



HAL
open science

Semi-Simplicial Set Models for Distributed Knowledge

Éric Goubault, Roman Kniazev, Jérémy Ledent, Sergio Rajsbaum

► **To cite this version:**

Éric Goubault, Roman Kniazev, Jérémy Ledent, Sergio Rajsbaum. Semi-Simplicial Set Models for Distributed Knowledge. *Logic in Computer Science*, Jun 2023, Boston (MA), United States. 10.1109/lics56636.2023.10175737 . hal-04336328

HAL Id: hal-04336328

<https://polytechnique.hal.science/hal-04336328v1>

Submitted on 11 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Semi-simplicial Set Models for Distributed Knowledge

Éric Goubault*, Roman Kniazev*[†], Jérémy Ledent[‡], Sergio Rajsbaum[§]

* LIX, CNRS, École Polytechnique, IP-Paris, Palaiseau Cedex, France

[†] Université Paris-Saclay, ENS Paris-Saclay, CNRS, LMF, Gif-sur-Yvette, France

[‡] MSP Group, University of Strathclyde, Glasgow, Scotland

[§] Universidad Nacional Autónoma de México (UNAM), Mexico D.F., Mexico and IRIF, Paris, France

Abstract—In recent years, a new class of models for multi-agent epistemic logic has emerged, based on simplicial complexes. Since then, many variants of these simplicial models have been investigated, giving rise to different logics and axiomatizations. In this paper, we present a further generalization, which encompasses all previously studied variants of simplicial models. Geometrically, this is achieved by generalizing beyond simplicial complexes, and considering instead semi-simplicial sets. By doing so, we define a new semantics for epistemic logic with distributed knowledge, where a group of agents may distinguish two worlds, even though each individual agent in the group is unable to distinguish them. As it turns out, these models are the geometric counterpart of a generalization of Kripke models, called “pseudo-models”. We show how to recover the previously defined variants of simplicial models as sub-classes of our models; and give a sound and complete axiomatization for each of them.

Index Terms—Epistemic logic, Simplicial sets, Distributed knowledge

I. INTRODUCTION

The usual semantics for multi-agent epistemic logic is based on the classic Kripke *possible worlds* relational structure [1]. However, the intimate relationship between distributed computing and algebraic topology [2] showed the importance of moving from the focus on global states represented by worlds, to local states, representing perspectives about possible worlds. Namely, moving from graph structures to simplicial complex structures. A formal semantics of multi-agent epistemic formulas in terms of simplicial models was presented [3], and shown to be equivalent to the usual Kripke model semantics for $S5_n$. Further work explored bisimilarity of simplicial models [4] and connections with covering spaces [5]. Remarkably, it has been shown by Yagi and Nishimura [6], [7] that the implicit topological information in Kripke models, exposed by the simplicial complex point of view, can be leveraged to produce a logical obstruction to the solvability of certain distributed computing problems. The notion of *distributed knowledge* [8] plays a crucial role there: in a sense, it is a higher-dimensional notion of knowledge.

Eric Goubault was partially funded by AID/CIEDS project FARO. Part of the work of Sergio Rajsbaum was performed while he was an invited professor at IRIF, Université Paris Cité, and at LIX, Ecole Polytechnique.

These first results assumed a finite, fixed set of agents, whose local states appear in every world. As a consequence, since a world with $n + 1$ agents is represented by a n -dimensional simplex, every facet of the simplicial model is of the same dimension. Such models are called *pure* simplicial models. However, in distributed systems, processes may fail, and it may happen that only a subset of the agents remain in some worlds. To model such situations, the categorical equivalence of [3] was extended by various authors to include simplicial models that may not be pure [9], [10], [11]. These situations have been thoroughly studied since early on in distributed computability, e.g. the seminal work of Dwork and Moses [12], where a complete characterization of the number of rounds required to reach simultaneous consensus was given in terms of common knowledge. By moving away from pure simplicial models, we also move away from the standard $S5_n$ epistemic logic. Indeed, there are a number of design choices to be made, both to define the models and their semantics. Should all simplices represent worlds, or only the facets? Can we have a world with no agents? How do we deal with formulas involving the knowledge of agents that are not present in the current world? All these choices can greatly influence the resulting logic: while [9] drops Axiom T and works with the logic $KB4_n$ (augmented with some extra axioms), other authors [11] take a completely different route and define a three-valued logic (where formulas may be undefined), with an axiom system called SS^{\square} .

In this paper, we aim to bring order to chaos by defining a class of simplicial models that encompasses all previous variants. Thus, the main objective of our paper is to *unify* previous work on simplicial complex models. For that purpose, we introduce a new class of models called *epistemic covering models*. As we will see in Section IV-B, the models studied in [3], [9], [10], [15] can all be viewed as sub-classes of epistemic covering models. This allows us to recover important theoretical results about each of those sub-classes, by proving them once and for all in our very general setting. Namely, our paper contains two central results: (i) Theorem 1 shows an equivalence of categories between epistemic covering models and their Kripke-style counterpart, called *generalized epistemic models*. Crucially, we also show (Lemma 3) that this equivalence can safely be restricted to sub-classes of epistemic


covering models; allowing us to recover similar equivalence results from [3], [9], [10]. (ii) The other essential result of our paper is a sound and complete axiomatization of epistemic covering models (Corollary 2). Once again, we are also interested in axiomatizing the various sub-classes of epistemic covering models, which we do in Section V-C. We show that there is a close correspondence between structural properties that define a sub-class, and axioms of the corresponding logic. These results are summed up in Fig. 1 below. The first two columns indicate the correspondence between properties of epistemic coverings, and axioms of the logic (they will be defined in Sections IV and V, respectively). The check marks in the other columns indicate how to recover previously-studied variants of simplicial models. One can check that the axiomatization provided here matches the one given in each of those papers (except for [10], which studies a different three-valued semantics).

Covering properties	Axioms	[3]	[9]	[10]	[15]
Proper	(P)	✓	✓	✓	✓
Pure	(T)	✓			✓
Minimal	(Min)	✓	✓		✓
Maximal	(Max)			✓	
No empty world	(NE)	✓	✓	✓	✓
Simplicial complex	no axiom	✓	✓	✓	

Fig. 1. Sub-classes of epistemic covering models that have been studied in previous papers, and their associated axioms.

One distinguishing feature of our models, which has not been considered previously in the literature, is that we generalize beyond simplicial complexes and consider instead semi-simplicial sets. Geometrically, this means that simplices (a.k.a. worlds) can be connected in more complex ways; for instance, two triangles might share two vertices but not the edge between those vertices (see Fig. 2, which is not a simplicial complex). This new feature enables us to model situations where a group of agents may distinguish two worlds, even though each individual agent in the group is unable to distinguish them.

A toy example of such a situation, in the realm of distributed computing is the following. Suppose we have three sensors, s_1 , s_2 and s_3 , in a sensor network, with overlapping visibility regions of the form of a unit disk, that can only count the number of targets within their visibility region. Suppose now that each of these three sensors detects exactly one target, indicated by a “ \times ” below. There are 5 possible configurations w_1, w_2, \dots, w_5 (from left to right below), with the total number of targets, ranging from 1 to 3:



Alone, no sensor can distinguish between these 5 situations. To disambiguate between configurations, we must indeed use an inclusion-exclusion principle, so we need to identify if the

same target has been detected by different sensors. For this, we suppose that, jointly, groups of sensors can determine if a target they saw was seen by all in that group. In this example, s_1 and s_2 together can for instance distinguish the worlds w_1 and w_2 , but not w_1 and w_4 . Only the three sensors together can distinguish between all five worlds. The framework we propose in this paper allows to formalize such applications.

We show the very important, and surprising fact that every semi-simplicial set model is bisimilar to a simplicial complex model (leading to Theorem 6). There are several reasons why, despite this fact, we still believe that semi-simplicial set models are worth studying. Indeed, some epistemic situations are described much more naturally and concisely using a semi-simplicial set model. This is for instance the case of the sensor model described previously. By turning it into a bisimilar simplicial complex model, we create a tree-like model with an infinite number of worlds, and where the underlying topology has disappeared. For some applications (e.g. model-checking), it might be crucial to keep a model which is as small as possible, let alone infinite. This situation is akin to classical topology: semi-simplicial sets and simplicial complexes describe the same spaces up to weak-homotopy. Yet, both structures are used in algebraic topology in their own right, according to the situation at hand. For some spaces, e.g. when dealing with the classification of combinatorial structures (posets, graphs, etc.), the simplicial complex approach is indeed very practical. However, in other situations (e.g. for non-triangulable spaces), the semi-simplicial set approach is much more convenient to use.

Plan of the paper: In Section II, we recall various mathematical notions that we will be using in this paper. In Section III, we define our two notions of models, generalized epistemic models and epistemic covering models; and prove that they are equivalent. We define various interesting sub-classes of these models in Section IV. In Section V, we prove that \mathbf{EC}_n is sound and complete with respect to these models (Theorem 2), and give sound and complete axiomatizations of interesting sub-classes (Theorems 3, 4, and 5). Finally, we prove that every simplicial set model is bisimulation equivalent to a simplicial complex model, bringing interesting questions about the topological nature of bisimulation (Theorem 6).

