spaces, but also investigating how to consider their subsets in order to treat non-compact models. Since submodels of IIS have been proved [CG24] to have the same kind of geometric computability characterization, we believe this should be achievable in the framework of this paper. It would also be important to determine if colored tasks could also be characterized using a similar spectral approach. They are many reasons to believe it could be achieved from the relationship between the space $\Pi^\infty(\mathcal{I})$ and the one proposed in [ACN23].

References

ACN23 Hagit Attiya, Armando Castañeda, and Thomas Nowak. Topological characterization of task solvability in general models of computation. In Rotem Oshman, editor, *DISC'23*, 2023.

AG13 Yehuda Afek and Eli Gafni. *Asynchrony from Synchrony*, pages 225–239. Number 7730 in Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2013.

BG93a E. Borowsky and E. Gafni. Immediate atomic snapshots and fast renaming. In *Proc. of the 12th Annual ACM Symposium on Principles of Distributed Computing*, 1993.

BG93b 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.

BG97 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.

CFQS12 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.

CG23 Yannis Coutouly and Emmanuel Godard. A topology by geometrization for sub-iterated immediate snapshot message adversaries and applications to set-agreement. In Rotem Oshman, editor, *37th International Symposium on Distributed Computing, DISC 2023, October 10-12, 2023, L'Aquila, Italy*, volume 281 of *LIPIcs*, pages 15:1–15:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.

CG24 Yannis Coutouly and Emmanuel Godard. A simple computability theorem for colorless tasks in submodels of the iterated immediate snapshot. In Dan Alistarh, editor, *38th International Symposium on Distributed Computing, DISC 2024, October 28 to November 1, 2024, Madrid, Spain*, volume 319 of *LIPIcs*, pages 16:1–16:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.

Cor74 William H. Cornish. n-normal lattices. *Proceedings of the American Mathematical Society*, 45(1):48–54, 1974.

DST19 Max Dickmann, Niels Schwartz, and Marcus Tressl. *Spectral Spaces*. New Mathematical Monographs. Cambridge University Press, 2019.

FHG25a Stephan Felber, Bernardo Hummes Flores, and Hugo Rincon Galeana. Brief announcement: A sheaf-theoretic characterization of tasks in distributed systems. In *SIROCCO*, volume 15671 of *Lecture Notes in Computer Science*, pages 425–430. Springer, 2025.

FHG25b Stephan Felber, Bernardo Hummes Flores, and Hugo Rincon Galeana. A sheaf-theoretic characterization of tasks in distributed systems. *CoRR*, abs/2503.02556, 2025.

GHK+25 Eric Goubault, Bernardo Hummes Flores, Roman Kniazev, Jérémy Ledent, and Sergio Rajsbaum. A categorical and logical framework for iterated protocols. *CoRR*, abs/2505.10071, 2025.

GP20 Emmanuel Godard and Eloi Perdereau. Back to the coordinated attack problem. *Math. Struct. Comput. Sci.*, 30(10):1089–1113, 2020.

GvG24 Mai Gehrke and Sam van Gool. *Topological Duality for Distributive Lattices*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2024.

HKR13 Maurice Herlihy, Dmitry N. Kozlov, and Sergio Rajsbaum. *Distributed Computing Through Combinatorial Topology*. Morgan Kaufmann, 2013.

Hoc69 M. Hochster. Prime ideal structure in commutative rings. *Transactions of the American Mathematical Society*, 142:43–60, 1969.

HRR12 Maurice Herlihy, Sergio Rajsbaum, and Michel Raynal. Computability in distributed computing: A tutorial. *SIGACT News*, 43(3):88–110, 2012.

HS99 Maurice Herlihy and Nir Shavit. The topological structure of asynchronous computability. *J. ACM*, 46(6):858–923, 1999.

Kei12 Idit Keidar. Distributed computing column 47: distributed computability. *SIGACT News*, 43(3):87, August 2012.

NSW19 Thomas Nowak, Ulrich Schmid, and Kyrill Winkler. Topological characterization of consensus under general message adversaries. In *PODC*, pages 218–227. ACM, 2019.

PP04 J.E. Pin and D. Perrin. *Infinite Words*, volume 141 of *Pure and Applied Mathematics*. Elsevier, 2004.

