# A Simplicial Complex Model for Dynamic Epistemic Logic to study Distributed Task Computability

Éric Goubault[a], Jérémy Ledent[b], Sergio Rajsbaum[c]

[a]*LIX, CNRS, École Polytechnique, Institute Polytechnique de Paris*
*1 rue Honoré d'Estienne d'Orves, 91120 Palaiseau, France*
[b]*Department of Computer and Information Sciences, University of Strathclyde*
*26 Richmond Street, Glasgow G1 1XH, United Kingdom*
[c]*Instituto de Matemáticas, UNAM*
*Ciudad Universitaria Mexico 04510, Mexico*

**Abstract**

The usual $\mathbf{S5_n}$ epistemic model for a multi-agent system is based on a Kripke frame, which is a graph whose edges are labeled with agents that do not distinguish between two states. We propose to uncover the higher dimensional information implicit in this structure, by considering a dual, *simplicial complex model.* We use dynamic epistemic logic (DEL) to study how an epistemic simplicial complex model changes after a set of agents communicate with each other. We concentrate on an action model that represents the so-called *immediate snapshot* communication patterns of asynchronous agents, because it is central to distributed computability (but our setting works for other communication patterns). There are topological invariants preserved from the initial epistemic complex to the one after the action model is applied, which determine the knowledge that the agents gain after communication. Finally, we describe how a distributed task specification can be modeled as a DEL action model, and show that the topological invariants determine whether the task is solvable. We thus provide a bridge between DEL and the topological theory of distributed computability, which studies task solvability in a shared memory or message passing architecture.

*Keywords:* Dynamic epistemic logic, Distributed computing, Simplicial complexes

## 1. Introduction

Modal epistemic logic has been widely studied to reason about multi-agent systems. The typical language extends propositional logic by adding $n$ modalities $K_i$ representing the knowledge of agent $i$, and typically a Kripke semantics is used.

### *1.1. Three goals*

The usual $\mathbf{S5_n}$ epistemic model for a multi-agent system is based on a *Kripke frame,* a graph whose edges are labeled with agents that do not distinguish between two states. A *Kripke model* is obtained by labelling the states of the graph with atomic propositions, representing the knowledge of the agents about a given situation. Our first goal is to expose the topological information implicit in a Kripke model, replacing it by its dual, a *simplicial complex model.* We prove that such a simplicial model is very closely related to the usual Kripke model: there is an equivalence of categories between the two structures. Thus, simplicial models retain the nice properties of Kripke models, such as soundness and completeness w.r.t. to the logic $\mathbf{S5_n}$.

Our second goal is to show that when agents communicate with each other, some topological invariants of the initial simplicial complex model are preserved. In this context, a very natural setting is dynamic epistemic logic (DEL) [1, 2] with *action models* [3]. We extend the duality to this setting by defining a simplicial version of action models and a corresponding product update operator. Thus, the product update of an initial simplicial model $\mathcal{I}$ and an action model $\mathcal{A}$ yields a simplicial model $\mathcal{I}[\mathcal{A}]$. The possible patterns of communication permitted by the action model determine the topological invariants of $\mathcal{I}$ that are preserved in $\mathcal{I}[\mathcal{A}]$.

A third goal is to show how DEL can be used as a specification tool. A *task* is the equivalent of a function in distributed computability [4]. Agents start with an input value, and after communicating with the others, produce an output value. The task defines the possible inputs to the agents, and for each set of inputs, it specifies the set of outputs that the agents may produce. An important example is the *consensus* task, where all the agents must agree on one of their input values. We use DEL in a novel way, to represent the task itself. A Kripke model $\mathcal{I}$ represents the possible initial states of the system. The task is specified by an action model $\mathcal{T}$, which describes the output values that the agents should be able to produce, as well as preconditions specifying which inputs are allowed to produce which outputs. The product update of the input model $\mathcal{I}$ with $\mathcal{T}$ yields an epistemic model $\mathcal{I}[\mathcal{T}]$ representing the knowledge that the agents should acquire in order to solve the task. Once the task is specified, given an action model $\mathcal{A}$ that represents some distributed protocol, the product update of $\mathcal{I}$ with $\mathcal{A}$ yields a Kripke model $\mathcal{I}[\mathcal{A}]$ that models how agents perceive the world after the protocol has been executed. The protocol $\mathcal{A}$ *solves* the task if there exists a morphism $\delta$ that makes the diagram on the right commute (Definition 2). This intuitively happens when there is sufficient knowledge in $\mathcal{I}[\mathcal{A}]$ to solve the task.

$$\begin{tikzcd}
 & \mathcal{I}[\mathcal{A}] \arrow[dl, "\pi_{\mathcal{I}}"'] \arrow[d, dashed, "\delta"] \\
\mathcal{I} & \mathcal{I}[\mathcal{T}] \arrow[l, "\pi_{\mathcal{I}}"]
\end{tikzcd}$$

### 1.2. Distributed computability

Beyond the three goals stated above, we are concerned with constructing a general framework that connects epistemic logic and *fault-tolerant distributed computability*, because its intimate relation to topology is well understood [5]. Also, DEL has applications to numerous research areas, but to the best of our knowledge it has not been used to study fault-tolerant distributed computing systems.

In one direction, uncovering the higher-dimensional topological structure hidden in Kripke models allows us to transport methods that have been used successfully in the algebraic topological approach to fault-tolerant distributed computability [5] to the realm of DEL. In particular, the knowledge gained by applying an action model is intimately related to how well it preserves the topology of the initial model. The benefit in the other direction is in providing a formal epistemic logic semantics to distributed task computability. This allows one to understand better the abstract topological arguments in terms of how much knowledge is necessary to solve a task.

Many distributed computing systems have been considered in the past e.g. [4, 6], and it is well-known that different systems have different computational power to solve tasks. The power of the system depends on the way in which processes communicate with each other (the shared memory or message passing primitives available), the failures that can happen (how many, and their severity, such as crash or Byzantine), and their relative speeds (synchronous, partially synchronous or asynchronous). And these features in turn induce the topological invariants that determine the computational power of the system. We concentrate on the most basic type of system, where *all* topological properties are preserved, which is why it is the most important case. We describe below one such model, where the communication is by simple read/write shared-memory operations, and asynchronous processes may fail by crashing. We select this system to illustrate our framework with a concrete example, but we stress that our framework can be applied to other systems considered in the past.

We concentrate on the *immediate snapshot* setting of asynchronous wait-free shared read/write memory, to illustrate the ideas with a concrete, simple example. The setting is simple, yet it is central to distributed computability, and has been thoroughly studied since papers such as [7, 8]. We define a corresponding *immediate snapshot action model*. This model indeed fully preserves the topology of the initial complex, as observed in earlier papers e.g. [7, 8]. This model corresponds to *wait-free* asynchronous processes, which means that the processes run at an arbitrary speed, independent from the others, and never wait for events to happen in other processes.

There are known equivalences between task solvability in the immediate snapshot model and other shared memory and message passing systems. The immediate snapshot model can be used as a basis to study task solvability in other more complex models, e.g. where the number of processes that can crash is bounded or even where Byzantine failures are possible [5].

### *1.3. Related work*

Modal logics of knowledge have been used as a formal tool for specifying and reasoning about multiagent systems in distributed computing [9], artificial intelligence [10], and economics [11]. The most common semantics is the relational semantics which were made popular by Kripke [12]. The idea of "possible worlds" to describe an ontology goes back to Leibniz. Topological semantics of modal logic is historically the first one and can be traced back to the 1930's and the subsequent influential paper [13]. The "possible worlds" correspond to points of a topological space, and there is a well known connection between topological spaces and Kripke models, e.g. [14]. Our approach seems to be new though, in the sense that instead of taking as the basic element a "possible world", we take as basic element the "perspectives of the agents" (local states instead of global states) about such a world. Thus, a possible world is represented by a *set* of compatible perspectives by the agents, about the world being represented. An accessibility relation is then induced when an agent has the same perspective of two different worlds. See Examples 1, 2, and 3 in the next section.

Work on knowledge and distributed systems is of course one of the inspirations of the present work [9], especially where connectivity [15, 16] is used.

In the *interpreted systems* used in this approach (a kind of computationally grounded semantics), a world is a vector composed of the local states of the agents and the environment's state. Interpreted systems and Kripke models for multiagent systems are closely related [17, 18]. The frame structure obtained is still a one-dimensional graph, while in our structure the local states are the vertices of a higher-dimensional structure, a simplicial complex (we do not consider the environment's state).

The authors know of no previous work using DEL [1, 2] to study such systems, and neither on directly connecting the combinatorial topological setting of [5] with Kripke models. Let us mention, though, that there are other categorical connections between Kripke frames and geometry e.g. [17]. Also, modal languages have been used to define regions in space (topological spaces) [14, 19]. Work on relating the interpreted systems of [9] with Kripke semantics, and applications to synchronous distributed systems communicating by broadcast [18], do not consider failures nor the underlying topological structure of the states of the system. In [20], the author proposes a variant of (non dynamic) epistemic logic for a restricted form of wait-free task specification that cannot account for important tasks such as consensus.

Similar to [21], we show that even though a problem may not explicitly mention the agents' knowledge, it can in fact be restated as knowledge gain requirements. Nevertheless, we exploit the "runs and systems" framework in an orthogonal way, and the knowledge requirements we obtain are about inputs; common knowledge in the case of consensus, but other forms of nested knowledge for other tasks. In contrast, the *knowledge of precondition principle* of [21] implies that common knowledge is a necessary condition for performing simultaneous actions.

DEL is often thought of as inherently being capable of modeling only agents that are synchronous, but as discussed in [22], this is not the case. More recently, [23] proposes a variant of public announcement logic for asynchronous systems that introduces two different modal operators for sending and receiving messages. A similar approach of asynchronous announcement has been taken in [24], furthermore showing that on multi-agent epistemic models, each formula in asynchronous announcement logic is equivalent to a formula in epistemic logic. We show here that DEL can naturally model the knowledge in an asynchronous distributed system, at least as far as it is concerned with task solvability. Further work is needed to study more in depth the knowledge that is represented in this way.

Finally, note that our formulation of carrier maps as products has been partially observed in [25].

*1.4. Organization*

The rest of this paper is organized as follows. Section 2 presents the new model for epistemic logic, based on chromatic simplicial complexes, and shows that the corresponding category is equivalent to the category of Kripke models. Also, we show how to reformulate the usual semantics of formulas in Kripke models, in terms of simplicial models. Section 3 describes our adaptation of Dynamic Epistemic Logic (DEL) to simplicial models, and the particular action model for distributed computing that we use to illustrate our ideas. Section 4 explains how to use DEL in an innovative way, to specify a distributed task. Then it presents several well-known task solvability results in the distributed computing literature, from the epistemic logic perspective, using our framework. Conclusions and open problems are in Section 5.

## 2. A simplicial model for epistemic logic

We describe here the new kind of model for epistemic logic, based on chromatic simplicial complexes. The link between $\mathbf{S5_n}$, DEL and distributed computing will be developed in the next sections. We begin in Section 2.1 developing the dual of Kripke frames in terms of simplicial complexes. Then we extend the ideas to models, proving the equivalence of the simplicial model and the Kripke model categories, in Section 2.2. Also, we reformulate the usual semantics of formulas in Kripke models, in terms of simplicial models. The equivalence is under a *locality* restriction, which is discussed and lifted in Section 2.3.

*Syntax.* Let $AP$ be a countable set of propositional variables and $A$ a finite set of agents. The language of epistemic logic formulas $\mathcal{L}_\mathcal{K}(A, AP)$, or just $\mathcal{L}_\mathcal{K}$ if $A$ and $AP$ are implicit, is generated by the following BNF grammar:

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_a\varphi \qquad p \in AP,\ a \in A$$

We are also interested in epistemic logic with *common knowledge* operators. We write $\mathcal{L}_{\mathcal{CK}}(A, AP)$ for the language of these formulas, defined by:

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_a\varphi \mid C_B\varphi \qquad p \in AP,\ a \in A,\ B \subseteq A$$

We denote by $\mathbf{S5_n}$ and $\mathbf{S5C_n}$ the deduction systems associated to $\mathcal{L}_\mathcal{K}$ and $\mathcal{L}_{\mathcal{CK}}$, respectively. The proof theory of epistemic logics can be found in [2]; in the rest of the paper, we focus mostly on studying models. In the following, we work with $n+1$ agents, and write $A = \{a_0, \ldots, a_n\}$.

*2.1. Simplicial complexes and Kripke frames*

Given a set $V$, a *simplicial complex* $C$ is a family of non-empty finite subsets of $V$ such that for all $X \in C$, $Y \subseteq X$ implies $Y \in C$. We say $Y$ is a *face* of $X$. Elements of $V$ (identified with singletons) are called *vertices*. Elements of $C$ are *simplexes*, and those which are maximal w.r.t. inclusion are *facets*. The set of vertices of $C$ is noted $\mathcal{V}(C)$, and the set of facets $\mathcal{F}(C)$. The *dimension* of a simplex $X \in C$ is $|X| - 1$, and a simplex of dimension $n$ is called an *$n$-simplex*. A simplicial complex $C$ is *pure* if all its facets are of the same dimension, $n$. In this case, we say $C$ is of dimension $n$. Given the set $A$ of agents (that we will represent as colors), a *chromatic simplicial complex* $\langle C, \chi\rangle$ consists of a simplicial complex $C$ and a coloring map $\chi : \mathcal{V}(C) \to A$, such that for all $X \in C$, all the vertices of $X$ have distinct colors.

Let $C$ and $D$ be two simplicial complexes. A *simplicial map* $f : C \to D$ maps the vertices of $C$ to vertices of $D$, such that if $X$ is a simplex of $C$, $f(X)$ is a simplex of $D$. A *chromatic simplicial map* between two chromatic simplicial complexes is a simplicial map that preserves colors. Let $\mathcal{S}_A$ be the category of pure chromatic simplicial complexes on $A$, with chromatic simplicial maps for morphisms.

A *Kripke frame* $M = \langle S, \sim\rangle$ over a set $A$ of agents consists of a set of states $S$ and a family of equivalence relations on $S$, written $\sim_a$ for every $a \in A$. Two states $u, v \in S$ such that $u \sim_a v$ are said to be *indistinguishable* by $a$. A Kripke frame is *proper* if any two states can be distinguished by at least one agent. Notice that being proper means that the intersection of all equivalence relations $\sim_a$ is the identity; this may reveal interesting parallels with distributed knowledge (a formula that is true in all states in the intersection relation), see e.g. [9]. Let $M = \langle S, \sim\rangle$ and $N = \langle T, \sim'\rangle$ be two Kripke frames. A *morphism* from $M$ to $N$ is a function $f$ from $S$ to $T$ such that for all $u, v \in S$, for all $a \in A$, $u \sim_a v$ implies $f(u) \sim'_a f(v)$. We write $\mathcal{K}_A$ for the category of proper Kripke frames, with morphisms of Kripke frames as arrows.

The following theorem states that we can canonically associate a proper Kripke frame with a pure chromatic simplicial complex, and vice versa. In fact, this correspondence extends to morphisms, and thus we have an equivalence of categories, meaning that the two structures contain the same information.

**Theorem 1.** *$\mathcal{S}_A$ and $\mathcal{K}_A$ are equivalent categories.*

*Proof.* We construct functors $F : \mathcal{S}_A \to \mathcal{K}_A$ and $G : \mathcal{K}_A \to \mathcal{S}_A$ as follows.

Let $C$ be a pure chromatic simplicial complex on the set of agents $A$. Its associated Kripke frame is $F(C) = \langle S, \sim \rangle$, where $S$ is the set of facets of $C$, and the equivalence relation $\sim_a$, for each $a \in A$, is defined as $X \sim_a Y$ (for $X$ and $Y$ facets of $C$) if $a \in \chi(X \cap Y)$.