II. PRELIMINARIES

A. Kripke semantics of distributed knowledge

Let A be a finite set of agents, and At a countable set of atomic propositions. We consider the language \mathcal{L}_D of epistemic logic with *distributed knowledge* [1], generated by the following BNF grammar:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid D_U \varphi \quad p \in \text{At}, U \subseteq A$$

The derived operators $\vee, \Rightarrow, \text{true}, \text{false}$, are defined as usual in propositional logic. We also use the following operators:

$$K_a \varphi := D_{\{a\}} \varphi \quad \widehat{D}_U \varphi := \neg D_U \neg \varphi$$

The operator $D_U\varphi$ is read “the group of agents U (collectively) knows φ ”, while its dual $\widehat{D}_U\varphi$ means that the group U considers possible the formula φ . $K_a\varphi$ is the standard knowledge operator of epistemic logic, “agent a knows φ ”.

Definition 1. A Kripke model is a structure $\mathcal{M} = \langle M, (\sim_a)_{a \in A}, L \rangle$, where:

- M is a set of *possible worlds*,
- For every agent $a \in A$, \sim_a is an equivalence relation,
- $L : M \rightarrow \mathcal{P}(\text{At})$ is a valuation function.

Given a Kripke model, the satisfaction relation $M, w \models \varphi$ is defined inductively as follows. We write $\sim_U = \bigcap_{a \in U} \sim_a$.

$$\begin{aligned} M, w \models p & \quad \text{iff} \quad p \in L(w) \\ M, w \models \neg\varphi & \quad \text{iff} \quad M, w \not\models \varphi \\ M, w \models \varphi \wedge \psi & \quad \text{iff} \quad M, w \models \varphi \text{ and } M, w \models \psi \\ M, w \models D_U\varphi & \quad \text{iff} \quad M, w' \models \varphi \text{ for all } w' \in M \\ & \quad \text{such that } w \sim_U w' \end{aligned}$$

B. Partial Equivalence Relations

Definition 2. A *Partial Equivalence Relation* (PER) on a set X is a relation $R \subseteq X \times X$ which is symmetric and transitive (but not necessarily reflexive).

The *domain* of a PER R is the set $\text{dom}(R) = \{x \in X \mid R(x, x)\} \subseteq X$, and it is easy to see that R is an equivalence relation on its domain, and empty outside of it. Thus, a PER on X is simply an equivalence relation on a subset of X . The *equivalence classes* of a PER R are defined as usual, when viewed as an equivalence relation on $\text{dom}(R)$: for $x \in \text{dom}(R)$, we write $[x]_R = \{y \in X \mid xRy\}$ for the equivalence class of x , and X/R for the set of equivalence classes.

C. Semi-simplicial sets

(Semi)-simplicial sets can be viewed as a higher-dimensional generalization of directed multigraphs, while simplicial complexes are a generalization of simple graphs. Thus, in dimension 1, simplicial sets allow loops and parallel edges; whereas simplicial complexes do not. Similarly, in higher dimensions, simplicial sets allow simplices to be assembled together in more complex ways, as in Fig. 2. Standard references for the theory of simplicial sets and some of their many uses in algebraic topology are [17], [18] and a more elementary introduction can be found in [19].




Fig. 2. A simplicial set which is not a simplicial complex.

Compared to simplicial complexes, simplicial sets exhibit interesting geometric and categorical features:

- In algebraic topology, they provide a combinatorial model for the homotopy theory of topological spaces, in fact they are Quillen equivalent to the standard Quillen model category of topological spaces [17].

- As a presheaf category, simplicial sets form a Grothendieck topos. Still, simplicial complexes are almost as nice, they are known as forming a quasitopos, see e.g. [20] for a discussion of the relationship between the quasitopos of simplicial complexes and the topos of simplicial sets.

Let Δ_{inj}^+ be the category of possibly empty linear orders with injective maps between them. We write $[n]$ for the $(n+1)$ -element linear order $[n] = \{0 < \dots < n\}$, and $[-1] = \emptyset$. The category of *augmented semi-simplicial sets* is defined as the presheaf category $\widehat{\Delta}_{\text{inj}}^+$. Thus, its objects are functors $F : (\Delta_{\text{inj}}^+)^{\text{op}} \rightarrow \mathbf{Set}$, and morphisms are natural transformations. The elements of $F([n])$ are called n -simplices, and we use the terms vertices, edges, and triangles for simplices of dimension 0, 1 and 2 respectively.

For example, the simplicial set depicted in Fig. 2 has 4 vertices, 6 edges and 2 triangles. So, accordingly, the sets $F([0])$, $F([1])$ and $F([2])$ are chosen to have 4, 6 and 2 elements, respectively. There are three face maps $\partial_i : F([2]) \rightarrow F([1])$ for $i = 0, 1, 2$ which assign to each triangle its three boundary edges. The maps $\partial_i : F([1]) \rightarrow F([0])$ for $i = 0, 1$ assign to each edge its source and target vertices.

III. GENERALIZED EPISTEMIC FRAMES AND EPISTEMIC COVERINGS

In this section, we define two classes of models for distributed knowledge: one generalizes Kripke frames, and the other generalizes simplicial models. We show that they are structurally equivalent. In the following sections, we will show that they subsume previous models found in the literature, and study their axiomatization.

A. Generalized epistemic frames

As in [9] we will consider Kripke frames where the accessibility relation \sim is not an equivalence relation but just a partial equivalence relation. But we generalize one step further by associating an accessibility relation not only for each agent, but also for any set of agents, making it possible to interpret distributed knowledge in a very general manner. Such models have been considered before in the epistemic logic literature, sometimes called “pseudo-models”, as an intermediate tool in completeness proofs involving distributed knowledge [13], [14]. One paper that used such models as the main object of study is [15], in order to model observability in quantum systems. Still, the models that we present here generalize further, by allowing worlds where not all agents are necessarily present, and as a consequence condition (b) below is new.

Definition 3. A *generalized epistemic frame* is a structure $\mathcal{M} = \langle M, \sim \rangle$, where:

- $M = \{w_0, w_1, \dots\}$ is a set of *possible worlds*,
- \sim is a function assigning for every group of agents $U \subseteq A$, a PER $\sim_U \subseteq M \times M$ called the *U -accessibility relation*. We write $[w]_U$ for the equivalence class of w with respect to \sim_U .
- The PERs \sim_U satisfy the following conditions.

(a) Compatibility:

$$\forall U' \subseteq U, w \sim_U w' \Rightarrow w \sim_{U'} w'$$

(b) Closure under union of the groups of alive agents:

$$\forall U, U' \subseteq A, (w \sim_U w \wedge w \sim_{U'} w) \Rightarrow w \sim_{U \cup U'} w.$$

The two conditions on \sim_U can be interpreted as follows. Condition (a) means that if a group of agents cannot distinguish between two worlds, all together, then there is no way a subgroup of agents can distinguish the same two worlds. Condition (b) implies that in each world w , there is a maximal group of agents U such that $w \sim_U w$. We call such a U the group of *alive agents* in world w , and denote it by \bar{w} . We say that an agent a is *alive* in w when $a \in \bar{w}$ or, equivalently, when $w \sim_{\{a\}} w$.

Usually, we write $w \sim_a w'$ as shorthand for $w \sim_{\{a\}} w'$. Additionally, we say that a world w is a *sub-world* of w' when $\bar{w} \subseteq \bar{w}'$ and $w \sim_{\bar{w}} w'$. The sub-world relation is a preorder, and is denoted by $w \leq w'$.

Example 1. An example of a generalized epistemic frame is given on Figure 3. It has seven worlds $M = \{w_0, \dots, w_5\} \cup \{w'_1\}$. Not all relations are shown, but only the generating ones. Sets of alive agents can be read off directly on the reflexive loops above a world; in w_5 , no agent is alive, that is $w_5 \not\sim_a w_5$ for all $a \in A$. The empty group can distinguish w_5 from other worlds: $w_5 \not\sim_{\emptyset} w_i, i \neq 5$. In world w_4 , agents a and c are alive, and they can distinguish it from all other worlds, but the empty group cannot. In w_2 and w_3 all three agents are alive. Agents b and c cannot individually distinguish w_2 from w_3 , however, together they can: we have $w_2 \sim_b w_3$ and $w_2 \sim_c w_3$, but $w_2 \not\sim_{\{b,c\}} w_3$. In w_1 and w'_1 , agents a and b are alive, but even together they cannot distinguish them, as we have $w_1 \sim_{\{a,b\}} w'_1$. World w_0 is a sub-world of both w_1 and w'_1 : only b is alive in w_0 and it cannot distinguish w_0 from w_1 or w'_1 .




Fig. 3. An example of a generalized epistemic frame.

Example 2. In the example of the introduction, w_1, \dots, w_5 are worlds, s_1, s_2, s_3 are agents, representing the sensors. For all i, j, k , we have $w_i \sim_{s_j} w_k$: no individual sensor s_j can distinguish between any pair of worlds w_i and w_k . We also have $w_1 \sim_{\{s_2, s_3\}} w_2$ since in these two worlds, s_2 and s_3 jointly see the same target. More generally, $w_1 \sim_{\{s_i, s_j\}} w_k$ if and only if some target is in $A_i \cap A_j$ in world w_k . But

the three sensors together can distinguish w_1 from all other worlds. Examining all other intersections in pairs of worlds gives the epistemic frame on Figure 4.