Raj10 Sergio Rajsbaum. Iterated shared memory models. In Alejandro López-Ortiz, editor, *LATIN 2010: Theoretical Informatics*, pages 407–416, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.

SZ00 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.

A Messages Adversaries

Here we detail the operational definitions of the main example of round-based distributed models : *message adversaries*. In complement from Section 3.3, we present other examples of message adversaries and full description of the operational model.

A.1 Iterated $t-$resilient models

We need the following definition to provide a notion of causality when considering infinite word over digraphs.

▶ **Definition 19** ([CFQS12]). *Let* $\mathbf{G}$ *a sequence* $G_1, G_2, \cdots, G_r, \cdots$. *Let* $p, q \in [n]$. *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 $p \overset{r}{\underset{\mathbf{G}}{\rightsquigarrow}} q$*. We also say that* p *is causally influencing* q *from round* r *in* $\mathbf{G}$*.*

▶ **Example 20.** Consider an Iterated Immediate Snapshot model where we remove all executions where a process is forever influenced by all others processes but t. This yields t-resilient Iterated Immediate Snapshot model $IR_{d,t}$.

Typically, we have for $d \geq 2, t \geq 1, IR_{d,t} \subsetneq IIS_d \subsetneq IS_d$

A.2 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 the initial state to all processes of $[n]$. Then the system progresses in rounds.

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 assumed that p always receives its own value, that is $(p, p) \in A(G)$ for all p and G. However, in examples, this might be implicit for clarity and brevity.

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 $\mathbf{s}_p(\iota.\varepsilon) = \iota(p)$ represents the initial state of p in ι, where ε denotes the empty word.

A.3 Protocol complexes of some message adversaries

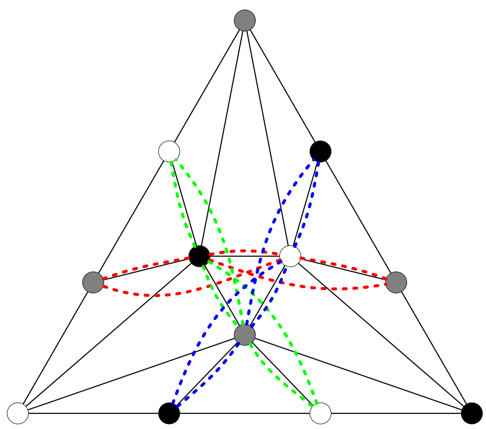

Figure 1 A representation of the complexes corresponding to a round of the Iterated Snapshot and Iterated Immediate Snapshot model for 3 processes.

In Figure 1, we present the protocol complex for one round of IIS_2 and one round of IS_2. The triangles with black edges belong to Π_{IIS}, and we represented in colored dashed lines the additional simplices corresponding to instant graphs of S_2 that are not in ImS_2. We emphasize that, in the figure, when the dashed line is very close to a black line, it actually represents the same simplex of dimension 1. This protocol complex can be iterated, however the corresponding complex is not as easily drawn as the one for the IIS model.

A.4 Protocols as functors with natural transformations

Here we justify our functorial approach to protocols, describing how the Yoneda extension allows us to glue together local information about the action of the distributed protocol and the associated carrier maps, thereby extending to an endofunctor on the category of simplicial complexes with an associated natural transformation as defined in Section 3.2.

As explained in Section 3.1 and in the sections above in the appendix, we consider round-based full-information protocols. What this means concretely is that given a global state, *i.e.* a set $\sigma = \{v_1, \dots, v_n\}$ of local states, a protocol Π returns the collection of possible global states $\sigma_1, \dots, \sigma_{k_n}$ after one round of execution. We denote by $\Pi(\sigma)$ the simplicial complex spanned by these sets, that is, the protocol complex.