For a morphism $f : C \to D$ in $\mathcal{S}_A$, we define $F(f) : F(C) \to F(D)$ that takes a facet $X$ of $C$ to its image $f(X)$, which is a facet of $D$ since $f$ is a chromatic map. Assume $X$ and $Y$ are facets of $C$ such that $X \sim_a Y$ in $F(C)$, that is, $a \in \chi(X \cap Y)$. So there is a vertex $v \in \mathcal{V}(C)$ such that $v \in X \cap Y$ and $\chi(v) = a$. Then $f(v) \in f(X) \cap f(Y)$ and $\chi(f(v)) = a$, so $a \in \chi(f(X) \cap f(Y))$. Therefore, $f(X) \sim_a f(Y)$, and $F(f)$ is a morphism of Kripke frames.

Conversely, consider a Kripke frame $M = \langle S, \sim \rangle$ on the set of agents $A = \{a_0, \ldots, a_n\}$. Intuitively, what we want to do is take one $n$-simplex $\{v_0^s, \ldots, v_n^s\}$ for each $s \in S$, and glue them together according to the indistiguishability relation. Formally, let $V = \{v_i^s \mid s \in S, 0 \leq i \leq n\}$, and equip it with the equivalence relation $R$ defined by $v_i^s \; R \; v_i^{s'}$ if and only if $s \sim_{a_i} s'$. Then define $G(M)$ whose vertices are the equivalence classes $[v_i^s] \in V/R$, and whose simplexes are of the form $\{[v_0^s], \ldots, [v_n^s]\}$ for $s \in S$, as well as their sub-simplexes. The coloring map is given by $\chi([v_i^s]) = a_i$. It is a well-defined chromatic simplicial complex since all elements of an equivalence class of $R$ have the same color. The facets are exactly the $\{[v_0^s], \ldots, [v_n^s]\}$ for $s \in S$, since the Kripke frame $M$ is proper, we never equate two facets.

Now let $f : M \to N$ be a morphism in $\mathcal{K}_A$. We define $G(f) : G(M) \to G(N)$ that maps a vertex $[v_i^s]$ of $G(M)$ to the vertex $[v_i^{f(s)}]$ of $G(N)$. This map is well-defined (i.e., the image of a vertex does not depend on the chosen representative) because $f$ is a morphism of Kripke frames, and thus it preserves the indistinguishability relations. It is easily checked that this is moreover a simplicial map.

Consider now a Kripke frame $M = \langle S, \sim \rangle$ in $\mathcal{K}_A$ with agent set $A$. $FG(M)$ is the Kripke frame $N = \langle T, \sim' \rangle$ such that $T$ is the set of facets of $G(M)$. But we have seen above that the facets of $G(M)$ are of the form $\{[v_0^s], \ldots, [v_n^s]\}$ (where $s \in S$), therefore, $T$ is in bijection with $S$. Finally, in $FG(M)$, $X \sim'_a Y$ if and only if $a \in \chi(X \cap Y)$, where $\chi$ is the coloring, in $G(M)$, of $X$ and $Y$ which are facets in $G(M)$. But facets in $G(M)$ are just in direct bijection with the worlds of $M$, i.e. $X = \{[v_0^s], \ldots, [v_n^s]\}$ and $Y = \{[v_0^t], \ldots, [v_n^t]\}$ where $s, t \in M$. Note that $\chi([v_i^s]) = a_i$ and $\chi([v_i^t]) = a_i$ so $a \in \chi(X \cap Y)$ means that $a = a_i$ for some $i$ and $v_i^s \ R \ v_i^t$. This can only be the case, by definition of $G(M)$ if $s \sim_{a_i} t$. This proves that $FG(M)$ and $M$ are isomorphic frames.

Conversely, let $C \in \mathcal{S}_A$ be a pure chromatic simplicial complex. Given a vertex $v \in \mathcal{V}(C)$ colored by $a_i$, we denote by $\widehat{v}$ the set of all the facets of $C$ which contain $v$. By definition, $\widehat{v}$ is an equivalence class of the relation $\sim_{a_i}$ in the Kripke frame $F(C)$. Therefore, it induces an $a_i$-colored vertex of $GF(C)$, which we write $f(v)$. It is easy to check that the map $f$ is a bijection between the vertices of $C$ and those of $GF(C)$. Moreover, the facets of $GF(C)$ are of the form $X_w = \{[v_0^w], \dots, [v_n^w]\}$, where $w$ is a world of $F(C)$, that is, a facet of $C$. So, we have a bijection between the facets of $C$ and those of $GF(C)$. This bijection is given precisely by $f$. Indeed, a facet $Y = \{v_0, \dots, v_n\}$ of $C$ is sent to $f(Y) = \{f(v_0), \dots, f(v_n)\}$. But since each vertex $v_i \in Y$, we have $\widehat{v_i} = [Y]_{a_i}$, and therefore $f(Y) = X_Y$. Hence $C$ and $GF(C)$ are isomorphic and therefore, $\mathcal{S}_A$ and $\mathcal{K}_A$ are equivalent categories. □

**Example 1.** The picture below shows a Kripke frame (left) and its associated chromatic simplicial complex (right). The three agents, named $b, g, w$, are represented as colors black, grey and white on the vertices of the simplicial complex. The three worlds of the Kripke frame correspond to the three triangles (i.e., 2-dimensional simplexes) of the simplicial complex. The two worlds indistinguishable by agent $b$, are glued along their black vertex; the two worlds indistinguishable by agents $g$ and $w$ are glued along the grey-and-white edge.

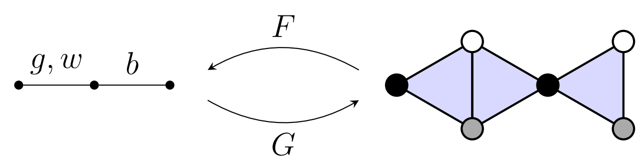


We now decorate our simplicial complexes with atomic propositions in order to get a notion of simplicial model.

*2.2. Simplicial models and Kripke models*

For technical reasons, we restrict to models where all the atomic propositions are saying something about some local value held by one particular agent. All the examples that we are interested in will fit in that framework. Let $\mathcal{V}$ be some countable set of values, and $AP = \{p_{a,x} \mid a \in A, x \in \mathcal{V}\}$ be the set of *atomic propositions*. Intuitively, $p_{a,x}$ is true if agent $a$ holds the value $x$. We write $AP_a$ for the atomic propositions concerning agent $a$.

A *simplicial model* $M = \langle C, \chi, \ell \rangle$ consists of a pure chromatic simplicial complex $\langle C, \chi \rangle$ of dimension $n$, and a labeling $\ell : \mathcal{V}(C) \to \mathscr{P}(AP)$ that associates with each vertex $v \in \mathcal{V}(C)$ a set of atomic propositions concerning agent $\chi(v)$, i.e., such that $\ell(v) \subseteq AP_{\chi(v)}$. Given a facet $X = \{v_0, \ldots, v_n\} \in C$, we write $\ell(X) = \bigcup_{i=0}^{n} \ell(v_i)$. A *morphism* of simplicial models $f : M \to M'$ is a chromatic simplicial map that preserves the labeling: $\ell'(f(v)) = \ell(v)$. We denote by $\mathcal{SM}_{A,AP}$ the category of simplicial models over the set of agents $A$ and atomic propositions $AP$.

A *Kripke model* $M = \langle S, \sim, L \rangle$ consists of a Kripke frame $\langle S, \sim \rangle$ and a function $L : S \to \mathscr{P}(AP)$. Intuitively, $L(s)$ is the set of atomic propositions that are true in the state $s$. A Kripke model is *proper* if the underlying Kripke frame is proper. A Kripke model is *local* if for every agent $a \in A$, $s \sim_a s'$ implies $L(s) \cap AP_a = L(s') \cap AP_a$, i.e., an agent always knows its own values. Let $M = \langle S, \sim, L \rangle$ and $M' = \langle S', \sim', L' \rangle$ be two Kripke models on the same set $AP$. A *morphism of Kripke models* $f : M \to M'$ is a morphism of the underlying Kripke frames such that $L'(f(s)) = L(s)$ for every state $s$ in $S$. We write $\mathcal{KM}_{A,AP}$ for the category of local proper Kripke models.

We can now extend the two maps $F$ and $G$ of Theorem 1 to an equivalence between simplicial models and Kripke models.

**Theorem 2.** *$\mathcal{SM}_{A,AP}$ and $\mathcal{KM}_{A,AP}$ are equivalent categories.*

*Proof.* We describe the functors $F : \mathcal{SM} \to \mathcal{KM}$ and $G : \mathcal{KM} \to \mathcal{SM}$. On the underlying Kripke frame and simplicial complex, they act the same as in the proof of Theorem 1.

Given a simplicial model $M = \langle C, \chi, \ell \rangle$, we associate the Kripke model $F(M) = \langle \mathcal{F}(C), \sim, L \rangle$ where the labeling $L$ of a facet $X \in \mathcal{F}(C)$ is given by $L(X) = \bigcup_{v \in X} \ell(v)$. This Kripke model is local since $X \sim_a Y$ means that $X$ and $Y$ share an $a$-colored vertex $v$, so $L(X) \cap AP_a = L(Y) \cap AP_a = \ell(v)$.

Conversely, given a Kripke model $M = \langle S, \sim, L \rangle$, the underlying simplicial complex of $G(M)$ is obtained by gluing together $n$-simplexes of the form $\{v_0^s, \ldots, v_n^s\}$, with $s \in S$. We label the vertex $v_i^s$ (colored by $a_i$) by $\ell(v_i^s) = L(s) \cap AP_{a_i}$. This is well defined because two vertices $v_i^s$ and $v_i^{s'}$ are identified whenever $s \sim_{a_i} s'$, so $L(s) \cap AP_{a_i} = L(s') \cap AP_{a_i}$ since $M$ is local.

The action of $F$ and $G$ on morphisms is the same as in Theorem 1. It is easy to check that the additional properties of morphisms between models are verified. Checking that $FG(M) \simeq M$ and $GF(M) \simeq M$ also works the same as in the previous theorem. $\square$

**Example 2.** The situation where each agent gets an input value from some base set of possible input values is often considered in distributed computing, e.g. [5]. Each agent knows its own input value, but doesn't know which value has been received by the other agents. Furthermore, any assignment of input values to the agents is possible. The *binary* case, where the set of possible input values is of size 2 is particularly interesting. The figure below shows the *binary input complex* and its associated Kripke model, for 2 and 3 agents. Notice that every possible combination of 0's and 1's is a possible world.

In the Kripke model, the agents are called $b, g, w$, and the labeling $L$ of the possible worlds is represented as a sequence of values, e.g., 101, representing the values chosen by the agents $b, g, w$ (in that order). In the 3-agents case, the labels of the dotted edges have been omitted to avoid overloading the picture, as well as other edges that can be deduced by transitivity.

In the simplicial model, agents are represented as colors (black, grey, and white). The labeling $\ell$ is represented as a single value in a vertex, e.g., "1" in a grey vertex means that agent $g$ has chosen value 1. The possible worlds correspond to edges in the 2-agents case, and triangles in the 3-agents case.

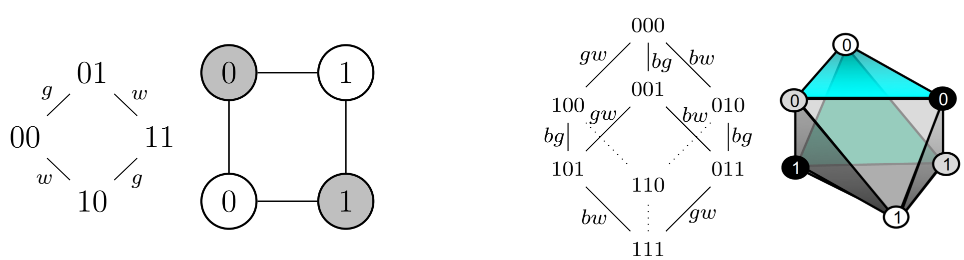


It is well known in the context of distributed computing [5] that the binary input simplicial complex for $n+1$ agents is homeomorphic to an $n$-dimensional sphere. Thus, it is sometimes called the *combinatorial $n$-sphere.*

**Example 3.** Consider the following situation. There are three agents black, grey and white, and a deck of four cards, $\{0, 1, 2, 3\}$. One card is given to each agent, and the last card is kept hidden. Each agent knows its own card, but not the other agents' cards. So, the atomic propositions indicate that "agent $a$ has card $x$". This situation also appears in distributed computing, related to *renaming* problems, e.g. [26], or Chapter 12 of [5].

The simplicial model corresponding to that situation is depicted below on the left. The color of vertices indicate the corresponding agent, and the labeling is its card. In the planar drawing, vertices that appear several times with the same color and value should be identified. The arrows $A$ and $B$ indicate how the edges should be glued together. What we obtain is a triangulated torus. If the deck of cards is $\{0, 1, 2\}$, we get the figure on the right, where the two white vertices (with card 0) should be identified, as well as the two black ones.

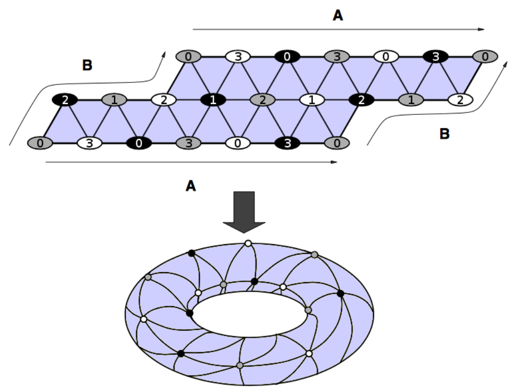


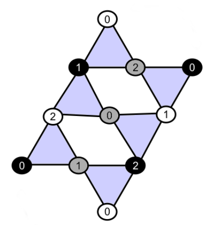

Thus, Theorem 2 says that simplicial models are closely related to Kripke models. Keeping that translation in mind, we can reformulate the usual semantics of formulas in Kripke models, in terms of simplicial models.

**Definition 1.** *We define the satisfaction relation $M, X \models \varphi$ determining when a formula $\varphi$ is* true *in some epistemic state $(M, X)$. Let $M = \langle C, \chi, \ell \rangle$ be a simplicial model, $X \in \mathcal{F}(C)$ a facet of $C$ and $\varphi \in \mathcal{L}_{CK}(A, AP)$.*

| | | |
|---|---|---|
| $M, X \models p$ | *iff* | $p \in \ell(X)$ |
| $M, X \models \neg\varphi$ | *iff* | $M, X \not\models \varphi$ |
| $M, X \models \varphi \wedge \psi$ | *iff* | $M, X \models \varphi$ *and* $M, X \models \psi$ |
| $M, X \models K_a\varphi$ | *iff* | *for all* $Y \in \mathcal{F}(C), a \in \chi(X \cap Y)$ *implies* $M, Y \models \varphi$ |
| $M, X \models C_B\varphi$ | *iff* | *for all* $Y \in \mathcal{F}(C), X \; R_B^* \; Y$ *implies* $M, Y \models \varphi$ |

*where $R_B^*$ denotes the transitive closure of the relation $R_B$ defined by $X \; R_B \; Y$ if there exists some agent $a \in B$ such that $a \in \chi(X \cap Y)$.*

We can show that this definition of truth agrees with the usual one (which we write $\models_{\mathcal{K}}$ to avoid confusion) on the corresponding Kripke model.