Fig. 4. Sensor network as an epistemic frame.

Morphisms between epistemic frames are structure-preserving functions between the sets of worlds.

Definition 4. Let $\mathcal{M} = \langle M, \sim \rangle$ and $\mathcal{N} = \langle N, \sim' \rangle$ be two generalized epistemic frames. A morphism from \mathcal{M} to \mathcal{N} is a function $f : M \rightarrow N$ such that for all $U \subseteq A$, for all $u, v \in M$, $u \sim_U v$ implies $f(u) \sim'_U f(v)$,

We write \mathbf{GEF}_A for the category of generalized epistemic frames with agents A . Later, in order to define the semantics of \mathcal{L}_D formulas, we will equip these frames with a valuation function, as in Section II-A. But first, let us first define the geometric counterpart of these frames: epistemic coverings.

B. Epistemic coverings

1) *Chromatic augmented semi-simplicial sets*: As in the case of simplicial models [3], our first step will be to decorate the vertices of a simplicial set with colors, representing the names of the agents in A . The resulting structure is called a *chromatic augmented semi-simplicial set*, or *cset* for short. We identify A with the linear order $[n] = \{0 < \dots < n\}$.

Let S_A denote the *standard* $(|A| - 1)$ -simplex, defined by:

- $(S_A)_k = \{(i_0, \dots, i_k) \mid 0 \leq i_0 < \dots < i_k \leq n\}$,
- $\partial_j(i_0, \dots, i_k) = (i_0, \dots, i_{j-1}, i_{j+1}, \dots, i_k) \in (S_A)_{k-1}$.

Given an augmented semi-simplicial set X , a coloring of X by the agents in A is simply a map $f : X \rightarrow S_A$. Note that, since we work with *semi-simplicial sets* here, without degeneracy maps, morphisms preserve the dimension of simplices. Thus, each simplex of X is well-colored, in the sense that all vertices in a simplex are labelled with distinct agents. We then define the category of chromatic augmented semi-simplicial sets to be the slice category $\hat{\Delta}_{inj}^+ / S_A$.

For the rest of the paper, we can either see this category as a slice category, or notice that, by the fundamental theorem of topos theory, the category of csets is once again a presheaf category on a site Γ made of simplices of the standard n -simplex. As with the site of semi-simplicial sets (see e.g., [21], [22]), Γ is the posetal category of subsets of A with the inclusion partial order, defined below.

in a facet of a simplicial complex. However, this construction is not functorial.

C. Isomorphism between epistemic frames and coverings

We want to show that the category of epistemic coverings and the category of generalized epistemic frames are isomorphic. We first define the functor $\kappa : \mathbf{eCov}_A \rightarrow \mathbf{GEF}_A$.

Let $f : E \rightarrow B$ be an epistemic covering. As E is projective, it can be represented as a disjoint union of disconnected standard simplices, that is $E = \bigcup_{U \subseteq A} \bigcup_{E_U} \Gamma[U]$, where E_U are some sets. Then the set W of worlds in $\kappa(f) = \langle W, \sim \rangle$ is exactly $\bigcup_{U \subseteq A} E_U$, that is every maximal simplex of E is interpreted as a world. To define the group indistinguishability relations \sim_U between the worlds of the corresponding frame, we proceed as follows. We think of two simplices in E as U -indistinguishable, if their images in B share a U -face. Formally, given two worlds $w_s, w_{s'}$ in $\kappa(f)$, corresponding to two maximal simplices s, s' in E , we let $w_s \sim_U w_{s'}$ if $\partial_U(f(s)) = \partial_U(f(s'))$.

Now we define this functor on morphisms. Suppose we are given a morphism of epistemic coverings $\langle \alpha_E, \alpha_B \rangle$. Then, the morphism of epistemic frames $\kappa(\langle \alpha_E, \alpha_B \rangle) : \kappa(f) \rightarrow \kappa(f')$ sends a world w_s to a world $w_{s'}$ if the image of the maximal simplex s in E is included in the maximal simplex s' in E' .

Next, we construct the inverse functor $\sigma : \mathbf{GEF}_A \rightarrow \mathbf{eCov}_A$. Given a generalized frame $\langle M, \sim \rangle$, we need to define two csets E_M and B_M together with a map $m : E_M \rightarrow B_M$. For every world $w \in M$ we associate a simplex of type \bar{w} to E_M , that is $E_M = \bigcup_{w \in M} \Gamma[\bar{w}]$.

It is slightly more intricate to build the base space B_M of the covering. We construct it as a presheaf on Γ , $B_M : \Gamma^{\text{op}} \rightarrow \mathbf{Set}$. For every group $U \subseteq A$ of agents, define $B_M(U) = M/\sim_U$, the quotient of M with respect to \sim_U . The restriction $\partial_{U,V} : M/\sim_V \rightarrow M/\sim_U$ sends the equivalence class $[w]_V$ to $[w]_U$. We need to verify that B_M is indeed a cset.

Lemma 1. *For any frame $\langle M, \sim \rangle$, B_M is a cset.*

Proof: First, let us show that restriction maps are correctly defined. Consider $\partial_{U,V}$ for $U \subseteq V$, and let $w' \in [w]_U$ be another representative of the equivalence class $[w]_U$. By monotonicity, we have $w \sim_V w' \Rightarrow w \sim_U w'$, so they also belong to the same equivalence class $[w]_U = [w']_U$. Thus, the function $\partial_{U,V}$ is correctly defined. To see that this is functorial, given sets of agents $U \subseteq V \subseteq W$, we need to prove that $\partial_{U,V} \circ \partial_{V,W} = \partial_{U,W}$, which is straightforward. ■


One can see that there is a canonical map $m : E_M \rightarrow B_M$ which sends a maximal simplex w of E_M , representing a world, to the corresponding equivalence class $[w]_{\bar{w}} \in B_M(\bar{w})$. This describes σ on objects by setting $\sigma(\langle M, \sim \rangle) = m$.

To define σ on morphisms, consider a morphism of epistemic frames $g : \mathcal{M} \rightarrow \mathcal{N}$. As both assignments of total spaces E_M, E_N and base spaces B_M, B_N are functorial one needs to check that the induced square commutes.

Proposition 1. *The functors κ and σ define an isomorphism of categories: $\kappa \circ \sigma = \text{id}_{\mathbf{GEF}_A}$ and $\sigma \circ \kappa = \text{id}_{\mathbf{eCov}_A}$.*

Proof: Consider an epistemic frame \mathcal{M} . Then $\kappa\sigma(\mathcal{M})$ has as its worlds the same worlds as in \mathcal{M} , as κ and σ just transfer this information. The relations \sim_U are also just the same: $w \sim_U w'$ in $\kappa\sigma(\mathcal{M})$ iff the simplices $\partial_U(w), \partial_U(w')$ in $E_{\sigma(\mathcal{M})}$ are sent to the same simplex in $B_{\sigma(\mathcal{M})}$, but this is the case exactly when $w \sim_U w'$ in \mathcal{M} . The same line of argument works for $\sigma \circ \kappa$, and extends to morphisms. ■

Example 4. Let us illustrate the isomorphism using on the epistemic frame \mathcal{M} of Example 1. We now describe its equivalent representation as an epistemic covering $\sigma(\mathcal{M})$. In fact, the base B of the covering is the one depicted in Example 3. Then, we need a projective cset E , which can be thought of as an “exploded view” of B , depicted below. Its maximal simplices, labelled $\{w_0, \dots, w_5\} \cup \{w'_1\}$, correspond to the worlds of the original frame \mathcal{M} .



The covering $f : E \rightarrow B$ maps each world of E to its geometric representation in the base B . Note that both w_1 and w'_1 are mapped to the same edge of B : this is how we model non-proper behavior in epistemic frames. Moreover, world w_0 is represented by a single vertex, because only one agent is alive; it is mapped to the top blue vertex of B . This is how we model sub-worlds: since w_0 is a sub-world of w_1 in \mathcal{M} , $f(w_0)$ is a sub-simplex of $f(w_1)$ in $\sigma(\mathcal{M})$.

D. Semantics of distributed knowledge

We now use generalized epistemic frames, and equivalently epistemic coverings, as a model for the logic of distributed knowledge. The missing piece of data is to label the worlds with *atomic propositions*, in order to specify which facts about the system are either true or false in any given world.

Definition 7. A (*generalized*) *epistemic model* $\mathcal{M} = \langle M, \sim, L \rangle$ over the set of agents A consists of a generalized frame $\langle M, \sim \rangle$ together with *valuation* function $L : M \rightarrow \mathcal{P}(\text{At})$. A morphism of epistemic models $f : \mathcal{M} \rightarrow \mathcal{N}$ is a morphism of underlying frames that preserves valuations, that is, if $p \in L_M(w)$, then $p \in L_N(f(w))$. The category of generalized epistemic models is denoted \mathcal{KM}_A .

Remark 2. Morphisms in \mathcal{KM}_A are also known as functional simulations [23]. They are different from the morphisms used in [3]: there the valuations of atomic propositions were preserved and reflected, that is $L_M(s) = L_N(f(s))$, whereas in our definition $L_M(s) \subseteq L_N(f(s))$. They are also different from morphisms in [9]: there, morphisms can be seen as relations, as they are maps $f : M \rightarrow \mathcal{P}(N)$.