Since $\Pi(\sigma)$ contains the possible states obtainable after one round of computation, we have $\bigcup_{\sigma' \subseteq \sigma} \Pi(\sigma') = \Pi(\sigma)$. For $\sigma_1, \sigma_2 \subseteq \sigma$, operationally, $\Pi(\sigma_1) \cap \Pi(\sigma_2)$ consists of all reachable states in which no process can know whether inputs were from σ_1 or σ_2. The only processes which can participate in such an execution are those having values in $\sigma_1 \cap \sigma_2$, because the others are able to distinguish between these two initial configurations. These executions are precisely those whose global states after one round of execution are in $\Pi(\sigma_1 \cap \sigma_2)$, so we have

$$\Pi(\sigma_1) \cap \Pi(\sigma_2) = \Pi(\sigma_1 \cap \sigma_2).$$

Thus, given a subset $\sigma' \subseteq \sigma$, we have $\Pi(\sigma') \subseteq \Pi(\sigma)$. Thus, sub-simplices of σ are mapped to sub-complexes of $\Pi(\sigma)$, defining an order-preserving map $\pi_\sigma : \mathcal{P}(\sigma) \to \mathcal{D}(\Pi(\sigma))$. Finally, since Π is oblivious, the values $v_1, \dots, v_n$ do not affect the outcome $\Pi(\sigma)$, meaning that we may consider two inputs σ, σ' equivalent when they are of the same size. In effect, we are thus in the category of finite sets and set maps.

Let $f : \sigma \to \sigma'$ be a mapping of sets. Since Π is full information, each of the vertices of $\Pi(\sigma)$ and $\Pi(\sigma')$ consist of views containing the values of σ and σ', respectively. Replacing each value $v \in \sigma$ by $f(v)$ in the vertices of $\Pi(\sigma)$ thus yields a vertex of $\Pi(\sigma')$. This yields a simplicial map since $\Pi(\sigma' \cap f[\sigma]) = \Pi(\sigma') \cap \Pi(f[\sigma])$ is a subcomplex. Notice that if f is the canonical inclusion $\sigma \subseteq \sigma'$, this mapping yields the canonical inclusion $\Pi(\sigma) \subseteq \Pi(\sigma')$.

Summing this up, and denoting by $\sigma_n = \{1, \dots, n\}$ the standard n-simplex, by $\Delta_n = \mathcal{P}(\sigma) \setminus \{\emptyset\}$ the standard n-simplex viewed as a complex, and by **Simp** the category of standard n-simplices and set-maps between them, we have a functor

$$\Pi : \mathbf{Simp} \to \mathbf{sComp},$$

along with a natural transformation $\pi : U \circ よ \Rightarrow V \circ \mathcal{D} \circ \Pi$, where $よ : \mathbf{Simp} \to \mathbf{sComp}$ is the Yoneda embedding $\sigma_n \mapsto \Delta_n$, and $U : \mathbf{sComp} \to \mathbf{Pos}_f$ and $V : \mathbf{DL} \to \mathbf{Pos}_f$ are forgetful functors. Moreover, each component of π_{σ_n} is a strict carrier map which is moreover *total*, *i.e.* for which the join of its images is the full complex $\Pi(\sigma_n)$.

Now we describe how to extend this information to the definition provided in Section 3.2 using the Yoneda extension. Given a simplicial complex $\mathcal{C}$ and denoting by $D_{\mathcal{C}}$ the diagram in **Simp** of inclusions of simplices of $\mathcal{C}$, we have $\mathcal{C} = \mathrm{colim}(よ(D_{\mathcal{C}}))$. Denoting by $\Pi(D_{\mathcal{C}})$

the image of this diagram under Π, and by $\Pi(\mathcal{C})$ its colimit, we have extended Π to an endofunctor of $\mathbf{sComp}$. Indeed, we have thereby defined an assignment on objects, and, given a simplicial map $f : \mathcal{C} \to \mathcal{C}'$, it comes from a map $V_{\mathcal{C}} \to V_{\mathcal{C}'}$ which defines a natural transformation from $D_{\mathcal{C}}$ to $D_{\mathcal{C}'}$. Indeed, for any simplex $\sigma = \{v_1, \dots, v_k\}$, we have a map of sets $f_{|\sigma} : \sigma \to f[\sigma]$, where $f[\sigma] = \{f(v_1), \dots, f(v_k)\}$. Applying Π, we see that all of this entails that $\Pi(\mathcal{C}')$ is a co-cone for $\Pi(D_{\mathcal{C}})$, which, by the universal property of the colimit, provides a unique map $\Pi(f) : \Pi(\mathcal{C}) \to \Pi(\mathcal{C}')$.