**Proposition 1.** *Given a simplicial model $M$ and a facet $X$, $M, X \models \varphi$ iff $F(M), X \models_{\mathcal{K}} \varphi$. Conversely, given a local proper Kripke model $N$ and state $s$, $N, s \models_{\mathcal{K}} \varphi$ iff $G(N), G(s) \models \varphi$, where $G(s)$ is the facet $\{v_0^s, \ldots, v_n^s\}$ of $G(N)$.*

*Proof.* This is straightforward by induction on the formula $\varphi$. $\square$

It is well-known that the axiom system $\mathbf{S5_n}$ is sound and complete with respect to the class of Kripke models [2]. Since we restrict here to local Kripke models, we need to add the following axiom (or axiom schema, if $\mathcal{V}$ is infinite), saying that every agent knows which values it holds:

$$\mathbf{Loc} \ = \bigwedge_{a \in A, x \in \mathcal{V}} K_a(p_{a,x}) \vee K_a(\neg p_{a,x})$$

Note that a similar locality condition has been introduced in [27], also axiomatized by a formula depending on atoms (and not just on the underlying frame structure). These formulas depending on atoms are rather rare in the literature, but can also be found in distributed computing related epistemic logics, where e.g. public announcements can be lost during transmission [28].

**Corollary 1.** *The axiom system* $\mathbf{S5_n} + \mathbf{Loc}$ *is sound and complete w.r.t. the class of simplicial models.*

*Proof.* We first adapt the proof of [2] for $\mathbf{S5_n}$, to show that $\mathbf{S5_n} + \mathbf{Loc}$ is sound and complete w.r.t. the class of local proper Kripke models. At the end, we translate this result into simplicial models.

*Soundness.* If $\varphi$ is provable in $\mathbf{S5_n} + \mathbf{Loc}$, then it is true in every state of every proper local Kripke model. Indeed, it is well known that all the deduction rules of $\mathbf{S5_n}$ are admissible in all Kripke models; we only have to check that *local* Kripke models satisfy the $\mathbf{Loc}$ axiom, which is straightforward.

*Completeness.* If $\varphi \in \mathcal{L}_K$ is true in every proper local Kripke model, then it is provable. The usual proof for $\mathbf{S5_n}$ [2] proceeds by contraposition: assuming $\varphi$ is not provable, we construct a model in which $\varphi$ is false. This model is called the *canonical model*.

A set $\Gamma$ of formulas is *consistent* if $\Gamma \not\vdash_{\mathbf{S5_n}+\mathbf{Loc}} \bot$, and *maximal* if there is no greater consistent set $\Gamma' \supsetneq \Gamma$. The canonical model $M^c$ is defined as $M^c = \langle S^c, \sim^c, L^c \rangle$, where:

- $S^c = \{\Gamma \mid \Gamma \text{ is a consistent and maximal set of formulas}\}$
- $\Gamma \sim_a^c \Delta$ iff $\{K_a\varphi \mid K_a\varphi \in \Gamma\} = \{K_a\varphi \mid K_a\varphi \in \Delta\}$
- $L^c(\Gamma) = \Gamma \cap AP$

The usual machinery on the canonical model shows that if $\varphi$ is not provable in $\mathbf{S5_n} + \mathbf{Loc}$, then there is $\Gamma$ such that $M^c, \Gamma \not\models \varphi$. We only have to check that $M^c$ is proper and local. First notice that if $\Gamma$ is consistent and maximal, then $p_{a,x} \in \Gamma \Leftrightarrow K_a(p_{a,x}) \in \Gamma$. The right-to-left direction follows from the truth axiom $K_a\varphi \Rightarrow \varphi$. For the converse, assume $p_{a,x} \in \Gamma$ and $K_a(p_{a,x}) \notin \Gamma$. Then we must have $K_a(\neg p_{a,x}) \in \Gamma$ because of the **Loc** axiom, which implies $\neg p_{a,x} \in \Gamma$, and thus $\Gamma$ would be inconsistent.

*$M^c$ is local:* assume $\Gamma \sim_a^c \Delta$, we show that $L^c(\Gamma) \cap AP_a = L^c(\Delta) \cap AP_a$. Indeed, $p_{a,x} \in \Gamma \Leftrightarrow K_a(p_{a,x}) \in \Gamma \Leftrightarrow K_a(p_{a,x}) \in \Delta \Leftrightarrow p_{a,x} \in \Delta$.

*$M^c$ is proper:* assume that for all $a$, $\Gamma \sim_a^c \Delta$. Then we show by a straightforward induction on $\varphi$ that $\varphi \in \Gamma \Leftrightarrow \varphi \in \Delta$, i.e., $\Gamma = \Delta$.

Using Proposition 1, we can now transpose this result into simplicial models. Suppose a formula $\varphi$ is true for every local proper Kripke model and any state. Then given a simplicial model and facet $(M, X)$, since by assumption $F(M), X \models_{\mathcal{K}} \varphi$, we also have $M, X \models \varphi$ by Proposition 1. So $\varphi$ is true in every simplicial model. Similarly, the converse also holds. $\square$

Remark that Corollary 1 works only for epistemic logic formulas in $\mathcal{L}_K$, without common knowledge operators. When we include common knowledge, it seems that the axiom system $\mathbf{S5C_n} + \mathbf{Loc}$ should be sound and complete, but the usual proof for Kripke models [2] is not as easy to adapt. So this remains an open question for now. The next theorem states that morphisms of simplicial models cannot "gain knowledge about the world". This will be useful in Section 4 when we formulate the solvability of a task as the existence of some morphism.

**Theorem 3** (knowledge gain). *Consider simplicial models $M = \langle C, \chi, \ell \rangle$ and $M' = \langle C', \chi', \ell' \rangle$, and a morphism $f : M \to M'$. Let $X \in \mathcal{F}(C)$ be a facet of $M$, $a$ an agent, and $\varphi \in \mathcal{L}_{CK}$ a positive formula, i.e. which does not contain negations except, possibly, in front of atomic propositions. Then, $M', f(X) \models \varphi$ implies $M, X \models \varphi$.*

*Proof.* We proceed by induction on $\varphi$. First, for $p$ an atomic proposition, since morphisms preserve valuations, we have $M', f(Y) \models p$ iff $M, Y \models p$. Thus the theorem is true for (possibly negated) atomic propositions. The case of the conjunction follows trivially from the induction hypothesis.

Suppose now that $M', f(X) \models K_a\varphi$. In order to show $M, X \models K_a\varphi$, assume that $a \in \chi(X \cap Y)$ for some facet $Y$, and let us prove $M, Y \models \varphi$.

Let $v$ be the $a$-colored vertex in $X \cap Y$. Then $f(v) \in f(X) \cap f(Y)$ and $\chi(f(v)) = a$. So $a \in \chi(f(X) \cap f(Y))$ and thus $M', f(Y) \models \varphi$. By induction hypothesis, we obtain $M, Y \models \varphi$. Finally, suppose that $M', f(X) \models C_B\varphi$. We want to show that $M, X \models C_B\varphi$, i.e., for every $Y$ reachable from $X$ following a sequence of simplexes sharing a $B$-colored vertex, $M, Y \models \varphi$. By the same reasoning as in the $K_a$ case, $f(Y)$ is $B$-reachable from $f(X)$, so $M', f(Y) \models \varphi$, and thus $M, Y \models \varphi$. $\square$

The restriction on $\varphi$ being positive forbids formulas saying something about what an agent does not know. Indeed, one can "gain" the knowledge that some agent does not know something; but this is not relevant information for solving the tasks that we have in mind.

*2.3. The case of non-local Kripke models*

Let us assume we are given a Kripke model, and we want to understand it geometrically. Theorem 2 gives us a way to translate a Kripke model into a simplicial model, with two restrictions: the original model must be *proper* and *local*. Being proper is not a very stong requirement, since most practical situations usually produce proper Kripke frames. However, the locality condition might seem quite restrictive. Indeed, we can think of many situations where some properties of the worlds do not take the form of a value held by one of the agents.

For instance, in Example 3, we could have wanted an atomic proposition representing the value of the remaining card that was not dealt to any of the agents. We cannot assign the hidden card to one of the agents, because none of them knows what its value is. Thus, this is a non-local situation. However, in this particular example, we did not need an atomic proposition for the hidden card, because this information would be redundant: we can recover it from the values of the cards that were dealt to the agents.

In this section, we will show that it is always possible to do this trick, as long as the underlying Kripke frame is *finite*. Starting with a finite (proper) non-local Kripke model $M$, we will construct a local Kripke model $M^{\text{loc}}$ (with a new set of atomic propositions), such that every formula $\varphi$ on $M$ can be translated to a formula $\varphi^{\text{loc}}$ on $M^{\text{loc}}$ which is equivalent.

Thus, in this section only, $AP$ denotes an arbitrary set of atomic propositions. Let $M = \langle S, (\sim_a)_{a \in A}, L \rangle$ be a proper Kripke model on $AP$, where the set $S$ is finite. For each agent $a \in A$, we write $X_a = S/\sim_a$ for the set of equivalence classes of $\sim_a$. For $s \in S$, we write $[s]_a \in X_a$ the equivalence class of $s$. Then we take one atomic proposition symbol for each agent and equivalence class, which gives us a new set of atomic propositions $AP^{\mathrm{loc}}$:

$$AP^{\mathrm{loc}} \;=\; \{E_{a,x} \mid a \in A \text{ and } x \in X_a\}$$

We now define a Kripke model $M^{\mathrm{loc}} = \langle S, (\sim_a)_{a\in A}, L^{\mathrm{loc}}\rangle$ on $AP^{\mathrm{loc}}$, with the same underlying Kripke frame $\langle S, \sim\rangle$ as $M$, and the labeling $L^{\mathrm{loc}} : S \to \mathscr{P}(AP^{\mathrm{loc}})$ is defined as $L^{\mathrm{loc}}(s) = \{E_{a,[s]_a} \mid a \in A\}$.

**Proposition 2.** *The Kripke model $M^{loc}$ is proper and local.*

*Proof.* It is proper since it has the same underlying Kripke frame as $M$, which is proper. To check locality, suppose $s \sim_a s'$. Then $L^{\mathrm{loc}}(s) \cap AP^{\mathrm{loc}}_a = \{E_{a,[s]_a}\} = L^{\mathrm{loc}}(s') \cap AP^{\mathrm{loc}}_a$, since $[s]_a = [s']_a$. $\square$

We now want to translate every formula $\varphi \in \mathcal{L}_{CK}(AP)$ to a formula $\varphi^{\mathrm{loc}} \in \mathcal{L}_{CK}(AP^{\mathrm{loc}})$, such that for every world $s \in S$, $M, s \models \varphi$ iff $M^{\mathrm{loc}}, s \models \varphi^{\mathrm{loc}}$. First, remark that, since $M$ is proper, every world $s \in S$ is uniquely determined by the set of equivalence relations $\{[s]_a \mid a \in A\}$. Thus, we can translate every atomic proposition $p \in AP$ as follows. Let $V_p = \{s \in S \mid p \in L(s)\}$ be the set of worlds where $p$ is true in $M$. We then define $p^{\mathrm{loc}} = \bigvee_{s\in V_p} \bigwedge_{a\in A} E_{a,[s]_a}$. This is well-defined because $V_p \subseteq S$ is finite. The above remark says that, in $M^{\mathrm{loc}}$, the formula $\bigwedge_{a\in A} E_{a,[s]_a}$ is true exactly in the world $s$. So, the formula $p^{\mathrm{loc}}$ is true exactly in the worlds of $V_p$, i.e., $M, s \models p$ iff $M^{\mathrm{loc}}, s \models p^{\mathrm{loc}}$. It is now straightforward to extend this translation to every formula $\varphi \in \mathcal{L}_{CK}(AP)$: we just replace every occurrence of an atomic proposition $p$ by $p^{\mathrm{loc}}$. This gives us a formula $\varphi^{\mathrm{loc}} \in \mathcal{L}_{CK}(AP^{\mathrm{loc}})$, and by a trivial induction we get:

**Theorem 4.** *For every $s \in S$, $M, s \models \varphi$ iff $M^{loc}, s \models \varphi^{loc}$.*

Therefore, through this translation, even the non-local Kripke models can be interpreted geometrically. The only really relevant condition is that the underlying Kripke frame must be proper. Notice however that the model $M^{\mathrm{loc}}$ might contain more information than $M$: given a formula $\psi \in \mathcal{L}_{CK}(AP^{\mathrm{loc}})$, there might be no equivalent formula $\varphi \in \mathcal{L}_{CK}(AP)$.

## 3. DEL via simplicial complexes

We describe here our adaptation of Dynamic Epistemic Logic (DEL) to simplicial models, and then we define one particular action model which is fundamental in distributed computing. We start by recalling some DEL notions in Section 3.1. Then in Section 3.2, we describe the distributed computing action model example that will be further studied in the next section. Some basic topological properties about this action model are included in Section 3.3.

### *3.1. DEL basic notions*

DEL is the study of modal logics of model change [1, 2]. A modal logic studied in DEL is obtained by using action models [3], which are relational structures that can be used to describe a variety of informational actions. An action can be thought of as an announcement made by the environment, which is not necessarily public, in the sense that not all agents receive these announcements. An action model describes all the possible actions that might happen, as well as how they affect the different agents. The key notion of DEL that we will use is the so-called product-update operation: given an epistemic model $M$ and an action model $\mathcal{A}$, the product update $M[\mathcal{A}]$ is a new model which describes all the new possible worlds after an action from $\mathcal{A}$ has occurred in $M$. In this section, we will define three flavors of this product update operation. First, we recall the usual notion of action model and product update, where $M$ is a Kripke model. Then, we define a hybrid version of the product update model, where $M$ is now a simplicial model, but $\mathcal{A}$ is the usual notion of action model. Finally, we describe a fully simplicial version, where both $M$ and $\mathcal{A}$ are given as simplicial complexes. We then prove formally that these three notions of product update agree.

*Dynamic Epistemic Logic.* An *action model* is a structure $\langle T, \sim, \mathsf{pre}\rangle$, where $T$ is a domain of *action points*, such that for each $a \in A$, $\sim_a$ is an equivalence relation on $T$, and $\mathsf{pre} : T \to \mathcal{L}_{\mathcal{K}}$ is a function that assigns a *precondition* formula $\mathsf{pre}(t)$ to each $t \in T$. For an initial Kripke model $M$, the effect of action model $\mathcal{A}$ is a Kripke model $M[\mathcal{A}]$. Let $M = \langle S, \sim, L\rangle$ be a Kripke model and $\mathcal{A} = \langle T, \sim, \mathsf{pre}\rangle$ be an action model. The *product update model* is $M[\mathcal{A}] = \langle S[\mathcal{A}], \sim^{[\mathcal{A}]}, L[\mathcal{A}]\rangle$, where each world of $S[\mathcal{A}]$ is a pair $(s,t)$ with $s \in S$, $t \in T$ such that $\mathsf{pre}(t)$ holds in $s$. Then, $(s,t) \sim_a^{[\mathcal{A}]} (s',t')$ whenever it holds that $s \sim_a s'$ and $t \sim_a t'$. The valuation $L[\mathcal{A}]$ at a pair $(s,t)$ is just as it was at $s$, i.e., $L[\mathcal{A}]((s,t)) = L(s)$.

**Proposition 3.** *Let $M$ be a local proper Kripke model and $\mathcal{A} = \langle T, \sim, \mathsf{pre}\rangle$ a proper action model, then $M[\mathcal{A}]$ is proper and local.*

*Proof.* $M[\mathcal{A}]$ is proper: let $(s,t)$ and $(s',t')$ be two distinct states of $M[\mathcal{A}]$. Then either $s \neq s'$ or $t \neq t'$, and in both cases, since $M$ and $\mathcal{A}$ are proper, at least one agent can distinguish between the two. Now, $M[\mathcal{A}]$ is local: suppose $(s,t) \sim_a^{[\mathcal{A}]} (s',t')$. Then in particular $s \sim_a s'$ and since $M$ is local, $L(s) \cap AP_a = L(s') \cap AP_a$. The same goes for $L[\mathcal{A}]$ since it just copies $L$. $\square$

*A simplicial complex version of the product update.* We now define a variant of DEL, where we start with a simplicial model $M$ instead of a Kripke model; and the product update model $M[\mathcal{A}]$ that we obtain at the end should also be a simplicial model. To understand the following definitions, the reader should keep in mind the translation between Kripke models and simplicial ones (see Theorem 2).