Given a generalized epistemic model, we can define the satisfaction relation as we did in Section II-A. Note however

that \sim_U is now part of the structure of the model, and might not be equal to $\bigcap_{a \in U} \sim_a$ in general.

$$M, w \models D_U \varphi \quad \text{iff} \quad M, w' \models \varphi \text{ for all } w' \in M \\ \text{such that } w \sim_U w'$$

Similarly, we can equip epistemic coverings with a valuation:

Definition 8. An *epistemic covering model* $X = \langle f : E \rightarrow B, \ell \rangle$ consists of an epistemic covering $f : E \rightarrow B$, together with a labelling $\ell : \max(E) \rightarrow \mathcal{P}(\text{At})$ that associates with each maximal simplex s of E a set of atomic propositions $\ell(s)$ that hold there. A morphism of epistemic covering models $\alpha : X \rightarrow Y$ is a morphism of epistemic coverings that preserves the labelling: $\ell(s) \subseteq \ell'(\uparrow \alpha(s))$, where $\uparrow \alpha(s)$ is the maximal simplex of E that contains $\alpha(s)$. We denote by \mathcal{EC}_A the category of epistemic covering models.

Given an epistemic covering model $X = \langle f, \ell \rangle$ together with a maximal simplex s in E , we can define the satisfaction relation $X, s \models \varphi$ by transporting what we did for generalized epistemic models via the isomorphism of Proposition 1.

$$X, s \models D_U \varphi \quad \text{iff} \quad X, s' \models \varphi \text{ for all maximal } s' \in E \\ \text{such that } \partial_U(f(s')) = \partial_U(f(s))$$

Proposition 1 can be readily extended to show that the categories of models are also isomorphic:

Theorem 1. *The category of epistemic covering models \mathcal{EC}_A is isomorphic to the category of generalized epistemic models \mathcal{KM}_A .*

Proof: We provide two functors κ, σ . On the underlying coverings and frames, they act as in Proposition 1. We just need to extend them to valuations. For $\kappa : \mathcal{EC}_A \rightarrow \mathcal{KM}_A$, for each maximal simplex s in E , we have an associated world w_s in $\kappa(f)$. We set $L_{\kappa(f)}(w_s) = \ell(s)$. Similarly, for $\sigma : \mathcal{KM}_A \rightarrow \mathcal{EC}_A$, there is an associated maximal simplex s_w for every world w . We set $\ell_{\sigma(M)}(s_w) = L(w)$. The rest of the proof is the same as in Proposition 1. ■

As expected, the satisfaction relations for both kinds of models yield the same result:

Lemma 2. *Given a pointed epistemic covering model (X, s) , we have $X, s \models \varphi$ iff $\kappa(X), w_s \models \varphi$. Conversely, given a pointed generalized epistemic model (M, w) , we have $M, w \models \varphi$ iff $\sigma(M), s_w \models \varphi$.*

Proof: We prove the first equivalence by induction on the structure of φ . The case of atomic propositions comes from the fact that we keep the labelling $L(w_s) = \ell(s)$. The case of boolean connectives is straightforward. For a formula of the form $D_U \varphi$ one can notice that we defined $w_s \sim_U w_{s'}$ iff $\partial_U(f(s)) = \partial_U(f(s'))$, which coincides with the semantics of D_U . The second equivalence follows from the first one, together with Theorem 1. ■

IV. PROPERTIES OF EPISTEMIC MODELS

The models presented in Section III-D are very versatile. Depending on what kind of applications we have in mind, we

might want to impose some extra properties on the structure of our models. For instance in distributed computing, the model is usually assumed to be a simplicial complex rather than a simplicial set, because a global state of the system is merely the sum of the local states of the agents, without any extra information. In this section, we define a number of interesting properties of epistemic frames, as well as their geometric counterpart, epistemic coverings. We will see how some of the results of previous papers on simplicial models arise as a special case of Proposition 1.

A. Properties of epistemic frames and coverings

Definition 9. An epistemic frame $\mathcal{M} = \langle M, \sim \rangle$ is said to


- have *trivial empty-group knowledge* if $\forall w, w' \in M. w \sim_{\emptyset} w'$;
- have *no empty worlds* if $\forall w \in M. \exists a \in A. w \sim_a w$;
- be *proper* if $(\overline{w} = \overline{w'} \wedge w \sim_{\overline{w}} w') \Rightarrow w = w'$;
- be *maximal* if $\forall w \in M. \forall U \subseteq \overline{w}. U \neq \emptyset \Rightarrow \exists w' \in M. (U = \overline{w'} \wedge w' \sim_U w)$;
- be *minimal* if $\forall w, w' \in M. (\overline{w} \subsetneq \overline{w'}) \Rightarrow w \not\sim_{\overline{w}} w'$;
- be *pure* if $\forall w \in M. \overline{w} = A$;
- have *standard group knowledge* if $\forall U \subseteq A. (\forall a \in U. w \sim_a w') \Rightarrow w \sim_U w'$.

Let us explain the meaning of these properties. In a frame with trivial empty-group knowledge, the empty group cannot distinguish any worlds. It models the idea that an empty group cannot measure anything, so all worlds have the same properties for it.

If a frame has no empty worlds, then there is an alive agent in every world, i.e., every possibility is observed by someone.

A frame is proper if every pair of worlds that has the same set of alive agents is distinguishable by some subgroup of agents. This also corresponds to the principle of observability: if even a maximal group cannot distinguish worlds, then they are the same. Notice that this allows sub-worlds.

Example 5. In the figure below, the leftmost frame has trivial empty-group knowledge, ($w_0 \sim_{\emptyset} w_1 \sim_{\emptyset} w_2$), but also has an empty world w_2 , and is not proper ($w_0 \sim_{\{a,b\}} w_1$). The middle frame has non-trivial empty-group knowledge ($w_1 \not\sim_{\emptyset} w_2$), has no empty worlds, and is not proper. The rightmost frame has non-trivial empty-group knowledge, has an empty world, but is proper ($w_0 \not\sim_b w_1$).




A frame is maximal if every world has a non-empty sub-world. A certain intuition comes from distributed computing: in a maximal frame, any number of agents may crash during the execution of a program, as long as at least one of them remains alive. Moreover, these crashes are undetectable.

A frame is minimal if there are no strict sub-worlds. Once again, this corresponds to a situation in distributed computing where crashes are detectable, that is when a process crashes, one of the remaining processes is aware of it.

A frame is pure if the set of alive agent in every world is the same. In such a situation, crashes are not allowed at all, and all agents always participate.

By “standard group knowledge”, we mean that group indistinguishability relations are generated by individual agents, that is, knowledge of the group is exactly the sum of individual agents’ knowledge. More formally, when $\sim_U = \bigcap_{a \in U} \sim_a$.

Example 6. In the figure below, the upper left frame is maximal, as it has all sub-worlds. The upper right frame is minimal as it does not have sub-worlds at all. The bottom left frame is pure as all its worlds have the same set of alive agents. The bottom right frame has non-standard group knowledge, contrary to all previous examples, since $w_0 \sim_a w_1$, $w_0 \sim_b w_1$, but $w_0 \not\sim_{a,b} w_1$.




Similarly, for epistemic coverings:

Definition 10. An epistemic covering $f : E \rightarrow B$ is said to


- have *trivial empty-group knowledge* if there is only one simplex of dimension -1 in B ;
- have *no empty worlds* if all maximal simplices of E have dimension ≥ 0 ;
- be *proper* if no two maximal simplices of E have the same image in B ;
- be *maximal* if every simplex in B is the image of a maximal simplex of E ;
- be *minimal* if the image of a maximal simplex of E is always a maximal simplex of B ;
- be *pure* if all maximal simplices of E have dimension $|A| - 1$;
- have *standard group knowledge* if B is a simplicial complex.

The intuition behind the definitions is exactly the same as in the case of frames. We give a few illustrative examples. In the picture below, the leftmost covering has trivial empty-group knowledge as there is only one (-1) -simplex in the base space (one dashed region); but it has empty worlds because of the maximal (-1) -simplex w_2 . The covering depicted in the middle has no empty worlds since all the maximal simplices have dimension 0 or 1. It, however, does not have trivial empty-group knowledge as it has two (-1) -simplices in the base space (depicted as two dashed regions). It is not proper either because both worlds w_0, w_1 label the same edge. The

rightmost covering is proper, as every simplex is annotated with at most one world. It has an empty world, w_2 , and does not have trivial empty-group knowledge.



In the examples below, the top left covering is maximal as every simplex is annotated, that is, every simplex has a maximal simplex that is sent to it. The top right covering is minimal because only maximal simplices are annotated. The bottom left covering is pure, as all annotated simplices are of the same dimension. All of the examples above have standard group knowledge since their base csets are in fact complexes. The bottom right covering has non-standard group knowledge as its base cset is not a simplicial complex.



As our terminology suggests, these properties of epistemic coverings are the geometric counterpart of the ones of epistemic frames that we defined previously.

Lemma 3. *The properties of Definition 9 agree with the ones of Definition 10 up to the equivalence in Proposition 1. Namely, if f is a covering of a certain type, then $\kappa(f)$ is of the same type, and conversely for σ .*

Proof: We only show two cases, as the proofs are very similar and just a matter of checking that we correctly translated the notions through the equivalence.