The argument extending the natural transformation is similar. Indeed, for a complex $\mathcal{C}$, noticing that forgetful functors preserve colimits, we have that $\mathcal{C}$ is the colimit in $\mathbf{Pos}_f$ of the diagram $U(\text{よ}(D_{\mathcal{C}}))$, and since $\mathcal{D}$ also preserves colimits, we have that $\mathcal{D}(\Pi(\mathcal{C}))$ is the colimit in $\mathbf{Pos}_f$ of the diagram $V(\mathcal{D}(\Pi(D_{\mathcal{C}})))$. By construction, the appropriate components of π are a natural transformation between $U(\text{よ}(D_{\mathcal{C}}))$ and $V(\mathcal{D}(\Pi(D_{\mathcal{C}})))$, thereby providing a unique $\pi_{\mathcal{C}} : U(\mathcal{C}) \to V(\mathcal{D}(\Pi(\mathcal{C})))$ by the property of the colimit. Combining this with the arguments used to prove that Π is functorial on $\mathbf{sComp}$, we deduce the commutativity of naturality squares. Finally, preservation of intersections, *i.e.* rigidity, and the join of images being the full complex, *i.e.* essential surjectivity, are preserved under these constructions.

Thus, we have shown that from a full-information oblivious protocol, we deduce a pair (Π, π), where $\Pi : \mathbf{sComp} \to \mathbf{sComp}$ is a functor and π is a natural transformation between $U : \mathbf{sComp} \to \mathbf{Pos}_f$, the forgetful functor, and $\mathcal{D} \circ \Pi : \mathbf{sComp} \to \mathbf{Pos}_f$ whose components are strict and essentially surjective carrier maps.

B Proofs

B.1 of Proposition 15

Since $|\mathcal{I}|$ is a Hausdorff space and Π is mesh-shrinking, the set $\bigcap |\sigma_n|$ is a singleton for every protocol sequence. This means that (σ) and (τ) are comparable in the point-wise order on protocol sequences if, and only if, $\bigcap \tau = \bigcap \sigma$. Indeed, if $\sigma_n \leq \tau_n$ for all n, then $|\sigma_n| \subseteq |\tau_n|$, whence $\bigcap \tau = \bigcap \sigma$. Conversely, suppose $\{y\} = \bigcap \tau \neq \bigcap \sigma = \{x\}$, and recall that since $|\mathcal{I}|$ is Hausdorff, there exist neighborhoods U and V of x and y, respectively, such that $U \cap V$ is empty. By Lemma 14, the topology on $|\mathcal{I}|$ is generated by $(\widehat{|\mathring{C}^z_n|})$, so there exists n such that $\widehat{|\mathring{C}^x_n|} \subseteq U$ and $\widehat{|\mathring{C}^y_n|} \subseteq V$, meaning that σ_n and τ_n are incomparable in $\Pi^n(\mathcal{I})$.

Since the $(\sigma^x_n$, for $x \in |\mathcal{I}|$ are each the minimum in the point-wise order among protocol sequences (σ) such that $\bigcap |\sigma_n| = \{x\}$, these are precisely the maximal points of $\Pi^\infty(\mathcal{I})$. Its subspace topology thus consists in restricting opens to points below such sequences. In Lemma 14, we showed that such opens suffice to generate the topology of $|\mathcal{I}|$, meaning that the latter is homeomorphic to $\max(\Pi^\infty(\mathcal{I}))$. Finally, for the continuity of the max map, observe that given a generating open V of $\max(\Pi^\infty(\mathcal{I}))$, its pre-image under max is $\downarrow V$, where the down-set is taken with respect to the specialization order of $\Pi^\infty(\mathcal{I})$. Note that if F is a prime filter corresponding to a point $x \in |\mathcal{I}|$, then $V \in F$, and moreover, if $F \subseteq F'$, then $V \in F'$. Since the specialization order is order-dual to the subset ordering on prime filters, we conclude that

$$\hat{V} = \{F \in pFilt(\Pi^\infty(\mathcal{I})) \mid V \in F\} = \downarrow V,$$

thus showing that $\max^{-1}(V) = \hat{V}$ is a generating open.