Let $M = \langle C, \chi, \ell \rangle$ be a simplicial model, and $\mathcal{A} = \langle T, \sim, \mathsf{pre} \rangle$ be a proper action model. The *product update simplicial model* $M[\mathcal{A}] = \langle C[\mathcal{A}], \chi[\mathcal{A}], \ell[\mathcal{A}] \rangle$ is defined as follows. Intuitively, the facets of $C[\mathcal{A}]$ should correspond to pairs $(X,t)$ where $X \in \mathcal{F}C$ is a facet of $M$ and $t \in T$ is an action of $\mathcal{A}$, such that $M, X \models \mathsf{pre}(t)$. Moreover, two such facets $(X,t)$ and $(Y,t')$ should be glued along their $a$-colored vertex whenever $a \in \chi(X \cap Y)$ and $t \sim_a t'$. Formally, the vertices of $C[\mathcal{A}]$ are pairs $(v,E)$ where $v \in \mathcal{V}(C)$ is a vertex of $M$; $E$ is an equivalence class of $\sim_{\chi(v)}$; and $v$ belongs to some facet $X \in \mathcal{F}(C)$ such that there exists $t \in E$ such that $M, X \models \mathsf{pre}(t)$. Such a vertex keeps the color and labeling of its first component: $\chi[\mathcal{A}](v,E) = \chi(v)$ and $\ell[\mathcal{A}](v,E) = \ell(v)$.

**Example 4.** Below is depicted an example with three agents, where $M$ consists of two worlds $X$ and $Y$, and $\mathcal{A}$ has two actions $t_1$ and $t_2$. The grey agent is the only one which cannot distinguish between the two actions. The precondition $\mathsf{pre}(t_1)$ is true in both $X$ and $Y$, but $\mathsf{pre}(t_2)$ is true only in $Y$.

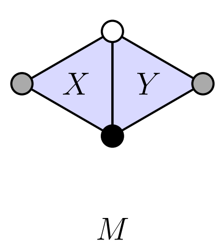


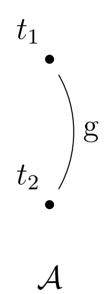


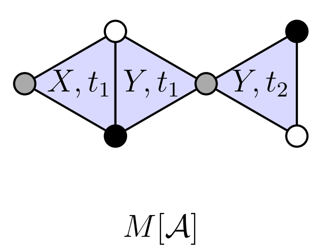


The advantage of this "hybrid" product update is that it mimics closely the classic definition; thus, it should be easier to understand for readers already familiar with DEL. In the following, we define yet another product update operation, where both $M$ and $\mathcal{A}$ are simplicial. This will help us uncover the geometric structure of this operation.

*A fully simplicial version of the product update.* We can extend the translation of DEL into simplicial complexes by noticing that the action model $\mathcal{A} = \langle T, \sim, \mathsf{pre} \rangle$ merely consists of a Kripke frame $\langle T, \sim \rangle$ along with precondition formulas. By applying Theorem 1 to the underlying Kripke frame, we obtain a simplicial counterpart of action models. A *simplicial action model* $\langle T, \chi, \mathsf{pre} \rangle$ consists of a pure chromatic simplicial complex $\langle T, \chi \rangle$, where the facets $\mathcal{F}(T)$ represent communicative *actions*, and $\mathsf{pre}$ assigns to each facet $X \in \mathcal{F}(T)$ a precondition formula $\mathsf{pre}(X)$ in $\mathcal{L}_{\mathcal{K}}$.

Before defining the corresponding product update operation, we first need to define the Cartesian products in the category of chromatic simplicial complexes. Given two pure chromatic simplicial complexes $C$ and $T$ of dimension $n$, the Cartesian product $C \times T$ is the following pure chromatic simplicial complex of dimension $n$. Its vertices are of the form $(u, v)$ with $u \in \mathcal{V}(C)$ and $v \in \mathcal{V}(T)$ such that $\chi(u) = \chi(v)$; the color of $(u, v)$ is $\chi((u, v)) = \chi(u) = \chi(v)$. Its simplexes are of the form $X \times Y = \{(u_0, v_0), \ldots, (u_k, v_k)\}$ where $X = \{u_0, \ldots, u_k\} \in C$, $Y = \{v_0, \ldots, v_k\} \in T$ and $\chi(u_i) = \chi(v_i)$.

Let $M = \langle C, \chi, \ell \rangle$ be a simplicial model, and $\mathcal{A} = \langle T, \chi, \mathsf{pre} \rangle$ be a simplicial action model. The *product update simplicial model* $M[\mathcal{A}] = \langle C[\mathcal{A}], \chi[\mathcal{A}], \ell[\mathcal{A}] \rangle$ is a simplicial model whose underlying simplicial complex is a subcomplex of the Cartesian product $C \times T$, induced by all the facets of the form $X \times Y$ such that $\mathsf{pre}(Y)$ holds in $X$, i.e., $M, X \models \mathsf{pre}(Y)$. The valuation $\ell : \mathcal{V}(C[\mathcal{A}]) \to \mathscr{P}(AP)$ at a pair $(u, v)$ is just as it was at $u$: $\ell[\mathcal{A}]((u, v)) = \ell(u)$.

Recall from Theorem 2 the two functors $F$ and $G$ that define an equivalence of categories between simplicial models and Kripke models. We have a similar correspondence between action models and simplicial action models, which we still write $F$ and $G$. On the underlying Kripke frame and simplicial complex they are the same as before; and the precondition of an action point is just copied to the corresponding facet. Proposition 4 says that the "hybrid" product update and the "fully simplicial" one give the same result.

**Proposition 4.** *Let $M$ be a simplicial model, and $\mathcal{A}$ an action model. $G(\mathcal{A})$ is the associated simplicial action model. Then the simplicial models $M[\mathcal{A}]$ and $M[G(\mathcal{A})]$ are isomorphic.*

*Proof.* By definition of the hybrid product update $M[\mathcal{A}]$, its vertices are of the form $(u, E)$ where $u$ is a vertex of $M$ and $E$ an equivalence class of $\sim_{\chi(u)}$. On the other hand, the vertices of the product update $M[G(\mathcal{A})]$ are of the form $(u, v)$ where $u$ is a vertex of $M$ and $v$ a vertex of $G(\mathcal{A})$. But by definition of $G$ (see the proof of Theorem 1), $v$ corresponds precisely to an equivalence class of $\sim_{\chi(u)}$. This gives us a bijection between the vertices of $M[\mathcal{A}]$ and those of $M[G(\mathcal{A})]$. There remains to check that both directions of the bijection are actually simplicial maps, which is straightforward. $\square$

The following proposition says that the "classic" product update agrees with the "fully simplicial" one.

**Proposition 5.** *Consider a simplicial model $M$ and simplicial action model $\mathcal{A}$, and their corresponding Kripke model $F(M)$ and action model $F(\mathcal{A})$. Then, the Kripke models $F(M[\mathcal{A}])$ and $F(M)[F(\mathcal{A})]$ are isomorphic. The same is true for $G$, starting with a Kripke model $M$ and action model $\mathcal{A}$.*

*Proof.* The main observation is that both constructions of product update model rely on a notion of Cartesian product (in the category of pure chromatic simplicial complexes for $M[\mathcal{A}]$, and in the category of Kripke frames for $F(M)[F(\mathcal{A})]$). These are both Cartesian products in the categorical sense, therefore they are preserved by the functor $F$ because it is part of an equivalence of category: given $C$ and $T$ the underlying chromatic simplicial complexes of $M$ and $\mathcal{A}$ respectively, we have $F(C \times T) \simeq F(C) \times F(T)$. Intuitively, a state of $F(C \times T)$ is a facet of $C \times T$ of the form $X \times Y$, which is entirely determined by the two facets $X \in F(C)$ and $Y \in F(T)$.

Then, $M[\mathcal{A}]$ is defined as the sub-complex of $C \times T$ consisting of all facets $X \times Y$ such that $\mathsf{pre}(Y)$ holds in $X$, that is, $M, X \models \mathsf{pre}(Y)$. On the other hand, $F(M)[F(\mathcal{A})]$ is defined as the sub-frame of $F(C) \times F(T)$ consisting of all worlds $(X, Y)$ such that $\mathsf{pre}(Y)$ holds in $X$, that is, $F(M), X \models_{\mathcal{K}} \mathsf{pre}(Y)$. By Proposition 1, these two conditions are equivalent, so the underlying Kripke frames of $F(M[\mathcal{A}])$ and $F(M)[F(\mathcal{A})]$ are isomorphic. Moreover, in both cases, the labeling $L$ of atomic propositions is just copied from the first component, so they are also isomorphic as Kripke models. $\square$

### *3.2. A basic action model for distributed computing*

We describe here the *immediate snapshot* action model $\mathcal{IS}$ for one communication exchange among asynchronous agents. As an action model, it is new and to the best of our knowledge it has not been studied from the DEL perspective. Immediate snapshot operations are important in distributed computing, and many variants of computational models based on them have been considered, including multi-round communication exchanges, see e.g. [5, 7]. For the point we want to make about using DEL, the main issues can be studied with this very simple action model, even in the one communication exchange case.

The situation we have in mind is the following. The $n+1$ agents correspond to $n+1$ concurrent processes. Initially, each process has some input value, and they communicate (only once) through a shared memory array in order to try to learn each other's input values. They use the following protocol: each process has a dedicated memory cell in the array, to which it writes its input value. Then, it reads one by one all the cells of the array, to see which other input values have been written. The processes are asynchronous, meaning that an execution consists of an arbitrary interleaving of the write and read operations of all the processes (one write per process, and $n+1$ reads per process).

We could describe the action model corresponding to this situation, and present all of our results using it. But, to illustrate more easily the basic ideas, we will actually restrict ourselves to a subset of all the executions described above, which will give rise to simplicial complexes with nicer topological properties. And we do so without loss of generality, because from the task computability perspective, they are known to be equivalent [7]. The resulting action model is what we call the immediate snapshot action model, $\mathcal{IS}$.

The interleavings that we consider can be represented by a sequence of concurrency classes, $c_1, c_2, \ldots, c_m$, where each $c_i$ is a set of agents. For each concurrency class $c_i$, all the agents in $c_i$ execute their write operations simultaneously, then all of them execute their read operations simultaneously, then we move on to the next concurrency class $c_{i+1}$. Thus, all the agents in $c_i$ see each other's values, as well as the values of the agents from the previous concurrency classes, but they do not see the values written by the agents in the subsequent concurrency classes. The set of values seen by an agent in an execution will be called its *view*. It is the only information that the agent gathers from the execution: thus, two executions will be indistinguishable by some agent if and only if it has the same view in both executions.

Let us define formally the action model corresponding to such immediate snapshot schedules. A *sequential partition* of agents $A$ is a sequence $c = c_1, c_2, \ldots, c_m$, of non-empty, disjoint subsets of $A$, whose union is equal to $A$. Each $c_i$ is called a *concurrency class*. Notice that $1 \leq m \leq |A|$, and when

$m = 1$ all agents take an immediate snapshot concurrently, while if $m = |A|$, all take immediate snapshots sequentially. The sequential partition $c$ represents one possible execution of the immediate snapshot operation.

The agents in a concurrency class $c_j$ learn the input values of all the agents in earlier concurrency classes $c_i$ for $i \leq j$, and which agent wrote which value. In particular, agents in $c_m$ learn the inputs of all agents (and there is always at least one such agent), and if $m = 1$, then all agents learn all the values.

**Example 5.** As a concrete example, suppose we have three agents $A = \{a, b, c\}$, and suppose that their input values are, respectively, 1, 2 and 3. We consider the execution $x = \{b\}, \{a, c\}$ with two concurrency classes. Initially, the shared array is empty, which we represent as $\begin{array}{|c|c|c|}\hline \bot & \bot & \bot \\ \hline\end{array}$, where the three cells correspond to $a$, $b$, $c$, in that order. First, process $b$, which is alone in the first concurrency class, writes its value and immediately reads. Thus, $b$ sees only its own value, and its view is $\begin{array}{|c|c|c|}\hline \bot & 2 & \bot \\ \hline\end{array}$. Then, in the second concurrency class, both $a$ and $c$ write their values simultaneously, and after that, they both read simultaneously. So, both $a$ and $c$ have the same view, which is the following array: $\begin{array}{|c|c|c|}\hline 1 & 2 & 3 \\ \hline\end{array}$. Another possible execution is $y = \{c\}, \{a\}, \{b\}$, where every concurrency class is a singleton. Here, the view of $c$ is $\begin{array}{|c|c|c|}\hline \bot & \bot & 3 \\ \hline\end{array}$, the view of $a$ is $\begin{array}{|c|c|c|}\hline 1 & \bot & 3 \\ \hline\end{array}$ and the view of $b$ is $\begin{array}{|c|c|c|}\hline 1 & 2 & 3 \\ \hline\end{array}$. Finally, if we consider the execution $z = \{a, b, c\}$, all three processes run concurrently, and all three agents will have the same view, $\begin{array}{|c|c|c|}\hline 1 & 2 & 3 \\ \hline\end{array}$.

Notice that, to define the view, we need to know not only the execution $c$, but also the input values of the agents. Thus, the action model $\mathcal{IS}$ that we are defining will actually be parameterized by an *input model*, which represents all the possible input values that we want to take into consideration. An *action* of $\mathcal{IS}$ will consist of an execution $c$, along with an assignment of an input value to each agent. This is in accordance with what is usually done in DEL: say we want to define an action model where, for example, an agent can reveal the value of its card to the other agents. Then we will not have only one action "reveal card", we will actually have many actions "reveal that the card is $x$", for each possible value of $x$. Then the preconditions will make sure that such an action happens in a world where the value of the card is actually $x$.

To make things more simple, let us fix one particular input model: the simplicial model of Example 2 where three agents $A = \{b, g, w\}$ each have a binary input value 0 or 1. Let $M = \langle C, \chi, \ell \rangle$ be the corresponding simplicial model, and denote a facet $X \in \mathcal{F}(C)$ by a binary sequence $b_0 b_1 b_2$, corresponding to the three values of $b, g, w$, in that order.

We define the *immediate snapshot* action model $\mathcal{IS} = \langle T, \sim, \mathsf{pre} \rangle$ as follows. An action $t \in T$ is given by the data $c, b_0, b_1, b_2$, where $c$ is a sequential partition of $A = \{b, g, w\}$, and $b_0$, $b_1$, $b_2$ are binary values 0 or 1. Such an action will be written $c^{b_0 b_1 b_2}$. The precondition $\mathsf{pre}(c^{b_0 b_1 b_2})$ of this action is a formula expressing the fact that the inputs of the agents $b, g, w$ are respectively $b_0, b_1, b_2$. Therefore, $\mathsf{pre}(c^{b_0 b_1 b_2})$ is true precisely in the facet $b_0 b_1 b_2$ of $M$. Formally, if $p_{a,x}$ is the atomic proposition expressing that agent $a$ has input value $x$, $\mathsf{pre}(c^{b_0 b_1 b_2}) = p_{b,b_0} \wedge p_{g,b_1} \wedge p_{w,b_2}$ The indistinguishability relation is defined as $t \sim_a t'$ iff $\mathsf{view}_a(t) = \mathsf{view}_a(t')$, where $\mathsf{view}_a(t)$ is defined as expected: if $c = c_1, \ldots, c_m$ is a sequential partition of $A$, and the agent $a$ is in $c_j$, then $\mathsf{view}_a(c^{b_0 b_1 b_2})$ is the vector obtained from $b_0 b_1 b_2$ by replacing the value $b_i$ by $\perp$ whenever the corresponding agent is not in $\bigcup_{i \leq j} c_i$.