Consider a proper covering $f : E \rightarrow B$. In the frame $\kappa(f)$, two worlds w_1, w_2 are indistinguishable by group $\overline{w_1} = \overline{w_2}$ if and only if the simplices in E that correspond to w_1 and w_2 are sent to the same simplex in B . But, as f is proper, no two simplices have the same image, thus the frame is proper too. Now, consider a proper frame \mathcal{M} . By construction of the functor σ , two simplices s_1, s_2 in $E_{\mathcal{M}}$ of the same color are sent to the same simplex in $B_{\sigma(\mathcal{M})}$, if and only if corresponding worlds w_1, w_2 in \mathcal{M} $w_1 \sim_{\overline{w_1}} w_2$. But \mathcal{M} is proper, so it is never the case, thus $\sigma(\mathcal{M})$ is proper too.

Consider a covering $f : E \rightarrow B$ with standard group knowledge. Suppose that for a pair w_1, w_2 in the frame $\kappa(f)$, $w_1 \sim_a w_2$ for all a in some U . It means that $\partial_a(s_1) = \partial_a(s_2)$, where s_1 corresponds to w_1 and s_2 to w_2 . As B is a simplicial

complex, it follows that $\partial_U(s_1) = \partial_U(s_2)$, so $\kappa(f)$ has standard group knowledge. Take now a frame M with standard group knowledge. Suppose there are two worlds w_1, w_2 in M such that $\overline{w_1} = \overline{w_2} = U$ and for all $a \in U$, $[w_1]_a = [w_2]_a$. Since M has standard group knowledge, $[w_1]_U = \bigcap_{a \in U} [w_1]_a$. Thus, $[w_1]_U = [w_2]_U$, which means precisely that if two simplices in $B_{\sigma(M)}$ have the same set of vertices, then they are equal. ■

B. Subclasses of epistemic frames and coverings

The properties of epistemic frames (or, equivalently, epistemic coverings) defined in the previous section can be combined in various ways in order to define particular sub-classes of interest. The only restriction is that *minimal* and *maximal* are mutually exclusive properties (except in degenerate cases where the model only has empty worlds). It is also possible, and perhaps sometimes desirable, to consider models that are neither maximal nor minimal. Moreover, a frame/covering which is *pure* must also be minimal and have no empty world, so there is no need to specify the latter if we already have the former.

We can recover variants of simplicial models that have been defined in previous papers. Since they are usually concerned with distributed computing, where agents represent processes that can reason about the system, there is little interest in studying empty worlds or empty groups of agents. Thus, they all have *trivial empty-group knowledge* and *no empty worlds*. They also have *standard group knowledge* since they work with simplicial complexes instead of semi-simplicial sets. Moreover, all previous instances of simplicial models were always *proper*, not out of necessity, but because without the notion of covering, it is not possible to model non-proper behaviors in the geometric approach. On top of that:

- The original paper introducing simplicial models [3] was working with *pure* simplicial models (and thus, minimal).
- In the sequel [9], the “pure” assumption is dropped, but models are still assumed to be minimal since formulas are only evaluated at the facets.
- On the other hand, [10] and [11] study simplicial models that may not be pure but are maximal: every simplex is a possible world.
- An example that is not concerned about simplicial models, but which studies non-standard group knowledge, is [15]. Their models are also pure, proper, and have trivial empty-group knowledge.

In light of Lemma 3, we can easily restrict the equivalence of categories proved in Proposition 1 to the various sub-categories of frames and coverings. For instance, we can reformulate the main result of [3] as follows:

Corollary 1 ([3]). *The category of pure proper epistemic frames with standard group knowledge is isomorphic to the category of pure proper epistemic coverings with standard group knowledge.*

Proof: By Lemma 3, the restrictions of κ and σ to those sub-categories is well-defined. Since we still have $\kappa \circ \sigma = \text{id}$

and $\sigma \circ \kappa = \text{id}$ as proved in Proposition 1, this is still an isomorphism of categories. ■

V. AXIOMATIZATION OF THE VARIOUS SUB-CLASSES

A. Reasoning about alive and dead agents

As in [9], we can express within the language \mathcal{L}_D the fact that agents can be dead or alive. For any agent $a \in A$ and group of agents $U \subseteq A$, we define the following formulas:

$$\begin{aligned} \text{dead}(a) &:= K_a \text{false} & \text{alive}(a) &:= \widehat{K}_a \text{true} \\ \text{dead}(U) &:= \bigwedge_{a \in U} \text{dead}(a) & \text{alive}(U) &:= \widehat{D}_U \text{true} \end{aligned}$$

It is easy to check that they have the expected semantics: for epistemic frames, we have $M, w \models \text{alive}(U)$ iff $w \sim_U w$; and for epistemic coverings, $X, s \models \text{alive}(U)$ iff s is a V -simplex with $U \subseteq V$.

B. Axiomatization of epistemic covering models

We rely on the usual axiomatization of normal modal logics, with all propositional tautologies, closure by modus ponens, and the necessitation rule. On top of that, we add the following five axioms:

- **(K)** $D_U(\varphi \Rightarrow \psi) \Rightarrow (D_U\varphi \Rightarrow D_U\psi)$
- **(4)** $D_U\varphi \Rightarrow D_U D_U\varphi$
- **(B)** $\varphi \Rightarrow D_U \neg D_U \neg \varphi$
- **(Mono)** for $U \subseteq U'$, $D_U\varphi \Rightarrow D_{U'}\varphi$
- **(Union)** for U, U' , $\text{alive}(U) \wedge \text{alive}(U') \Rightarrow \text{alive}(U \cup U')$

We abbreviate **KB4_n + Mono + Union** as **EC_n**, which stands for the logic of epistemic coverings (as we will see). Notice that the difference between **KB4_n** and the more standard multi-agent epistemic logic **S5_n** is the absence of axiom **T**: $D_U\varphi \Rightarrow \varphi$. Here are a few examples of valid formulas in **EC_n** related to the life and death of agents.

- **EC_n** $\vdash \text{dead}(a) \Rightarrow K_a\varphi$: dead agents know everything. More generally, for $a \in U$, **EC_n** $\vdash \text{dead}(a) \Rightarrow D_U\varphi$.
- **EC_n** $\vdash \text{alive}(a) \Rightarrow K_a \text{alive}(a)$: Alive agents know that they are alive. The same holds for a group U of agents.
- **EC_n** $\vdash \text{alive}(U) \Rightarrow (D_U\varphi \Rightarrow \varphi)$: Axiom **T** is verified when restricted to groups of agents that are alive.

We are going to show soundness and completeness of **EC_n** with respect to generalized epistemic models. In order to show completeness, we use the standard approach: by providing a canonical model, see for example [23].

Definition 11. The *canonical generalized epistemic model* $M^c = \langle W^c, \sim, L \rangle$ is defined as follows:

- $W^c = \{\Gamma \mid \Gamma \text{ is a maximal consistent set of formulas}\}$.
- $\Gamma \sim_U \Delta$ iff $D_U\varphi \in \Gamma$ implies $\varphi \in \Delta$.
- $L(\Gamma) = \Gamma \cap \text{At}$.

Lemma 4 (Truth Lemma). *For any formula φ and any maximal consistent set of formulas $\Gamma \in M^c$, we have $\varphi \in \Gamma$ iff $M^c, \Gamma \models \varphi$.*

Proof: We proceed by induction on φ . The base case of atomic propositions holds by definition of M^c . For the boolean connectives, the proof is trivial.

Let us do the case of $D_U\varphi$. Assume that $D_U\varphi$ is in Γ and let Δ be an element of M^c such that $\Gamma \sim_U \Delta$. By definition of \sim , we have $\varphi \in \Delta$, so by induction hypothesis $M^c, \Delta \models \varphi$. As Δ is arbitrarily chosen, we have $M^c, \Gamma \models D_U\varphi$. Conversely, assume that $M^c, \Gamma \models D_U\varphi$ and suppose by contradiction that $D_U\varphi \notin \Gamma$. Then the set $\Delta^- = \{\neg\varphi\} \cup \{\psi \mid D_U\psi \in \Gamma\}$ is consistent. Indeed, suppose Δ^- is inconsistent. Then we have a proof of $\vdash \psi_1 \wedge \dots \wedge \psi_k \Rightarrow \varphi$, where $D_U\psi_i \in \Gamma$ for every i . Then, by axiom **K**, we can prove $\vdash D_U\psi_1 \wedge \dots \wedge D_U\psi_k \Rightarrow D_U\varphi$. As Γ is maximal consistent, this implies that $D_U\varphi \in \Gamma$, which contradicts the assumption. Thus, Δ^- is consistent, and by Lindenbaum's Lemma, we can extend it to a maximal consistent set $\Delta \supseteq \Delta^-$. By construction, $\Gamma \sim_U \Delta$, and by induction hypothesis, $M^c, \Delta \not\models \varphi$. This contradicts the assumption that $M^c, \Gamma \models D_U\varphi$. Therefore, $D_U\varphi \in \Gamma$, and this concludes the proof. ■

Lemma 5. *In the canonical generalized epistemic model M^c , for any $\Gamma \in M^c$, $U \subseteq \bar{\Gamma}$ iff $\text{alive}(U) \in \Gamma$.*

Proof: Suppose $U \subseteq \bar{\Gamma}$, $\Gamma \sim_U \Gamma$. Hence, $\Gamma \models \widehat{D}_U\text{true}$, as there is Γ which is U -accessible from Γ , which means that $\text{alive}(U)$ is in Γ by the Truth Lemma. Conversely, assume $\text{alive}(U) \in \Gamma$. Then, by Truth Lemma, $\Gamma \models \widehat{D}_U\text{true}$, so there is Δ , such that $\Delta \sim_U \Gamma$. By symmetry and transitivity of \sim_U , we have $\Gamma \sim_U \Delta$, so $U \subseteq \bar{\Gamma}$. ■

Theorem 2. *The logic EC_n is sound and complete with respect to the family of generalized epistemic models.*

Proof: Soundness is straightforward to check: the axioms **K**, **B** and **4** correspond exactly to the fact that the indistinguishability relations are PERs. Axiom **Mono** corresponds to condition (a) and axiom **Union** corresponds to condition (b) in the definition of generalized epistemic frames.