B.2 of Lemma 17

We proceed by induction on $n \geq 0$. For $n = 0$, we, without loss of generality, write $\sigma = \{v_1, \ldots, v_d\}$ where $V(\mathcal{I}) = \{v_1, \ldots, v_{d+1}\}$. The latter is the unique maximal simplex that contains σ, and clearly $|\sigma| \subseteq \partial|\mathcal{I}|$. For $n \geq 0$, $\sigma \in \mathbf{Ch}^{n+1}(\mathcal{I})$ of co-dimension one and let $\sigma' = m_n(\sigma)$. By the induction hypothesis, there exists unique

- $\tau'_1, \tau'_2 \in \mathbf{Ch}^n(\mathcal{I})$ covering σ' and $|\sigma'|$ is contained in the interior of $|\mathcal{I}|$. Thus $|\sigma| \subseteq |\sigma'|$ is too. Now, since τ'_i is a d-simplex, without loss of generality we write $\{v_1, \ldots, v_d, v_{d+1,i}\}$ for its set of vertices, and $\sigma' = \{v_1 \ldots, v_d\}$. Since $\sigma \in b_n(\sigma')$, we can write $\sigma = (V_1, \ldots, V_d)$ where $V_1 \leq \cdots \leq V_d$ is a maximal chain in $\mathcal{P}(\sigma')$. Up to permutation, we can assume without loss of generality that $V_k = v_1, \ldots, v_k$. A simplex $\tau \in b_n(\tau'_i)$ is a maximal chain $(U_1, \ldots, U_{d+1})$ in $\mathcal{P}(\tau'_i)$, its faces being $(U_1, \ldots, \hat{U}_i U_{d+1})$, where $\hat{U}_i$ means omitting U_i from the chain. Thus, the only simplex in $b_n(\tau'_i)$ covering σ is $\tau_i = \{V_1, \ldots, V_d, V_d \cup \{v_{d+1,i}\}\}$.
- $\tau' \in \mathbf{Ch}^n(\mathcal{I})$ covering σ' and $|\sigma| \subseteq |\sigma'| \subseteq \partial|\mathcal{I}|$. Using the same notations as above for σ, σ', and writing $\tau' = \{v_1, \ldots, v_d, v_{d+1}\}$, a similar argument shows that the unique simplex in $b_n(\tau)$ covering σ is $\tau = \{V_1, \ldots, V_d, V_d \cup \{v_{d+1}\}\}$.

B.3 of Proposition 18

We proceed case by case:

- ($x \in C_1$) In this case Γ_n^x is maximal in each $\mathbf{Ch}^n(\mathcal{I})$, so its principal in the specialization order is a singleton.
- ($x \in C_3$) By Lemma 17, each σ_n^x is covered by either exactly one or two simplices depending on whether $x \in \partial|\mathcal{I}|$ or not, respectively. Since the m_n are order and dimension preserving, these form sequences (τ) or (τ_1) and (τ_2), respectively. Since each τ_n or the $(\tau_i)_n$ are unique cover of σ_n^x, these are the unique sequence above (σ^x) in the point-wise order, which is the order-dual of the specialization.
- ($x \in C_\infty$) Let N so that $\mathrm{codim}(\sigma_n^x) \geq 2$ for all $n \geq N$. There exists τ_N of codimension one such that $\sigma_N^x \leq \tau_N$. The latter is covered by two maximal simplices τ_N^1 and τ_N^2. By a similar argument to that given in Lemma 17, we show that σ_{N+1}^x is below by at least four maximal simplices $\tau_{N+1}^{11}, \tau_{N+1}^{12}, \tau_{N+1}^{21}, \tau_{N+1}^{22}$ in $\mathbf{Ch}^{N+1}(\mathcal{I})$, where $\tau_{N+1}^{i1}, \tau_{N+1}^{i2} \in b_N(\tau_N^i)$. By induction we thus conclude that σ_{N+k}^x is covered by at least 2^{k+1} maximal simplices in $\mathbf{Ch}^{N+k}(\mathcal{I})$, and that we have constructed 2^{k+1} maximal sequences that are above $(\sigma_N^x, \ldots \sigma_{N+k}^x)$ in the point-wise order on $\prod_{i=N}^{N+k} \mathbf{Ch}^i(\mathcal{I})$. This shows that taking k to infinity, we obtain an infinite amount of sequences above (σ_n^x).