We can also describe the *simplicial action model* $G(\mathcal{IS})$, which is the simplicial counterpart of $\mathcal{IS}$ given by Theorem 1. It contains exactly the same data as $\mathcal{IS}$, but it is translated in the language of chromatic simplicial complexes, which allows us to visualise it better. Formally, we have $G(\mathcal{IS}) = \langle \hat{T}, \chi, \mathsf{pre} \rangle$ where $\langle \hat{T}, \chi \rangle$ is a chromatic simplicial complex whose vertices are $\mathcal{V}(\hat{T}) = \{\langle a, \mathsf{view}_a(c^{b_0 b_1 b_2}) \rangle \mid a \in A, c \text{ is an execution, and } b_0, b_1, b_2 \in \{0,1\}\}$; and whose facets are of the form:

$$X = \{\langle b, \mathsf{view}_b(c^{b_0 b_1 b_2}) \rangle, \langle g, \mathsf{view}_g(c^{b_0 b_1 b_2}) \rangle, \langle w, \mathsf{view}_w(c^{b_0 b_1 b_2}) \rangle\}$$

The precondition of such a facet is $\mathsf{pre}(X) = p_{b,b_0} \wedge p_{g,b_1} \wedge p_{w,b_2}$.

The picture below illustrates (part of) the simplicial action model $G(\mathcal{IS})$. The two triangles on the left represent two facets of the input model $M$, with input values 000 (green) and 100 (yellow). On the right are the corresponding facets of $G(\mathcal{IS})$. Each of the two input triangles have been subdivided into 13 smaller triangles: one for each possible sequential partition of $A = \{b, g, w\}$. Four of these sequential partitions are depicted in the bubbles $X, Y, Z, W$. The tables in the bubbles show the scheduling of the execution from top to bottom: for example, in execution $Z$, process $b$ goes first and sees only itself; then process $g$ goes second and sees both $b$ and $g$; then process $w$ goes last and sees everyone. The colors black, grey, white of the vertices correspond respectively to agents $b$, $g$, $w$. The view of each vertex is written next to it; when two (or three) neighboring vertices have the same view, it is written only once, on the edge (or triangle) between the two (or three) vertices. The precondition of all the green facets on the right is true exactly in the green facet of the input model $M$ on the left, and similarly for the yellow facets.

Notice that, for example, neither $b$ nor $w$ distinguish between actions $Y$ and $Z$, because their views are equal in $Y$ and $Z$: the view of $b$ consists of itself with value 0 and the view of $w$ consists of the three agents with value 0. Therefore, the two triangles corresponding to actions $Y$ and $Z$ are glued along their white and black vertices. Finally, let us look at what happens on the boundary shared by both subdivisions. For example, the two facets in the middle of the figure correspond to the sequential partition $\{gw\}\{b\}$; neither $w$ nor $g$ have seen $b$, so they cannot tell whether the input of $b$ is 0 or 1.

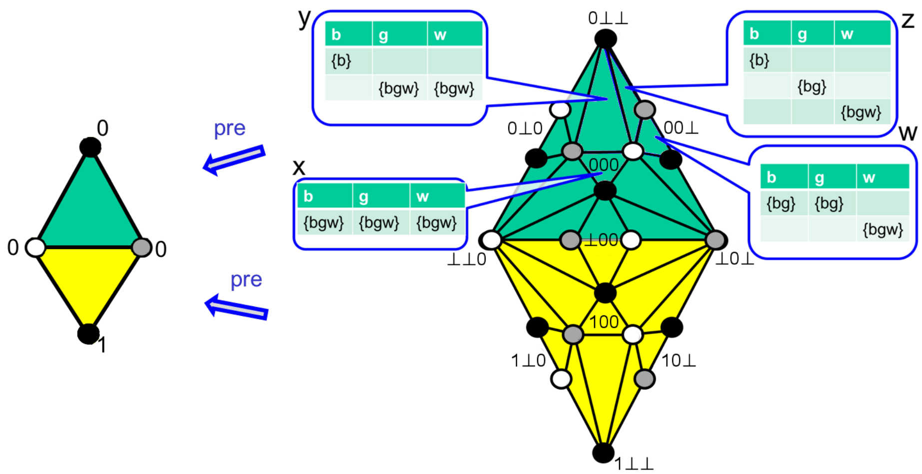


One last thing to notice about this picture is that, when we compute the product update model $M[\mathcal{IS}]$, we obtain a simplicial model whose underlying simplicial complex is the same as the one of $G(\mathcal{IS})$, depicted on the right. So, starting from the input model $M$, the effect of applying $\mathcal{IS}$ is to subdivide each facet of the input. The same thing happens for any input model $M$. Remarkably, the topology of the input simplicial complex is preserved: if $M$ is a sphere as in Example 2, then $M[\mathcal{IS}]$ is still a sphere.

In the rest of the paper, since by Proposition 4 $M[G(\mathcal{IS})]$ and $M[\mathcal{IS}]$ are isomorphic, we will drop the distinction between regular and simplicial action models, and just write $\mathcal{IS}$ even for the simplicial version.

*Multi-round communication.* In the $\mathcal{IS}$ model, each agent executes a single immediate snapshot. Iterating this model gives rise to the *iterated immediate snapshot model* $\mathcal{IS}^r$ [5, 29], where each agent executes $r$ consecutive immediate snapshots, on $r$ consecutive shared memory arrays. Starting from an input model $M$, the effect of applying the iterated immediate snapshot protocol is to subdivide each facet of the input complex $r$ times. So, once again, the topology of $M$ is preserved in $M[\mathcal{IS}^r]$. In the non-iterated version, the $r$ immediate snapshots are executed on the same memory. A subdivision is still obtained, but it is more complex [7]. If all schedules are considered, not only immediate snapshot schedules, then the topology is still preserved, even-though the resulting complex is no longer a subdivision, see e.g. [30] for more precise meaning about these claims and further discussion.

### *3.3. Some topological properties of $\mathcal{IS}$*

In this section, we recall some topological properties of the immediate-snapshot simplicial complex $\mathcal{IS}$. Even though we reformulate them using the "action model" vocabulary introduced in this paper, all the notions presented here can be found in [5]. Our main goal is to introduce Sperner's Lemma, which will be needed in out last application of DEL to distributed computing, in Section 4.4. In order to be able to apply Sperner's Lemma, we need $\mathcal{IS}$ to be a *pseudomanifold*, so we also define this notion and prove this fact.

A simplicial action model is *uniform* if its set of actions (facets) can be partitioned into $k$ copies of a complex $C$, called *components*, such that all actions in $C_i$ have the same precondition, which is true in exactly one facet $X_i$ of the simplicial model $M$. The action model $\mathcal{IS}$ is indeed uniform, and its components are isomorphic to a simplicial complex $C$, called the *standard chromatic subdivision*. Two components (in green and yellow) are depicted in the figure on page 24. It is clear from the picture that $C$ is a subdivision, but for an arbitrary number of agents, the proof is not simple [5, 31]. It has been shown to have several other topological properties, such as being collapsible [30]. But in fact, for many applications such as consensus and set agreement, it is sufficient to observe that $\mathcal{IS}$ is a *pseudomanifold*. For a more thorough and more general account of this topic, see chapter 9 of [5].

Let $C$ be a pure simplicial complex of dimension $n$. The complex $C$ is said to be *strongly connected* if for every pair $(X, Y)$ of $n$-simplices of $C$, there is a sequence $X_0, X_1, \ldots, X_k$ of $n$-simplices in $C$ such that $X_0 = X$, $X_k = Y$ and for every $i$, $X_i$ and $X_{i+1}$ share a common $(n-1)$-dimensional face. A strongly connected pure simplicial complex $C$ of dimension $n$ is a *pseudomanifold* if each $(n-1)$-simplex of $C$ is a face of either one or two $n$-simplices of $C$. The $(n-1)$-simplices of $C$ which are faces of precisely one $n$-simplex form a (possibly empty) subcomplex of $C$, called its *boundary*.

**Lemma 1.** *Each component of $\mathcal{IS}$ is a pseudomanifold. If $M$ is a pseudomanifold, then so is $M[\mathcal{IS}]$.*

*Proof.* For completeness purposes, we give a sketch of the proof that can be found in [5]. First, we prove the following simple observation, see [7]. Consider a sequential partition $c = c_1, c_2, \ldots, c_m$, and one of its concurrency classes, $c_i$, with $|c_i| > 1$. If we replace $c_i$ by splitting it in two non-empty concurrency classes $c'_i, c''_i$, in this order, we obtain another sequential partition, $c'$. Then only the agents in $c'_i$ have a different view. This result follows because for every $a \notin c'_i$, $\mathsf{view}_a(c) = \mathsf{view}_a(c')$. Only the agents $a \in c'_i$ have a different view, because they do not get information from the agents in $c''_i$.

Now we can prove Lemma 1. The components of $\mathcal{IS}$ are isomorphic to the complex with one facet for each sequential partition, and whose vertices are glued when they have the same view. Notice from the first observation we made that, if we take a singleton for $c'_i$, only one process's view has changed: two $n$-simplices share a common $(n-1)$-face.

*Strong connectivity:* let $c$ and $c'$ be two sequential partitions, corresponding to two $n$-simplices in a component of $\mathcal{IS}$. Then by splitting the concurrency classes of $c$ one by one, we can can go from $c$ to a sequential partition with only singletons. From there, we can regroup them again to obtain a sequential partition where all the agents are in the same class. Doing the same thing with $c'$, we have shown strong connectivity.

*Pseudomanifold:* consider a sequential partition $c$, and $X$ an $(n-1)$-face of $c$. Let $a$ be the agent which is not in $X$. We distinguish three cases:

- Either $a$ belongs to a concurrency class $c_i$ with $|c_i| > 1$. By the observation above, we can split it and replace it by $\{a\}, (c_i \setminus \{a\})$, so that only agent $a$ has a different view. So, $X$ belongs to exactly two $n$-simplices.

- Or the concurrency class of $a$ is $\{a\}$, and it is not the last one. We can merge it with the next concurrency class, once again, only agent $a$ has a different view. So, $X$ belongs to exactly two $n$-simplices.

- Or the last concurrency class of $c$ is $\{a\}$, in which case no other sequential partition $c'$ can have $X$ as a face. So, $X$ is on the boundary.

Now suppose $M$ is a pseudomanifold. In $M[\mathcal{IS}]$, we take one component of $\mathcal{IS}$ for each facet of $M$, and glue them along their boundary (see the drawing on page 24). According to what was said above, the boundary of a component of $\mathcal{IS}$ concerning agents $A \setminus \{a\}$ corresponds to a scheduling where $a$ is alone in the last concurrency class. So none of the other agents have seen $a$, and their view remains the same independently of $a$'s input value. This ensures that the components are glued properly.

Therefore, a $(n-1)$-simplex of $M[\mathcal{IS}]$ belongs to exactly two facets if it is in the interior of a component, or if it is on the boundary of a component but in the interior of $M$. It belongs to exactly one facet if it is on the boundary of a component, which is itself on the boundary of $M$. $\square$

Sperner's Lemma is a very useful and well-known result in combinatorial topology. It is the combinatorial counterpart of Brouwer's fixed point theorem. We just state the simplest version, for a subdivided triangle in two dimensions. The proof for a more general version of Sperner's Lemma that works in any pseudomanifold of dimension $n$ can be found in many places, e.g. [7, 32] or in [5, Chapter 9], and the excellent exposition in [33, Chapter 1].

Note that Sperner's lemma involves a coloring of the vertices of a simplicial complex. But here the colors in a simplex are not required to be distinct. When we apply it for set agreement, the colors will not correspond to agents but to decision values: therefore, we call them $\{0, 1, 2\}$ to avoid confusion.

**Lemma 2** (Sperner's Lemma)**.** *Let $T$ be a triangle with vertices $A, B, C$, and $S$ be a subdivision of $T$ into smaller triangles. The vertices of $S$ are colored according to the following rules:*

- *The corner vertices $A, B, C$ have color $0, 1, 2$, respectively.*
- *Each vertex along the subdivided boundary connecting two corner vertices is colored with one of the two colors of the corners. E.g., every vertex on the boundary connecting $A$ and $B$ has color $0$ or $1$.*
- *The vertices inside of the subdivision (not on the boundary) can have any of the colors $0, 1, 2$.*

*A coloring satisfying these properties is called a Sperner coloring. Then the Lemma says that given any Sperner coloring, there is an odd number of triangles in the subdivision $S$ whose vertices have distinct colors. In particular, there is at least one.*

## 4. A DEL semantics for distributed task computability

A goal of this section is to show that indeed the higher dimensional topological structure exposed by simplicial models (as opposed to the implicit information in the 1-dimensional structure of a Kripke model) has direct implications about the knowledge that processes can gain by communicating with each other. To do so, we first describe a way of specifying a target for knowledge gaining using DEL, in Section 4.1. We describe this target in a way that corresponds to the notion of a task, and so establish a bridge between epistemic logic semantics and distributed computability, and present it in Section 4.2. In Section 4.3 and 4.4 we present implications about the solvability of specific, well-studied tasks in distributed computability.

### *4.1. Tasks*

Consider the situation where a set of agents $A$ starts in an initial global state, defined by input values given to each agent. The values are local, in the sense that each agent knows its own initial value, but not necessarily the values given to other agents. The agents communicate to each other their initial values, via the immediate snapshot action model $\mathcal{IS}$ of Section 3.2. Then, based on the information each agent has after communication, the agent produces an output value. A task specifies the output values that the agents may decide, when starting in a given input state. Tasks have been studied since early on in distributed computability [34]. Here we provide, for the first time, a DEL semantics for tasks. Namely, we use DEL in a novel way to provide a specification of the problem that a set of agents should solve.

Consider a simplicial model $\mathcal{I} = \langle I, \chi, \ell \rangle$ called the *initial simplicial model*. Each facet of $I$, with its labeling $\ell$, represents a possible initial configuration. Let us fix $\mathcal{I}$ to be the binary inputs model of Example 2, to illustrate the ideas, and because it appears frequently in distributed computing.

A *task* for $\mathcal{I}$ is a simplicial action model $\mathcal{T} = \langle T, \chi, \mathsf{pre} \rangle$ for agents $A$, where each facet is of the form $X = \{\langle b, d_b \rangle, \langle g, d_g \rangle, \langle w, d_w \rangle\}$, where the values $d_b, d_g, d_w$ are taken from an arbitrary domain of *output values*. Each such $X$ has a precondition that is true in one or more facets of $\mathcal{I}$, interpreted as "if the input configuration is a facet in which $\mathsf{pre}(X)$ holds, and every agent $a \in A$ decides on the value $d_a$, then this is a valid execution".

**Example 6.** The most important task in distributed computing is *binary consensus*. Assume we have three agents $A = \{b, g, w\}$, and the input model $\mathcal{I}$ is still the binary input model from Example 2. This means that every combination of 0's and 1's is a possible initial configuration. At the end of the computation, each agent must decide on an output value, which can be either 0 or 1. We write $d_b, d_g, d_w$ for the decision value of agent $b, g, w$ respectively. The goal of the task is to achieve the following properties:

- *Agreement:* the agents must decide on the same value, i.e. $d_b = d_g = d_w$.
- *Validity:* the agreed value must be one of the three inputs.

Say, for example, that the three inputs are $(0, 1, 0)$: then the three outputs can be either $(1, 1, 1)$ or $(0, 0, 0)$. On the other hand, if the three inputs are $(1, 1, 1)$, then the only possible output is $(1, 1, 1)$, because agreeing on value 0 would contradict the validity condition.