For completeness, consider the canonical model $M^c = \langle W^c, \sim, L \rangle$ from Definition 11. Axioms **K**, **B** and **4** ensure that \sim_U is a PER, as in the standard treatment of completeness. Similarly, axiom **Mono** ensures that the generated family of PERs is monotone: assume $U \subseteq U'$, $\Gamma \sim_{U'} \Delta$ and $D_U\varphi \in \Gamma$, then by **Mono**, $D_{U'}$ is also in Γ . As we assumed $\Gamma \sim_{U'}$, $\varphi \in \Delta$, so $\Gamma \sim_U \Delta$. Axiom **Union** ensures that the condition (b) is satisfied. Assume that $\Gamma \sim_U \Gamma$, $\Gamma \sim_{U'} \Gamma$ and $D_{U \cup U'}\varphi \in \Gamma$. First, we have that if $\Gamma \sim_U \Gamma$, then $\widehat{D}_U\text{true} \in \Gamma$. Indeed, by contraposition we have that for any ψ , $\psi \notin \Gamma$ entails $D_U\psi \notin \Gamma$. In particular, as Γ is consistent, $\text{false} \notin \Gamma$, so $D_U\text{false} \notin \Gamma$. By standard reasoning, $\neg D_U\text{false} \in \Gamma$, as we intended. Hence, $\text{alive}(U) \wedge \text{alive}(U') \in \Gamma$. By **Union** and modus ponens, $\text{alive}(U \cup U') \in \Gamma$. As the formula $\text{alive}(V) \Rightarrow (D_V\varphi \Rightarrow \varphi)$ is deducible when $V = U \cup U'$, we have $D_{U \cup U'}\varphi \Rightarrow \varphi \in \Gamma$. As we assumed $D_{U \cup U'}\varphi \in \Gamma$, by modus ponens $\varphi \in \Gamma$, which shows that the canonical model is indeed a generalized epistemic frame.

Applying the Lindenbaum Lemma and the Truth Lemma, any consistent formula φ holds in some state of W^c , and thus φ is satisfiable. ■

The following corollary follows from Lemma 2.

Corollary 2. *The logic EC_n is sound and complete with respect to epistemic covering models.*

C. Axiomatization of other sub-classes

We now show how to axiomatize the various sub-classes of epistemic models, as discussed in Section IV-B. As in the previous section, we first establish soundness and completeness for epistemic frames, and then we get the same result for epistemic coverings thanks to Lemmas 2 and 3. The proof is modular: each property of the frames (Definition 9) corresponds to an extra axiom. We denote by U^c the complement of the set of agents U , that is $A \setminus U$.

- **(NE)** $\bigvee_{a \in A} \text{alive}(a)$;
- **(P)** $\text{alive}(U) \wedge \text{dead}(U^c) \wedge \varphi \Rightarrow D_U(\text{dead}(U^c) \Rightarrow \varphi)$;
- **(Max)** for $U \neq \emptyset$, $\text{alive}(U) \Rightarrow \neg D_U\neg\text{dead}(U^c)$;
- **(Min)** $\text{alive}(U) \wedge \text{dead}(U^c) \Rightarrow D_U\text{dead}(U^c)$;
- **(Pure)** $\text{alive}(A)$.

However, there are no axioms related to trivial empty-group knowledge and standard group knowledge. This is because those two properties cannot be expressed in the language \mathcal{L}_D . More formally, we will see in Theorem 3 and Theorem 6 that EC_n is complete with respect to both classes of structures. As a consequence, they can be assumed “for free”, without additional axiom. Indeed, to show that the logic EC_n is sound and complete with respect to frames with trivial empty-group knowledge, we can use the strategy from [15]. The proof that standard group knowledge requires no extra axiom is the subject of Section V-D.

Remark 3. There are several interesting relationships between those axioms. Axiom **Pure**, which says that all agents are alive in all worlds, has many consequences. It entails the axioms **NE**, **Min**, and **Union**. Axiom **P** is greatly simplified and becomes $\varphi \Rightarrow D_A\varphi$, where A is the set of all agents. Furthermore, $\text{EC}_n + \text{Pure}$ together entail Axiom **T**, so that the logic KB4_n becomes S5_n when **Pure** is assumed. Another possible interaction is **P** + **Min**, which can be reformulated together as $\text{alive}(U) \wedge \text{dead}(U^c) \wedge \varphi \Rightarrow D_U\varphi$. This axiom appears in [9], in the particular case of $U = \{a\}$.

Now we prove soundness and completeness.

Theorem 3. *The logic $\text{EC}_n + \text{NE} + \text{P} + \text{Max}$ is sound and complete with respect to proper maximal epistemic models with trivial empty-group knowledge and no empty worlds.*

Proof: For brevity, call an epistemic model with the properties from the theorem statement a *good maximal* epistemic model. First, we show soundness. Suppose we are given a good maximal model \mathcal{M} . Axioms of EC_n hold in \mathcal{M} as it is an epistemic model in particular. Axiom **P** holds in \mathcal{M} because it is proper: assume that $\mathcal{M}, w \models \text{alive}(U) \wedge \text{dead}(U^c) \wedge \varphi$, that is, $\mathcal{M}, w \models \varphi$ and $\bar{w} = U$. Let $w' \in M$ such that $w' \sim_U w$. Assume then that $\mathcal{M}, w' \models \text{dead}(U^c)$, thus $\bar{w}' = U$, and by properness of \mathcal{M} , $w = w'$ and $\mathcal{M}, w' \models \varphi$. For axiom **NE**, as every world is non-empty, for some $a \in \bar{w}$, $K_a\varphi \Rightarrow \varphi$ holds. For axiom **Max**, suppose for some world $w \in M$, $\mathcal{M}, w \models \text{alive}(U)$, that is $U \subseteq \bar{w}$. As \mathcal{M} is maximal, there

is always a world w' , such that $w' \sim_U w$ and $\bar{w}' = U$. So $\neg D_U \neg \text{dead}(U^c)$ holds in w since $\text{dead}(U^c)$ holds in w' .

Now we show completeness. We define the canonical model M^c as in Definition 11, except that consistency now refers to the new logic. The proof of the Truth Lemma (Lemma 4) can easily be adapted. Let φ_0 be a consistent formula. We shall show that φ_0 is satisfiable in some good epistemic model. Let M^c be the canonical model for the logic, as in Theorem 2. We need to show that the canonical model is a good model. Recall that in the canonical model $\Gamma \sim_U \Delta$ iff for any φ , $D_U \varphi \in \Gamma \Rightarrow \varphi \in \Delta$. The model M^c is:

Proper: Suppose $\bar{\Gamma} = \bar{\Delta} = U$ and $\Gamma \sim_U \Delta$. We need to show that $\Gamma = \Delta$, that is for every formula φ , $\varphi \in \Gamma$ iff $\varphi \in \Delta$. Suppose $\varphi \in \Gamma$. Then, by axiom **P**, $D_U(\text{dead}(U^c) \Rightarrow \varphi)$ is in Γ . By the definition of \sim_U , it means that $\text{dead}(U^c) \Rightarrow \varphi \in \Delta$. By modus ponens, $\varphi \in \Delta$. Similarly, we can show that $\varphi \in \Delta \Rightarrow \varphi \in \Gamma$. Hence, $\Gamma = \Delta$ and the canonical model is proper.

Maximal: Let $U \subsetneq \bar{\Gamma}$. We want to exhibit a sub-world Δ of Γ , such that $\bar{\Delta} = U$. By Lemma 5, $\text{alive}(U) \in \Gamma$, so by axiom **Max**, maximality and consistency, $\neg D_U \neg \text{dead}(U^c)$ is in Γ . Then the set $\Delta^- = \{\text{dead}(U^c)\} \cup \{\psi \mid D_U \psi \in \Gamma\}$ is consistent, by the same reasoning as in the proof of the Truth Lemma. By Lindenbaum's Lemma, Δ^- can be extended to a maximal consistent set Δ . Moreover, $\Delta \sim_U \Gamma$ by construction, so $U \subseteq \bar{\Delta}$. Also, as $\text{dead}(U^c) \in \Delta$, $\bar{\Delta} \subseteq U$. Hence, $\bar{\Delta} = U$ and Δ is a sub-world of Γ .

No empty worlds: By axiom **NE**, maximality and consistency of any Γ , there is an agent $a \in A$, such that $\text{alive}(a) \in \Gamma$. It entails that $a \in \bar{\Gamma}$.

Trivial empty-group knowledge: The canonical model M^c , however, does not have trivial empty-group knowledge. Nevertheless, for every φ_0 , we can extract a sub-model, M_0^c , which consists of the set $\{\Gamma \in M^c \mid \Gamma \sim_\emptyset \Gamma_0\}$, where Γ_0 is some maximal consistent theory that contains φ_0 . By monotonicity of \sim , this restriction preserves all worlds U -accessible from Γ_0 and properties of \sim . Moreover, M_0^c has trivial empty-group knowledge. Thus, φ_0 holds in at Γ_0 in M_0^c , which is proper, maximal, has no empty worlds and trivial empty-group knowledge. ■

Inspecting the proof, one can see that the axioms and properties of the model are pairwise independent. Hence, any combination of those axioms yields a sound and complete axiom system for the corresponding class of models. Another example, with minimal models:

Theorem 4. *The logic $\text{EC}_n + \text{NE} + \text{P} + \text{Min}$ is sound and complete with respect to proper minimal epistemic models with trivial empty-group knowledge and no empty worlds.*

Proof: Call an epistemic that satisfies properties from the statement of the theorem a good minimal model. The proof of all clauses, except the correspondence between axiom **Min** and minimality, is the same as in Theorem 3. So for soundness, let us show the validity of axiom **Min** in any good minimal model. Consider a good minimal model \mathcal{M} and a world $w \in \mathcal{M}$. Assume that $\text{alive}(U) \wedge \text{dead}(U^c)$ holds in w . It means




Fig. 5. Variants of epistemic logic

that \bar{w} is exactly U . Then by minimality of \mathcal{M} , as w is not a subworld of any other world, in any w' , such that $w \sim_U w'$, we have $\bar{w}' = U$. Hence, $\mathcal{M}, w' \models \text{dead}(U^c)$ and $\mathcal{M}, w \models D_U \text{dead}(U^c)$, so axiom **Min** is valid.

For completeness, it remains only to verify that the canonical model M^c for this logic is minimal. Consider $\Gamma, \Delta \in M^c$, such that $\bar{\Gamma} \subsetneq \bar{\Delta}$ and $\bar{\Gamma} = U$. Suppose that $\Gamma \sim_U \Delta$. By Lemma 5, $\text{alive}(U) \wedge \text{dead}(U^c)$ belongs to Γ . By modus ponens and axiom **Min**, $D_U \text{dead}(U^c)$ is in Γ . By definition of \sim_U , $\text{dead}(U^c) \in \Delta$. But $\bar{\Delta}$ is strictly bigger than U , thus there $a \in U^c$ such that $a \in \bar{\Delta}$, so $\text{alive}(a) \in \Delta$ which contradicts with consistency of Δ . Thus, $\Gamma \not\sim_U \Delta$ as required. ■

For the case of pure models:

Theorem 5. *The logic $\text{EC}_n + \text{NE} + \text{P} + \text{Pure}$ is sound and complete with respect to proper pure epistemic models with trivial empty-group knowledge and no empty worlds.*

Proof: Extending the previous theorem, for soundness we just need to check validity of axiom **Pure** in a pure model. Take an epistemic model \mathcal{M} that has properties from the statement of the theorem. The formula $\bar{D}_U \text{true}$ holds in any world w because $w \sim_a w$ as \mathcal{M} is pure.

For completeness, we need to verify that the canonical model M^c is pure. As $\text{alive}(A)$ is in any Γ in M^c , by Lemma 5 we have $\bar{\Gamma} = A$, so M^c is pure. ■

The relationship between all the axiom systems that we consider is summarized in Figure 5.

D. Standard group knowledge

The last property that we did not axiomatize is the standard group knowledge, or, geometrically, the distinction between simplicial sets and simplicial complexes. In this section, we show that restricting to simplicial complexes does not require any additional axiom. For that purpose, we use a construction called *unraveling*, which turns a generalized epistemic frame M into a new frame $U(M)$ that is bisimilar to M and has standard group knowledge.

Let M be a generalized epistemic model. A *history* in M is a finite sequence of the form $h = (w_0, U_1, w_1, \dots, U_k, w_k)$ for some $k \geq 0$, such that $w_{i-1} \sim_{U_i} w_i$ for all $1 \leq i \leq k$ and U_i is a maximal such subgroup of agents. Notice that since epistemic models do not necessarily have standard

group knowledge, every pair of worlds w, w' can have several groups U maximal with respect to inclusion such that $w \sim_U w'$. For example, if $w \sim_U w'$, $w \sim_{U'} w'$ and both U, U' are maximal, then one has, in particular, two different histories (w, U, w') and (w, U', w') . We write $\text{last}(h) = w_k$ for the last element of a history, and we write $h \rightarrow_U h'$ if $h' = (h, U, w_{k+1})$ with $U \subseteq U'$.

Definition 12. The *unraveling* of M is a generalized epistemic model $U(M) = (H, \sim^u, L^u)$ defined as follows:

- H is the set of histories of M ,
- \sim_U^u is the transitive and symmetric closure of \rightarrow_U , i.e., $\sim_U^u = (\rightarrow_U \cup \leftarrow_U)^*$,
- $L^u(h) = L(\text{last}(h))$.

The unraveling construction is similar to the tree unraveling from [24], adjusted to the case of multiple relations related by monotonicity.

Lemma 6. *Let M be a generalized epistemic model. Then its unraveling $U(M)$ is a generalized epistemic model.*

Proof: It is easy to see that every \sim_U^u is a PER, as it is transitive and symmetric closure. If $U \subseteq U'$, $\rightarrow_{U'} \subseteq \rightarrow_U$ by definition, so $\sim_{U'}^u \subseteq \sim_U^u$, that is \sim^u satisfies the compatibility condition. Now, suppose that $h \sim_U^u h$ and $h \sim_{U'}^u h$. By definition of \sim^u , it means that there is a sequence of worlds in M such that $\text{last}(h) \sim_U \dots \sim_U \text{last}(h)$, so, by transitivity of \sim_U , $\text{last}(h) \sim_U \text{last}(h)$. Similarly, $\text{last}(h) \sim_{U'} \text{last}(h)$. As M is a generalized epistemic model, it follows that $\text{last}(h) \sim_{U \cup U'} \text{last}(h)$. Thus, there is a history $h' = (h, U \cup U', \text{last}(h))$, and $h \sim_{U \cup U'}^u h'$. By transitivity and symmetry of $\sim_{U \cup U'}^u$, we have $h \sim_{U \cup U'}^u h$, thus concluding that $U(M)$ is closed under union of the groups of alive agents. ■

Lemma 7. *For every model M , its unraveling $U(M)$ has standard group knowledge.*

Proof: We need to show that for any two h, h' in $U(M)$, if $h \sim_U^u h'$ and $h \sim_{U'}^u h'$, then $h \sim_{U \cup U'}^u h'$. There are two cases: if $h = h'$ and $h \neq h'$. For the first case, the statement holds since $U(M)$ is a generalized epistemic model by Lemma 6.

For the second case, notice first that \rightarrow_V respects the ordering of histories with respect to the prefix relation: if $h \rightarrow_V h'$, then h is a prefix of h' . The prefix relation forms a tree on the set of histories, and it implies that if $h \sim_V^u h'$, then there is a unique non-redundant path $h \leftarrow_V \dots \leftarrow_V h'' \rightarrow_V \dots \rightarrow_V h'$ from h to h' that witnesses it, where h'' is the common prefix of h and h' . Moreover, this path is the same for any $V \subseteq A$. We can write $h = (h'', V_1, w_1, \dots, V_n, w_n)$ and $h' = (h'', V'_1, w'_1, \dots, V'_m, w'_m)$. As it is the same path for both U and U' , $U \cup U' \subseteq V_i$ and $U \cup U' \subseteq V'_j$ for all i, j . In particular, we have that $\text{last}(h) = w_n \sim_{U \cup U'} \dots \sim_{U \cup U'} \text{last}(h'') \sim_{U \cup U'} \dots \sim_{U \cup U'} w'_m = \text{last}(h')$. Thus, $h \sim_{U \cup U'}^u h'$, and $U(M)$ has standard group knowledge. ■

Remark 4. In the proof we have also shown that if $h \sim_U^u h'$, then $\text{last}(h) \sim_U \text{last}(h')$.

Given two generalized epistemic models M, N , we say that morphism $p : M \rightarrow N$ is a *functional bisimulation* if the following conditions hold:

- (atoms) for any $w \in M$, $L_N(f(w)) = L_M(w)$;
- (forth) for all $U \subseteq A$, if $w \sim_U^M w'$, then $f(w) \sim_U f(w')$;
- (back) for all $U \subseteq A$, if $f(w) \sim_U^N v'$, then there is $w' \in M$ such that $f(w') = v'$ and $w \sim_U^M w'$.

The definition we give is an extension of the standard notion of bisimulation, which links structural similarity of models with validity of formulas. In particular, we have the following proposition by adapting the standard construction (see [23]).

Proposition 2. *If $f : M \rightarrow N$ is a functional bisimulation, then for any formula φ , we have that $M, w \models \varphi$ if and only if $N, f(w) \models \varphi$.*

This is the essential proposition that allows us to show that \mathbf{EC}_n is sound and complete with respect to frames with standard group knowledge.