Formally, this task is described by an action model $\mathcal{T} = \langle T, \chi, \mathsf{pre}\rangle$. The simplicial complex $T$ has only two facets, corresponding to the two possible kinds of outputs: $X_0$ where all decisions are 0, and $X_1$ where all decisions are 1. These two facets are disjoint: this means, in classical DEL terms, that we have two actions $X_0$ and $X_1$, and there is no indistinguishability relation between them. The precondition $\mathsf{pre}(X_0)$ must be true in all the facets of $\mathcal{I}$ where the output $(0, 0, 0)$ is valid, i.e., whenever at least one of the agent has input value 0. If $p_{a,x}$ is the atomic formula saying that agent $a$ has input value $x$, we take $\mathsf{pre}(X_0) = p_{b,0} \vee p_{g,0} \vee p_{w,0}$. Similarly, $\mathsf{pre}(X_1)$ is true whenever at least one agent has input 1, i.e., $\mathsf{pre}(X_1) = p_{b,1} \vee p_{g,1} \vee p_{w,1}$.

In the picture below, a portion of the binary input complex is represented on the left, and the two facets $X_0$ and $X_1$ of the task action model are pictures on the right. The arrows labeled "$\mathsf{pre}$" indicate in which facets of the input model the precondition relations of $X_0$ and $X_1$ are satisfied.

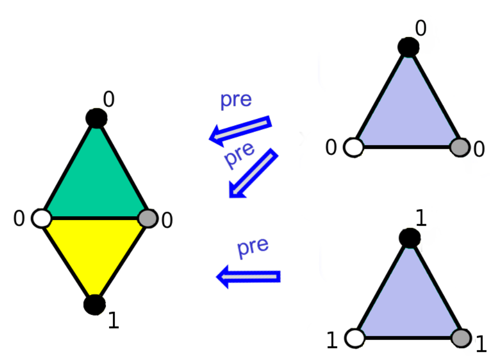

**Example 7.** The *k-set agreement* task is a weaker version of consensus. For this task, we have a set of $k+1$ distinct input values: in the initial configuration, each agent has one input among $\{0, \ldots, k\}$. After the computation, each process must decide on an output value such that:

- $k$-agreement: the set of decision values is of size at most $k$.
- Validity: all the decision values are among the inputs.

This generalization of consensus has been well studied in distributed computability [5]. We will prove in section 4.4 that 2-set agreement among 3 processes is not solvable using immediate snapshots.

*4.2. Semantics of task solvability*

Given the simplicial input model $\mathcal{I}$ and a communication model $\mathcal{A}$ such as $\mathcal{IS}$, we get the *simplicial protocol model* $\mathcal{I}[\mathcal{A}]$, that represents the knowledge gained by the agents after executing $\mathcal{A}$. To solve a task $\mathcal{T}$, each agent, based on its own knowledge, should produce an output value, such that the vector of output values corresponds to a facet of $\mathcal{T}$, respecting the preconditions of the task.

The following gives a formal epistemic logic semantics to task solvability. Recall that a morphism $\delta$ of simplicial models is a chromatic simplicial map that preserves the labeling: $\ell'(\delta(v)) = \ell(v)$. Also recall that the product update model $\mathcal{I}[\mathcal{A}]$ is a sub-complex of the Cartesian product $\mathcal{I} \times \mathcal{A}$, whose vertices are of the form $(i, ac)$ with $i$ a vertex of $\mathcal{I}$ and $ac$ a vertex of $\mathcal{A}$. We write $\pi_{\mathcal{I}}$ for the first projection on $\mathcal{I}$, which is a morphism of simplicial models.

**Definition 2.** *A task $\mathcal{T}$ is* solvable *in $\mathcal{A}$ if there exists a morphism $\delta : \mathcal{I}[\mathcal{A}] \to \mathcal{I}[\mathcal{T}]$ such that $\pi_I \circ \delta = \pi_I$, i.e., the diagram of simplicial complexes below commutes.*

$$\begin{tikzcd} & \mathcal{I}[\mathcal{A}] \arrow[dl, "\pi_{\mathcal{I}}"'] \arrow[d, dashed, "\delta"] \\ \mathcal{I} & \mathcal{I}[\mathcal{T}] \arrow[l, "\pi_{\mathcal{I}}"] \end{tikzcd}$$

The justification for this definition is the following. A facet $X$ in $\mathcal{I}[\mathcal{A}]$ corresponds to a pair $(i, ac)$, where $i \in \mathcal{F}(\mathcal{I})$ represents input value assignments to all agents, and $ac \in \mathcal{F}(\mathcal{A})$ represents an action, codifying the communication exchanges that took place. The morphism $\delta$ takes $X$ to a facet $\delta(X) = (i, dec)$ of $\mathcal{I}[\mathcal{T}]$, where $dec \in \mathcal{F}(\mathcal{T})$ is the set of decision values that the agents will choose in the situation $X$. Moreover, $\mathsf{pre}(dec)$ holds in $i$, meaning that $dec$ corresponds to valid decision values for input $i$. The commutativity of the diagram expresses the fact that both $X$ and $\delta(X)$ correspond to the same input assignment $i$. Now consider a single vertex $v \in X$ with $\chi(v) = a \in A$. Then, agent $a$ decides its value solely according to its knowledge in $\mathcal{I}[\mathcal{A}]$: if another facet $X'$ contains $v$, then $\delta(v) \in \delta(X) \cap \delta(X')$, meaning that $a$ has to decide on the same value in both situations.

The diagram above has two illuminating interpretations. First, by Theorem 3, we know that the knowledge of each agent can only decrease (or stay constant) along the $\delta$ arrow. So, any (positive) formula which is known in $\mathcal{I}[\mathcal{T}]$ should already be known in $\mathcal{I}[\mathcal{A}]$. In other words, the goal of the agents is to improve knowledge through communication, by going from $\mathcal{I}$ to $\mathcal{I}[\mathcal{A}]$, in order to match the knowledge required by $\mathcal{I}[\mathcal{T}]$. Secondly, the possibility of solving a task depends on the existence of a certain simplicial map from the complex of $\mathcal{I}[\mathcal{A}]$ to the complex of $\mathcal{I}[\mathcal{T}]$. Recall that a simplicial map is the discrete equivalent of a continuous map, and hence task solvability is of a topological nature.

*4.3. Applications*

Here we describe how to use our DEL setting to analyze solvability in the immediate-snapshot model of three well-studied distributed computing tasks: consensus, approximate agreement, and set agreement. Their solvability is already well-understood; in particular, Theorems 5, 6 and 7 are already known. The novelty here is that we give new proofs based on logical arguments. The aim of these proofs is to understand the epistemic logic content of the known topological arguments that are used to show unsolvability.

*Consensus.* Let $\mathcal{I} = \langle I, \chi, \ell \rangle$ be the initial simplicial model for binary input values (see Example 2), and $\mathcal{T} = \langle T, \chi, \mathsf{pre} \rangle$ be the action model for binary consensus (see Example 6). Thus, $T$ has only two facets, $X_0$ where all decisions are 0 and $X_1$, where all decisions are 1. The underlying complex of $\mathcal{I}[\mathcal{T}]$ consists of two disjoint simplicial complexes: $I_0 \times X_0$ and $I_1 \times X_1$, where $I_0$ consists of all input facets with at least one 0, and $I_1$ consists of all input facets with at least one 1. Notice that, in fact, each of the two complexes $I_i \times X_i$, for $i \in \{0, 1\}$, is isomorphic to $I_i$, since $X_i$ consists of just one facet.

To show that binary consensus cannot be solved by the immediate snapshot protocol, we must prove that the map $\delta : \mathcal{I}[\mathcal{A}] \to \mathcal{I}[\mathcal{T}]$ of Definition 2 does not exist. The usual proof of impossibility uses a topological obstruction to the existence of $\delta$. Here, instead, we exhibit a logical obstruction.

**Theorem 5.** *The binary consensus task is not solvable by $\mathcal{IS}$.*

*Proof.* We first state some required knowledge at $\mathcal{I}[\mathcal{T}]$ to solve the task. Let $\varphi_i$ be the formula denoting that at least one agent has input $i$. More precisely, if we write $\mathsf{input}_i^a$ for the atomic proposition "agent $a$ has input value $i$", we take $\varphi_i := \bigvee_{a \in A} \mathsf{input}_i^a$. Then, for each $i \in \{0,1\}$, at any facet $Y$ of $I_i \times X_i$, there is common knowledge that at least one input is $i$, that is, $\mathcal{I}[\mathcal{T}], Y \models C_A \varphi_i$. Indeed, every facet $Z$ in the connected component of $Y$ must still be in $I_i \times X_i$, and by definition of $I_i$ there is at most one agent with input $i$, so $\mathcal{I}[\mathcal{T}], Z \models \varphi_i$.

Now, consider the simplicial model $\mathcal{I}[\mathcal{IS}]$ for the immediate snapshot action model. By Lemma 1 the underlying complex of $\mathcal{I}[\mathcal{IS}]$ is strongly connected, therefore there is a path from the facet with valuation indicating that all inputs are 0 to the facet where all inputs are 1. Therefore, at any facet $X$ of $\mathcal{I}[\mathcal{IS}]$, we have $\mathcal{I}[\mathcal{IS}], X \not\models C_A \varphi_i$, for both $i = 0$ and $i = 1$.

Finally, we know that morphisms of simplicial models cannot "gain knowledge about the world" from Theorem 3, and so, there cannot be a morphism $\delta$ from $\mathcal{I}[\mathcal{IS}]$ to $\mathcal{I}[\mathcal{T}]$, by the two previous properties. $\square$

Two observations. First, notice that the proof argument holds for any other model, instead of $\mathcal{IS}$, which is connected. This is the case for any number of communication rounds by wait-free asynchronous agents [35], and even if only one may crash in a message passing system [36], which are the classic consensus impossibility results. Secondly, the usual topological argument for impossibility is the following: because simplicial maps preserve connectivity, $\delta$ cannot send a connected simplicial complex into a disconnected simplicial complex. Notice how in both the logical and the topological proofs, the main ingredient is a connectedness argument.

*Approximate agreement.* We now discuss a weaker version of the consensus task, where agents are required to decide on values which are close to each other, not necessarily equal. It turns out that no matter how close to each other one requires the agents to decide, this task is solvable in the immediate snapshot (multi-round) model. Many versions of this task have been considered. We present here a simple one, for two agents, $g$ and $w$.

The input complex is the binary input complex for two agents, depicted on the left of Example 2: so, every possible combination of 0 and 1 can be assigned to the two agents. Their goal will be to output real values in the interval $[0,1]$, such that:

- if their input is the same (i.e., 00 or 11), they should both choose their input value as output.
- if their input is different (i.e., 01 or 10), they should decide on values $d_g$ and $d_w$ such that $|d_g - d_w| \leq \varepsilon$, for some fixed parameter $\varepsilon \in [0, 1]$.

In order to be able to work with finite models, we define a discrete version of this task, $N$-*approximate agreement*. The output values are only allowed to be of the form $k/N$ for $0 \leq k \leq N$. The two decision values should be within $1/N$ of each other: $|d_g - d_w| \leq 1/N$.

Let $\mathcal{I} = \langle I, \chi, \ell \rangle$ be the input simplicial model for two agents with binary inputs, and $\mathcal{T} = \langle T, \chi, \mathsf{pre} \rangle$ be the following action model. The set of vertices of $T$ is $\mathcal{V}(T) = \{(a, k/N) \mid a \in A \text{ and } 0 \leq k \leq N\}$. The facets of $T$ are edges $X_{k,k'} = \{(g, k/N), (w, k'/N)\}$ with $|k - k'| \leq 1$. The color of a vertex is $\chi(a, k/N) = a$. The precondition $\mathsf{pre}(X_{0,0})$ is true in the worlds 00, 01 and 10 of $\mathcal{I}$; the precondition $\mathsf{pre}(X_{N,N})$ is true in the worlds 11, 01 and 10; and all the other preconditions $\mathsf{pre}(X_{k,k'})$ are true in the worlds 01 and 10. In the figure below are depicted the input model $\mathcal{I}$ (left) and the action model $\mathcal{T}$ (right), for $N = 5$.

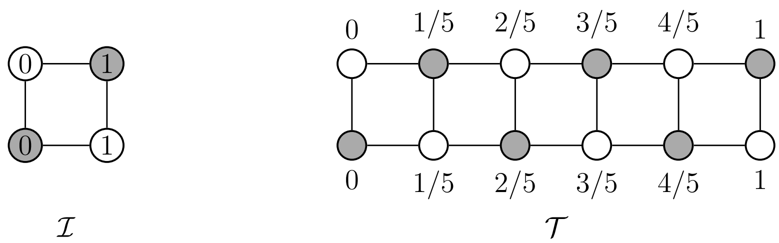


The product update simplicial model $\mathcal{I}[\mathcal{T}]$ is depicted in the next figure. Since its shape differs slightly depending on the parity of $N$, two cases are depicted: $N = 4$ and $N = 5$. The numbers depicted in the nodes are the atomic propositions describing the input values from $\mathcal{I}$. The decision values (of the form $k/5$ or $k/4$) are implicit, the first column of nodes corresponds to the decision value 0, the second column is decision value $1/5$, and so on. For example, the world marked $X$ on the figure corresponds to the situation where $g$ started with value 1 and decided value $2/5$, $w$ started with value 0 and decided value $3/5$. So, $X$ represents a correct execution of the 5-approximate agreement task. The world $Y$ is another possible execution, where $g$ decides $3/5$ and $w$ decides $2/5$.

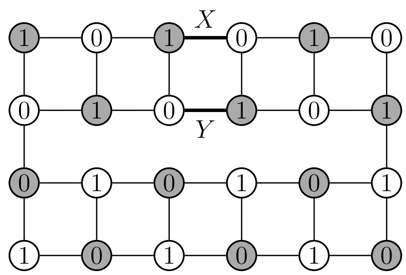


$\mathcal{I}[\mathcal{T}]$ for $N = 5$

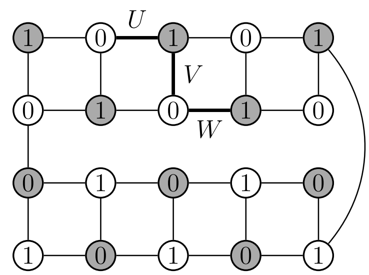


$\mathcal{I}[\mathcal{T}]$ for $N = 4$

Note that the world $X$ on the figure is the one with the most knowledge in the following sense. We write $\varphi_{01}$ for the formula expressing that the two inputs are different, and $E\varphi = K_g\varphi \wedge K_w\varphi$ for the group knowledge of $\varphi$ among the agents $\{g, w\}$. Then, we have $\mathcal{I}[\mathcal{T}], X \models E^3\varphi_{01}$, where $E^3$ denotes three nested $E$ operators. On the other hand, the world $Y$ has less knowledge: we have $\mathcal{I}[\mathcal{T}], Y \models E^2\varphi_{01}$, but $\mathcal{I}[\mathcal{T}], Y \not\models E^3\varphi_{01}$. Similarly, in the case where $N = 4$, the three worlds $U, V, W$ satisfy the formula $E^2\varphi_{01}$, but they do not satisfy $E^3\varphi_{01}$.

**Lemma 3.** *In the simplicial model $\mathcal{I}[\mathcal{T}]$ for the $N$-approximate agreement task, there are worlds $X, Y$ (if $N$ is odd) or $U, V, W$ (if $N$ is even) which satisfy the formula $E^k\varphi_{01}$, for $k = \lfloor N/2 \rfloor$.*

*Proof.* We choose the worlds in the "middle" of the model $\mathcal{I}[\mathcal{T}]$, as shown in the pictures. More formally, recall that the vertices of $\mathcal{I}[\mathcal{T}]$ are defined as tuples $(a, i, d)$ where $a$ is an agent, $i$ its input value and $d$ its decision value. For instance, the world $X$ is defined as the edge $\{(g, 1, \lfloor N/2 \rfloor), (w, 0, \lfloor N/2 \rfloor + 1)\}$, and so on. Checking that the formula is satisfied in these worlds simply consists in computing the length of the shortest path to one of the 00 or 11 edges on the sides. Note that $X$ also satisfies the formula $E^{k+1}\varphi_{01}$, but we will not use that fact. $\square$

We now study the solvability of this task in the $r$-round iterated immediate snapshot model $\mathcal{IS}^r$. In dimension 1, each iteration of the immediate snapshot subdivides each edge into three parts. The picture below shows the input model $\mathcal{I}$, the model $\mathcal{I}[\mathcal{IS}]$ after one round, and the model $\mathcal{I}[\mathcal{IS}^2]$ after two rounds.