Lemma 8. *For every generalized epistemic model M , its unraveling $U(M)$ is bisimilar to M .*

Proof: We shall show that $\text{last} : U(M) \rightarrow M$ is a functional bisimulation. First, by the definition of L^u , $L^u(h) = L(\text{last}(h))$, so the atomic proposition are preserved.

Suppose now that $h \sim_U^u h'$. By Remark 4, $\text{last}(h) \sim_U h'$, so (forth) condition is satisfied.

For (back) condition, suppose that $\text{last}(h) \sim_U w'$. Then there is a history $h' = (h, U', w')$, such that $U \subseteq U'$. Clearly, $h \rightarrow_U h'$, so $h \sim_U^u h'$, and the (back) condition holds too, thus concluding that $\text{last} : U(M) \rightarrow M$ is a bisimulation. ■

Theorem 6. *The logic \mathbf{EC}_n is sound and complete with respect to models with standard group knowledge.*

Proof: Soundness is straightforward, as models with standard group knowledge are generalized epistemic models in particular.

Consider the canonical model M^c for \mathbf{EC}_n from Theorem 2. It is shown that M^c is a generalized epistemic frame. By Lemma 6, $U(M^c)$ is a generalized epistemic frame, and by Lemma 7 it has standard group knowledge. For any formula φ , there is Γ in M^c , such that $M^c, \Gamma \models \varphi$. Since $\text{last} : U(M^c) \rightarrow M^c$ is a bisimulation by Lemma 8, for any $h \in U(M^c)$ such that $\text{last}(h) = \Gamma$, we have $U(M^c), h \models \varphi$. Thus, any formula φ is valid in a model with classic group knowledge, which concludes completeness. ■


E. Topological interpretation

The completeness result with respect to models with standard group knowledge points us towards the study of expressivity of modal logic with respect to topological properties of csets. In particular, a natural question is what kind of similarity between csets is induced by a bisimulation? In this subsection, we give a partial answer to this question.

Let us restrict our attention to good maximal epistemic models. By the isomorphism we have shown, it is the same as to consider csets where every simplex is a unique world. Let

$\alpha : X \rightarrow Y$ be a functional bisimulation. We can translate the definition of a bisimulation in terms of epistemic coverings: it means that the valuation of a simplex $s \in X$ coincides with the valuation of its image $\alpha(s) \in Y$; if two simplices s, s' in X have a common U -subsimplex, then $\alpha(s)$ and $\alpha(s')$ have a common U -subsimplex; if $\alpha(s)$ has a common U -subsimplex with t' in Y , then there is a simplex s' in X that shares a U -subsimplex with s , and $\alpha(s') = t'$. The second condition is always satisfied by any morphism of csets. The last condition is more interesting, as it can be interpreted as a certain type of *lifting condition* that is ubiquitous in topology. This point of view can be traced back to [25].

Without giving a proof, we remark that the forgetful functor from csets to simplicial sets (which just forgets colors of vertices) sends functional bisimulations of csets with *unique lifting* (corresponding to having a unique simplex s' with $\alpha(s') = t'$ in the condition above) to *simplicial coverings* (see [26]). An example of such a bisimulation that can be seen as a topological covering is depicted. It represents the classical example of a covering of a circle by an infinite helix.



VI. CONCLUSIONS AND FUTURE WORK

We have introduced a very general class of epistemic models based on generalized epistemic frames, and their geometric counterpart based on simplicial sets, epistemic covering models. These models subsume many variants of simplicial models found in the literature. We made good use of this generality by establishing a close connection between properties of the models and axioms of the logic. This yields soundness and completeness for a variety of logics (Fig. 5).

A notable fact that we prove is that every model based on simplicial sets is equivalent modulo bisimulation to a model based on simplicial complexes. This unveils interesting geometric considerations: epistemic modal logics do not distinguish simplicial sets from their coverings because of the local nature of the knowledge operators. Could we define a logic that is able to capture better the global geometry of the model?

In future work, we aim at putting these logics in action for, in particular, distributed computing applications such as in e.g. [3], [4], [27], [7], and bridge the gap between the geometric interpretations [28] of distributed problems such as sensor networks, as touched upon in Example 2, with their logical interpretations.

REFERENCES

[1] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi, *Reasoning About Knowledge*. Cambridge, MA, USA: MIT Press, 2003.
 [2] M. Herlihy, D. Kozlov, and S. Rajsbaum, *Distributed Computing Through Combinatorial Topology*. Francisco, CA, USA: Morgan Kaufmann, 2013.

[3] É. Goubault, J. Ledent, and S. Rajsbaum, “A simplicial complex model for dynamic epistemic logic to study distributed task computability,” *Inf. Comput.*, vol. 278, 2021.
 [4] H. v. Ditmarsch, É. Goubault, J. Ledent, and S. Rajsbaum, “Knowledge and simplicial complexes,” in *Philosophy of Computing*, B. Lundgren and N. A. Nuñez Hernández, Eds., vol. 143. Springer 2022, p. 1–50.
 [5] H. van Ditmarsch, É. Goubault, M. Lazić, J. Ledent, and S. Rajsbaum, “A dynamic epistemic logic analysis of equality negation and other epistemic covering tasks,” *J. Log. Algebraic Methods Program.*, vol. 121, 2021.
 [6] K. Yagi and S. Nishimura, “Logical obstruction to set agreement tasks for superset-closed adversaries,” *CoRR*, vol. abs/2011.13630, 2020.
 [7] S. Nishimura, “Proving unsolvability of set agreement task with epistemic mu-calculus,” *CoRR*, vol. abs/2205.06452, 2022.
 [8] J. Y. Halpern and Y. Moses, “Knowledge and common knowledge in a distributed environment,” *J. ACM*, vol. 37, no. 3, pp. 549–587, 1990.
 [9] É. Goubault, J. Ledent, and S. Rajsbaum, “A simplicial model for KB4_n: Epistemic logic with agents that may die,” in *STACS 2022*, 2022, p. 33:1–33:20.
 [10] H. van Ditmarsch, “Wanted dead or alive: Epistemic logic for impure simplicial complexes,” in *WoLLIC 2021*, Lecture Notes in Computer Science, vol. 13038. Springer, 2021, pp. 31–46.
 [11] R. F. Randrianomensoa, H. van Ditmarsch, and R. Kuznets, “Impure simplicial complexes: Complete axiomatization,” *CoRR*, vol. abs/2211.13543, 2022.
 [12] C. Dwork and Y. Moses, “Knowledge and common knowledge in a byzantine environment: Crash failures,” *Inf. Comput.*, vol. 88, no. 2, pp. 156–186, 1990.
 [13] J. Gerbrandy, “Distributed knowledge,” in *Formal Semantics and Pragmatics of Dialogue*, vol. 98, 1998, pp. 111–124.
 [14] A. Baltag and S. Smets, “Learning what others know,” in *LPAR 2020: EPiC Series in Computing*, vol. 73, 2020, pp. 90–119.
 [15] —, “Correlated Knowledge: An Epistemic-Logic View on Quantum Entanglement,” vol. 49, no. 12, pp. 3005–3021.
 [16] J. Garson, “Modal Logic,” in *The Stanford Encyclopedia of Philosophy*, Summer 2021 ed., E. N. Zalta, Ed. Metaphysics Research Lab, Stanford University, 2021.
 [17] P. Goerss and J. Jardine, *Simplicial Homotopy Theory*, ser. Modern Birkhäuser Classics. Birkhäuser Basel, 2009.
 [18] J. May, *Simplicial Objects in Algebraic Topology*, ser. Chicago Lectures in Mathematics. University of Chicago Press, 1993.
 [19] G. Friedman, “An elementary illustrated introduction to simplicial sets,” 2008, V7 25 May 2021.
 [20] J. C. Baez and A. E. Hoffnung, “Convenient categories of smooth spaces,” 2008.
 [21] S. Eilenberg and J. A. Zilber, “Semi-simplicial complexes and singular homology,” *Annals of Mathematics*, pp. 499–513, 1950.
 [22] E. Riehl, “A leisurely introduction to simplicial sets.”
 [23] P. Blackburn, M. de Rijke, and Y. Venema, *Modal Logic*, ser. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press.
 [24] H. Sahlgqvist, “Completeness and Correspondence in the First and Second Order Semantics for Modal Logic,” in *Studies in Logic and the Foundations of Mathematics*, Elsevier, vol. 82, pp. 110–143.
 [25] A. Joyal, M. Nielsen, and G. Winskel, “Bisimulation from Open Maps,” vol. 127, no. 2, pp. 164–185.
 [26] P. Gabriel and M. Zisman, *Calculus of Fractions and Homotopy Theory*, Ergebnisse Der Mathematik Und Ihrer Grenzgebiete. 2. Folge. Springer-Verlag.
 [27] É. Goubault, M. Lazić, J. Ledent, and S. Rajsbaum, “Wait-free solvability of equality negation tasks,” in *33rd International Symposium on Distributed Computing, DISC 2019, October 14-18, 2019, Budapest, Hungary*, LIPIcs, J. Suomela, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019, pp. 21:1–21:16.
 [28] P. Dlotko, R. Ghrist, M. Juda, and M. Mrozek, “Distributed computation of coverage in sensor networks by homological methods,” *Applicable Algebra in Engineering, Communication and Computing*, 2012.