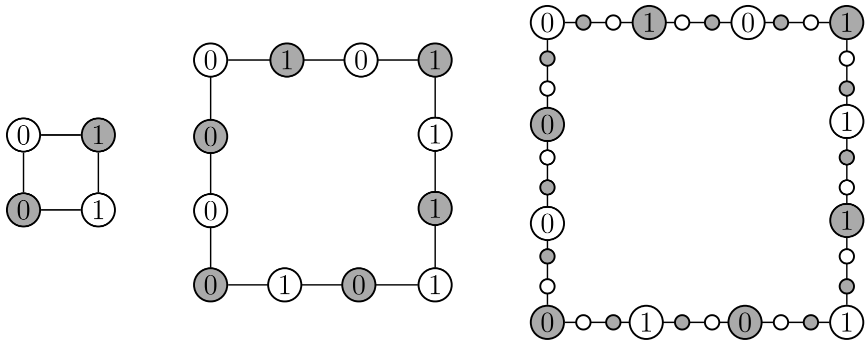

$\mathcal{I}$ $\quad\quad$ $\mathcal{I}[\mathcal{IS}]$ $\quad\quad$ $\mathcal{I}[\mathcal{IS}^2]$

**Lemma 4.** *In the $r$-round immediate snapshot model $\mathcal{I}[\mathcal{IS}^r]$, there is no world $X$ such that $\mathcal{I}[\mathcal{IS}^r], X \models E^k\varphi_{01}$, for $k = \lceil 3^r/2 \rceil$.*

*Proof.* After $r$ rounds of immediate snapshots, each of the four edges of the input model $\mathcal{I}$ have been subdivided into $3^r$ edges. Thus, every world is at a distance at most $k-1$ from the nearest world with inputs 00 or 11. □

Putting the two lemmas together, we get the following result:

**Theorem 6.** *The $N$-approximate agreement task is not solvable in the $r$-round iterated immediate-snapshot model, when $N \geq 3^r + 1$.*

*Proof.* Assume by contradiction that the task is solvable. Then, we would have a map $\delta : \mathcal{I}[\mathcal{IS}^r] \to \mathcal{I}[\mathcal{T}]$. Our goal is to find a contradiction using Lemma 3. To achieve this, we should find a formula $\varphi$ and a world $Z$ of $\mathcal{I}[\mathcal{IS}^r]$, such that $\varphi$ is false in $Z$ but true in $\delta(Z)$. We choose the formula $\varphi := E^k\varphi_{01}$ for $k = \lfloor N/2 \rfloor$. Since $N \geq 3^r + 1$ implies $\lfloor N/2 \rfloor \geq \lceil 3^r/2 \rceil$, we know by Lemma 4 that this formula is false in every world $Z$ of $\mathcal{I}[\mathcal{IS}^r]$. All that remains to do is prove that there exists a world of $\mathcal{I}[\mathcal{T}]$, *which is in the image of* $\delta$, and where the formula $\varphi$ is true.

Since $\mathcal{I}[\mathcal{IS}^r]$ is connected, and simplicial maps preserve connectedness, its image $\delta(\mathcal{I}[\mathcal{IS}^r])$ is connected too. Moreover, the world 00 and the world 11 of $\mathcal{I}[\mathcal{T}]$ must both be in the image of $\delta$, because of the commutative diagram of Definition 2. By connectedness, at least one of the middle worlds $X, Y$ or $U, V, W$ must belong to the image of $\delta$. By Lemma 3, this world satisfies $\varphi$, which concludes the proof.

□

Conversely, it is known in distributed computing that $N$-approximate agreement is solvable in $\mathcal{IS}^r$ whenever $N \leq 3^r$ [5]. The proof of the above theorem sheds light on the required knowledge to solve approximate agreement: while consensus is about reaching common knowledge, approximate agreement is about reaching some finite level of nested knowledge.

*Set agreement.* Let $\mathcal{I} = \langle I, \chi, \ell \rangle$ be the initial simplicial model for $A = \{b, w, g\}$, and three possible input values, $\{0, 1, 2\}$. Let $\mathcal{T} = \langle T, \chi, \mathsf{pre} \rangle$ be the action model for 2-set agreement (see Example 7), requiring that each agent decides on one of the input values, and at most 2 distinct values are decided. More precisely, the vertices of $T$ are of the form $v_d^a$ with $a \in A$ and $d \in \{0, 1, 2\}$, and the facets of $T$ are $X_{d_0,d_1,d_2} = \{v_{d_0}^b, v_{d_1}^g, v_{d_2}^w\}$, for each vector $d_0, d_1, d_2$, such that $d_i \in \{0, 1, 2\}$ and $|\{d_0, d_1, d_2\}| \leq 2$. The preconditions are $\mathsf{pre}(X_{d_0,d_1,d_2}) = \varphi_{d_0} \wedge \varphi_{d_1} \wedge \varphi_{d_2}$, where the formula $\varphi_i$ expresses that at least one agent has input $i$.

To get an intuition about what $\mathcal{I}$ and $\mathcal{T}$ look like, first remark that $\mathcal{T}$ can be decomposed into three spheres: $T_{01}$, $T_{12}$ and $T_{02}$, where $T_{ij}$ is the subcomplex of $T$ whose decision values are in $\{i, j\}$. Each of these $T_{ij}$ is isomorphic to the binary input complex of Example 2, because it allows every possible combination of the two decision values $i$ and $j$. Moreover, the sphere $T_{01}$ and the sphere $T_{12}$ share a facet, namely, the one where all processes have decision value 1. Similarly, $T_{12}$ and $T_{02}$ are glued along the facet where everyone decides 2; and $T_{02}$ and $T_{01}$ are glued along the facet where everyone decides 0. Thus, the complex $T$ is a "necklace" of three 2-dimensional spheres. In particular, it has a 1-dimensional hole in the middle, so $T$ is not simply connected.

On the other hand, the input complex $\mathcal{I}$ allows every possible combination of the three values $0, 1, 2$. So, $\mathcal{I}$ contains a subcomplex isomorphic to $T$, consisting of all the input configurations where there are at most two distinct input values. But it also allows configurations where the three inputs are different (there are six such configurations), which fill the hole in $T$. Hence, $\mathcal{I}$ is isomorphic to a wedge of spheres. In particular, $\mathcal{I}$ is simply connected. Using these topological facts about $\mathcal{I}$ and $\mathcal{T}$, one can prove the following Theorem (see [5]).

**Theorem 7.** *The* 2*-set agreement task is not solvable by* $\mathcal{IS}$*.*

Our goal here is to prove Theorem 7 using our dynamic epistemic logic framework. If we tried to mimic what we did above for consensus and approximate agreement, the idea of the proof would be the following. Assume by contradiction that the task is solvable, i.e., there is a suitable map $\delta : \mathcal{I}[\mathcal{IS}] \to \mathcal{I}[\mathcal{T}]$. Then, we need to find a positive formula $\varphi$ such that:

- $\varphi$ is false in some world $X$ of $\mathcal{I}[\mathcal{IS}]$,
- and $\varphi$ is true in the world $\delta(X)$ of $\mathcal{I}[\mathcal{T}]$.

This would be a contradiction according to Theorem 3. Unfortunately, we have not been able to find such a formula. In fact, we conjecture that no suitable formula $\varphi$ exists. In a recent paper [37], we exhibited a task which is known to be unsolvable, but where no epistemic logic formula $\varphi$ can witness it. For the set agreement task, this question is still open.

In the next section, to get around this problem, we will strengthen the language of our logic by adding new atomic propositions. This will allow us to express more properties of the models.

### *4.4. A logical proof of impossibility for set agreement*

In the 2-set agreement example of Section 4.3, the atomic propositions on the input model $\mathcal{I}$ express what the input value of each process is. We write $\mathsf{input}^a_i$ for the atomic proposition saying that "agent $a$ has input value $i$"; thus the set of atomic propositions is $AP = \{\mathsf{input}^a_i \mid a \in \{b, g, w\}, i \in \{0, 1, 2\}\}$. As we mentioned in the previous section, it seems that epistemic formulas in the language $\mathcal{L}_{\mathcal{CK}}(A, AP)$ are too weak to express the reason why set agreement is not solvable. Indeed, the specification of the set agreement task relies heavily on the relationship between the inputs and the outputs; but these formulas only talk about the inputs. To express the knowledge which must be achieved in order to solve the task, we might want to write a formula like the following one, translating the two informal conditions of Example 7:

$$\begin{aligned}\varphi \quad = \quad & \forall d_0, d_1, d_2.\ \mathsf{decide}^b_{d_0} \wedge \mathsf{decide}^g_{d_1} \wedge \mathsf{decide}^w_{d_2} \\ & \qquad \Longrightarrow \quad |\{d_0, d_1, d_2\}| \leq 2 \\ & \qquad\qquad \wedge\, \exists a.\ \mathsf{input}^a_{d_0} \wedge \exists a.\ \mathsf{input}^a_{d_1} \wedge \exists a.\ \mathsf{input}^a_{d_2}\end{aligned}$$

The quantifiers $\forall$ and $\exists$ are just shortcuts for conjunctions and disjunctions ranging over all possible values for $d_0, d_1, d_2$ and $a$ (there are only finitely many of them). The condition $|\{d_0, d_1, d_2\}| \leq 2$ is not in the language $\mathcal{L}_{\mathcal{CK}}(A, AP)$, but we could simply replace it by "true" or "false" in each case, depending on the values of $d_0, d_1, d_2$. The only thing that we really cannot express in this language is the $\mathsf{decide}^a_d$ formulas, whose intuitive meaning would be: "the agent $a$ decides on the value $d$".

In this section, we describe an ad-hoc construction whose sole purpose is to make sense of the above formula $\varphi$. Namely, we modify the product-update model $\mathcal{I}[\mathcal{T}]$ by arbitrarily adding new atomic propositions of the form $\mathsf{decide}^a_d$. Thus, we obtain a new model $\widehat{\mathcal{I}[\mathcal{T}]}$ on which the formula $\varphi$ can be interpreted. This will allow us to prove the impossibility of solving set agreement (Theorem 7), using a logical argument similar to the ones of the previous section.

As pointed out by one of the reviewers, it turns out that our construction is in fact very close to *DEL with factual change* [38, 39]. Indeed, in the standard DEL framework, given a model $M$ and an action model $\mathcal{A}$, the worlds of $M[\mathcal{A}]$ are by definition pairs $(w,t)$ where $w$ is a world of $M$ and $t$ an action of $\mathcal{A}$. What we will do in this section amounts to considering new atomic propositions $p_t$ for each $t$, allowing us to speak about which actions have occurred. This information could alternatively be captured using factual change, by considering that the atomic propositions $p_t$ are already present in the original model $M$, but all set to $\mathsf{false}$. Then, the product update operation is able to *change* the value of these propositions in order to record the actions that happened. Simplicial action models with factual change are described in a sequel to this paper [40]; but reformulating the proof below in terms of factual change is left for future work.

Recall that, by definition, the worlds of $\mathcal{I}[\mathcal{T}]$ are pairs $(x, X_{d_0,d_1,d_2})$ where $x$ is a world of $\mathcal{I}$ and $X_{d_0,d_1,d_2}$ is a world of $\mathcal{T}$, whose intuitive meaning is "agent $b$ decides $d_0$, agent $g$ decides $d_1$, and agent $w$ decides $d_2$". Our new set of atomic propositions is written $\widehat{AP} = AP \cup \{\mathsf{decide}^a_d \mid a \in A, d \in \{0,1,2\}\}$. We define the *extended task model* $\widehat{\mathcal{I}[\mathcal{T}]}$ as follows.

- Its vertices are triples $(a,i,d)$ with $a \in A$ is an agent, $i \in \{0,1,2\}$ is an input value and $d \in \{0,1,2\}$ is a decision value. The facets are of the form $\{(b,i_b,d_b),(g,i_g,d_g),(w,i_w,d_w)\}$ such that $|\{d_b,d_g,d_w\}| \leq 2$ and $\{d_b,d_g,d_w\} \subseteq \{i_b,i_g,i_w\}$.
- The coloring map is $\chi(a,i,d) = a$.
- The atomic proposition labeling is $\ell(a,i,d) = \{\mathsf{input}^a_i, \mathsf{decide}^a_d\}$.

In other words, $\widehat{\mathcal{I}[\mathcal{T}]}$ is just a copy of $\mathcal{I}[\mathcal{T}]$, where we moreover labeled the vertices with the $\mathsf{decide}_d^a$ atomic propositions accordingly. We can easily check that the formula $\varphi$ defined earlier is true in every world of $\widehat{\mathcal{I}[\mathcal{T}]}$.

Now, we would also like the formula $\varphi$ to make sense in the protocol complex $\mathcal{I}[\mathcal{IS}]$. This model does not have any information about decision values: it only describes the input values, and which execution has occurred. However, it is precisely the role of the simplicial map $\delta : \mathcal{I}[\mathcal{IS}] \to \mathcal{I}[\mathcal{T}]$ to assign decision values to each world of $\mathcal{I}[\mathcal{IS}]$. Thus, given such a map $\delta$, we can lift it to a map $\widehat{\delta} : \widehat{\mathcal{I}[\mathcal{IS}]} \to \widehat{\mathcal{I}[\mathcal{T}]}$ as the following lemma states.

**Lemma 5.** *Let $M = \langle C, \chi, \ell \rangle$ be a simplicial model over the set of agents $A$ and atomic propositions $AP$, and let $\delta : M \to \mathcal{I}[\mathcal{T}]$ be a morphism of simplicial models. Then there is a unique model $\widehat{M} = \langle C, \chi, \widehat{\ell} \rangle$ over $\widehat{AP}$, where $\widehat{\ell}$ agrees with $\ell$ on $AP$, such that $\widehat{\delta} : \widehat{M} \to \widehat{\mathcal{I}[\mathcal{T}]}$ is still a morphism of simplicial models.*

*Proof.* All we have to do is label the worlds of $M$ with the $\mathsf{decide}_d^a$ atomic propositions, such that the resulting $\widehat{\delta}$ is a morphism of simplicial models. Thus, we define $\widehat{\ell} : \mathcal{V}(C) \to \mathscr{P}(\widehat{AP})$ as $\widehat{\ell}(v) = \ell(v) \cup \{\mathsf{decide}_d^a\}$, whenever $\delta(v) = (a, i, d) \in \mathcal{I}[\mathcal{T}]$. Then $\delta$ is still a chromatic simplicial map (since we did not change the underlying complexes nor their colors), and moreover we have $\widehat{\ell}(v) = \ell_{\widehat{\mathcal{I}[\mathcal{T}]}}(\delta(v))$ for all $v$. The model $\widehat{M}$ is unique since any other choice of $\widehat{\ell}(v)$ would have broken this last condition, so $\delta$ would not be a morphism of simplicial models. $\square$

We can finally prove that 2-set agreement is not solvable by the immediate snapshot protocol.

*Proof of Theorem 7.* We assume by contradiction that it is solvable, i.e., that there exists a morphism $\delta : \mathcal{I}[\mathcal{IS}] \to \mathcal{I}[\mathcal{T}]$ that makes the diagram of Definition 2 commute. By Lemma 5, $\delta$ can be lifted to a map $\widehat{\delta} : \widehat{\mathcal{I}[\mathcal{IS}]} \to \widehat{\mathcal{I}[\mathcal{T}]}$. Our goal is to contradict Theorem 3, using the formula $\varphi$ defined in the beginning of this section. As we remarked earlier, $\varphi$ is true in every world of $\widehat{\mathcal{I}[\mathcal{T}]}$. Thus, all we have to do is prove that there exists one world $X$ of $\widehat{\mathcal{I}[\mathcal{IS}]}$ where the formula is false.

Consider one facet of the input model $\mathcal{I}$ where all the agents start with distinct values, for example, $I_{012} = \{(b, 0), (g, 1), (w, 2)\} \in \mathcal{I}$. This facet induces a subcomplex of $\mathcal{I}[\mathcal{IS}]$, which we write $I_{012}[\mathcal{IS}]$. It consists of all the worlds of the form $(I_{012}, c^{012})$, where $c$ is a sequential partition of the agents. The figure below represents the subcomplex $I_{012}[\mathcal{IS}]$; the values written inside the nodes represent the $\mathsf{input}_i^a$ atomic propositions.

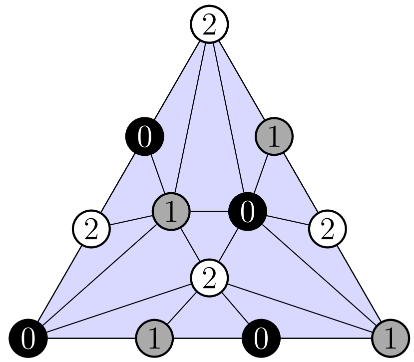


According to Lemma 1, $I_{012}[\mathcal{IS}]$ is a pseudomanifold with boundary. Moreover, in the extended model $\widehat{\mathcal{I}[\mathcal{IS}]}$, the decision values induce a Sperner coloring on its boundary. Indeed, let us first look at the extremal vertices of the triangle, for example the bottom left $(b, 0)$ vertex. The view of this vertex is $0\bot\bot$, so, it also belongs to the subcomplex $I_{000}[\mathcal{IS}]$ where everyone started with value 0. In this subcomplex, all the decision values must be 0, otherwise the validity condition of the formula $\varphi$ would be broken (which would conclude our proof). So, this vertex must have the decision value 0. Simularly, the bottom right vertex $(g, 1)$ must decide 1, and the top vertex $(w, 2)$ must decide 2. We can reason similarly with the edge vertices. For example, the view of the vertices on the bottom edge is $01\bot$. So, they also belong to the subcomplex $I_{010}[\mathcal{IS}]$ where white has value 0. In this subcomplex, all the decision values must be either 0 or 1, otherwise the validity condition would not be satisfied.

Thus, by Sperner's Lemma, there must be one world $X$ of $I_{012}[\mathcal{IS}]$ with three different decision values. Then, it is easily checked that $\widehat{\mathcal{I}[\mathcal{IS}]}, X \not\models \varphi$, which concludes the proof. $\square$

Note that, for simplicity, we illustrated the above proof with the case of 2-set agreement for 3 agents. The exact same proof would hold to prove the impossibility of solving $k$-set agreement among $k+1$ agents, in the iterated immediate snapshot model.

## 5. Conclusions

We have made a first step in several directions. The first one is the proposal of using simplicial models to expose the higher-dimensional topological information which is implicit in Kripke models. We have argued that this topological information is of interest, for two reasons, which are well-known in distributed computing, but seem to be novel in epistemic logic. First, when agents communicate with each other, there are topological invariants that are preserved, from the initial simplicial model $\mathcal{I}$, to the simplicial model $\mathcal{I}[\mathcal{IS}]$ representing the knowledge after communication has taken place. Secondly, these topological invariants in turn determine the coordinated decision that the agents are able to take. Namely, the solvability of a distributed task for a set of agents depends on such topological information.

We have also made a first step into defining a version of fault-tolerant multi-agent DEL using simplicial complexes as models, providing a different perspective from the classical knowledge approach based on Kripke frames. DEL is convenient in this setting, but we expect that other formalisms could have been used, such as directly using interpreted systems.

Usually in logic, a computational problem is given by a formula that specifies it. We have made a first step in defining problem specifications using a DEL action model, instead of a formula. There are interesting open problems in this direction too. We have seen that the impossibility of consensus or approximate agreement can be expressed as a formula, but in sequel work [37], we define a notion of bisimulation for simplicial models and use it to prove that there are tasks for which no such formula exists.

We have thus established a bridge between the theory of distributed computability and epistemic logic. We illustrated the setting with a simple one-round communication action model $\mathcal{IS}$, that corresponds to a well-studied situation in distributed computing, but many other models can be treated similarly.

The research of this paper was initiated with the preliminary report [41]. In another preliminary report [42], we explore knowledge about the model of distributed computing via DEL. A thorough account of the relationship between DEL, simplicial models and distributed computing can be found in [43, Chapter 3]. Many interesting questions are left for future work. For instance, we have developed here all our theory on pure simplicial complexes, where all the facets are of the same dimension. Extending it to complexes with lower dimensional facets would allow us to model situations where the number of agents may vary, which corresponds in the context of distributed computing to detectable failures. In a sequel to this paper [40], we explore in more detail various other issues like bisimulation, locality and belief.

It is known to be undecidable whether a task is solvable in the immediate snapshot model, even for three processes [44, 45], and hence the connection we establish with DEL implies that it is undecidable whether certain knowledge has been gained in multi-round immediate snapshot action models, but further work is needed to study this issue. Future work is needed to study bisimulations and their relation to the simulations studied in task computability [46]. It would be of interest to study other distributed computing settings, especially those which have stronger communication objects available, and which are known to yield complexes that might not preserve the topology of the input complex.

*Acknowledgments.* This work was partially supported by UNAM-PAPIIT IN109917 and IN106520, and France-Mexico ECOS 207560(M12-M01), and part of it was done while S. Rajsbaum was at École Polytechnique, partially funded by the Ecole Polytechnique 2016-2017 and 2017-2018 Visiting Scholar Program. Éric Goubault and Jérémy Ledent were partially supported by the DGA project "Validation of Autonomous Drones and Swarms of Drones" and the academic chair "Complex Systems Engineering" of Ecole Polytechnique-ENSTA-Télécom-Thalès-Dassault-Naval Group-DGA-FX-FDO-Fondation ParisTech. We thank Hans van Ditmarsch, Yoram Moses, David Rosenblueth, Carlos Velarde, and the reviewers for interesting and numerous discussions, remarks and complements to the literature.

## References

[1] A. Baltag, B. Renne, Dynamic epistemic logic, in: The Stanford Encyclopedia of Philosophy, see https://plato.stanford.edu/archives/win2016/entries/dynamic-epistemic/, Metaphysics Research Lab, Stanford University, 2016.

[2] H. v. Ditmarsch, W. van der Hoek, B. Kooi, Dynamic Epistemic Logic, Springer, 2007. doi:10.1007/978-1-4020-5839-4.

[3] A. Baltag, L. Moss, S. Solecki, The logic of common knowledge, public announcements, and private suspicions, in: TARK VII, 1998, pp. 43–56. doi:10.1007/978-3-319-20451-2_38.

[4] H. Attiya, J. Welch, Distributed Computing: Fundamentals, Simulations, and Advanced Topics, 2nd Edition, Wiley, 2004. doi:10.1002/0471478210.

[5] M. Herlihy, D. Kozlov, S. Rajsbaum, Distributed Computing Through Combinatorial Topology, Elsevier-Morgan Kaufmann, 2013. doi:10.1016/C2011-0-07032-1.

[6] N. A. Lynch, Distributed Algorithms, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1996.

[7] H. Attiya, S. Rajsbaum, The combinatorial structure of wait-free solvable tasks, SIAM J. Comput. 31 (4) (2002) 1286–1313. doi:10.1137/S0097539797330689.

[8] M. Saks, F. Zaharoglou, Wait-Free k-Set Agreement is Impossible: The Topology of Public Knowledge, SIAM J. Comput. 29 (5) (2000) 1449–1483.

[9] Y. M. R. Fagin, J. Halpern, M. Vardi, Reasoning About Knowledge, MIT Press, 1995.

[10] J. McCarthy, Artificial Intelligence, Logic and Formalizing Common Sense, Springer Netherlands, Dordrecht, 1989, pp. 161–190. doi:10.1007/978-94-009-2448-2_6.
URL https://doi.org/10.1007/978-94-009-2448-2_6

[11] A. Rubinstein, A. Wolinsky, On the logic of agreeing to disagree type results, Journal of Economic Theory 51 (1) (1990) 184 – 193. doi:https://doi.org/10.1016/0022-0531(90)90057-Q.
URL http://www.sciencedirect.com/science/article/pii/002205319090057Q

[12] S. A. Kripke, A completeness theorem in modal logic, The Journal of Symbolic Logic 24 (1) (1959) 1–14.
URL http://www.jstor.org/stable/2964568

[13] J. C. C. McKinsey, A. Tarski, The algebra of topology, Annals of Mathematics 45 (1) (1944) 141–191.
URL http://www.jstor.org/stable/1969080

[14] J. van Benthem, G. Bezhanishvili, Modal Logics of Space, Springer Netherlands, Dordrecht, 2007, pp. 217–298. doi:10.1007/978-1-4020-5587-4_5.
URL https://doi.org/10.1007/978-1-4020-5587-4_5

[15] A. Castañeda, Y. A. Gonczarowski, Y. Moses, Unbeatable set consensus via topological and combinatorial reasoning, in: PODC, ACM, 2016, pp. 107–116. doi:10.1145/2933057.2933120.

[16] A. Castañeda, Y. A. Gonczarowski, Y. Moses, Unbeatable consensus, in: DISC, no. 8784 in LNCS, Springer, 2014, pp. 91–106. doi:10.1007/978-3-662-45174-8_7.

[17] T. Porter, Interpreted systems and kripke models for multiagent systems from a categorical perspective, Theoretical Computer Science 323 (1) (2004) 235 – 266. doi:10.1016/j.tcs.2004.04.005.

[18] A. R. Lomuscio, R. van der Meyden, M. Ryan, Knowledge in multiagent systems: Initial configurations and broadcast, ACM Trans. Comput. Logic 1 (2) (2000) 247–284. doi:10.1145/359496.359527.
URL http://doi.acm.org/10.1145/359496.359527

[19] A. Baltag, A. Özgün, A. L. Vargas Sandoval, Topo-logic as a dynamic-epistemic logic, in: A. Baltag, J. Seligman, T. Yamada (Eds.), Logic, Rationality, and Interaction, Springer Berlin Heidelberg, Berlin, Heidelberg, 2017, pp. 330–346.

[20] Y. Hirai, An intuitionistic epistemic logic for sequential consistency on shared memory, in: LPAR, Springer Berlin Heidelberg, 2010, pp. 272–289. doi:10.1007/978-3-642-17511-4_16.

[21] Y. Moses, Relating knowledge and coordinated action: The knowledge of preconditions principle, in: TARK, EPTCS, 2015, pp. 231–245. doi:10.4204/EPTCS.215.17.

[22] C. Dégremont, B. Löwe, A. Witzel, The synchronicity of dynamic epistemic logic, in: TARK XIII, ACM, 2011, pp. 145–152. doi:10.1145/2000378.2000395.

[23] S. Knight, B. Maubert, F. Schwarzentruber, Reasoning about knowledge and messages in asynchronous multi-agent systems, Mathematical Structures in Computer Science (2017) 1–42doi:10.1017/S0960129517000214.

[24] P. Balbiani, H. van Ditmarsch, S. Fernandez Gonzalez, Asynchronous Announcements, arXiv e-prints arXiv:1705.03392.

[25] J. Havlicek, Computable obstructions to wait-free computability, Distributed Computing 13 (2) (2000) 59–83. doi:10.1007/s004460050068.

[26] A. Castañeda, S. Rajsbaum, M. Raynal, The renaming problem in shared memory systems: An introduction, Computer Science Review 5 (3) (2011) 229–251. doi:10.1016/j.cosrev.2011.04.001.
URL https://doi.org/10.1016/j.cosrev.2011.04.001

[27] A. Lomuscio, M. Ryan, On the relation between interpreted systems and kripke models, in: W. Wobcke, M. Pagnucco, C. Zhang (Eds.), Agents and Multi-Agent Systems Formalisms, Methodologies, and Applications, Springer Berlin Heidelberg, Berlin, Heidelberg, 1998, pp. 46–59.

[28] T. Bolander, H. van Ditmarsch, A. Herzig, E. Lorini, P. Pardo, F. Schwarzentruber, Announcements to attentive agents, Journal of Logic, Language and Information 25 (1) (2016) 1–35.

[29] S. Rajsbaum, Iterated shared memory models, in: LATIN, Vol. 6034 of LNCS, Springer, 2010, pp. 407–416. doi:10.1007/978-3-642-12200-2_36.

[30] F. Benavides, S. Rajsbaum, Collapsibility of read/write models using discrete morse theory, Journal of Applied and Computational Topology (2018) 1–32doi:10.1007/s41468-018-0011-7.

[31] D. N. Kozlov, Chromatic subdivision of a simplicial complex, Homology Homotopy Appl. 14 (2) (2012) 197–209. doi:10.4310/HHA.2012.v14.n2.a12.

[32] D. Kozlov, Combinatorial Algebraic Topology, Springer, 2007. doi:10.1007/978-3-540-71962-5.

[33] J. A. Bondy, U. S. R. Murty, Graph Theory with Applications, 5th Edition, North Holland, 1982.

[34] O. Biran, S. Moran, S. Zaks, A Combinatorial Characterization of the Distributed 1-Solvable Tasks, J. Algorithms 11 (3) (1990) 420–440. doi:10.1016/0196-6774(90)90020-F.

[35] M. P. Herlihy, Impossibility and universality results for wait-free synchronization, in: PODC, ACM, 1988, pp. 276–290. doi:10.1145/62546.62593.

[36] M. Fischer, N. A. Lynch, M. S. Paterson, Impossibility Of Distributed Commit With One Faulty Process, Journal of the ACM 32 (2) (1985) 374–382. doi:10.1145/3149.214121.

[37] É. Goubault, M. Lazić, J. Ledent, S. Rajsbaum, A dynamic epistemic logic analysis of the equality negation task, in: Dynamic Logic. New Trends and Applications - Second International Workshop, DaLí 2019, Proceedings, 2019, pp. 53–70. doi:10.1007/978-3-030-38808-9_4.

[38] J. van Benthem, J. van Eijck, B. Kooi, Logics of communication and change, Information and Computation 204(11) (2006) 1620–1662.

[39] H. van Ditmarsch, B. Kooi, Semantic results for ontic and epistemic change, in: Proc. of 7th LOFT, Texts in Logic and Games 3, Amsterdam University Press, 2008, pp. 87–117.

[40] H. van Ditmarsch, E. Goubault, J. Ledent, S. Rajsbaum, Knowledge and simplicial complexes, arXiv e-prints arXiv:2002.08863.

[41] E. Goubault, S. Rajsbaum, A simplicial complex model of dynamic epistemic logic for fault-tolerant distributed computing, Tech. rep., arXiv:1703.11005 (2017).

[42] E. Goubault, S. Rajsbaum, Models of fault-tolerant distributed computation via dynamic epistemic logic, Tech. rep., arXiv:1704.07883 (2017).

[43] J. Ledent, Geometric semantics for asynchronous computability, Ph.D. thesis, École Polytechnique, Palaiseau, France (2019).

[44] E. Gafni, E. Koutsoupias, Three-processor tasks are undecidable, SIAM J. Comput. 28 (3) (1999) 970–983. doi:10.1137/S0097539796305766.

[45] M. Herlihy, S. Rajsbaum, The decidability of distributed decision tasks, in: STOC, ACM, 1997, pp. 589–598. doi:10.1145/258533.258652.

[46] M. Herlihy, S. Rajsbaum, Simulations and reductions for colorless tasks, in: PODC, ACM, 2012, pp. 253–260. doi:10.1145/2332432.2332483.