

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

Éric Goubault

Jérémy Ledent

Sergio Rajsbaum

LIX, École Polytechnique
Palaiseau, France

Instituto de matemáticas, UNAM
Ciudad Universitaria Mexico 04510, Mexico

{goubault, jeremy.ledent}@lix.polytechnique.fr

sergio.rajsbaum@gmail.com

The usual epistemic $\mathbf{S5}_n$ 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.

1 Introduction


The usual epistemic logic 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. A Kripke $\mathbf{S5}_n$ model represents 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 these simplicial models are very closely related to the usual Kripke models: 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. (a slightly modified version of) the logic $\mathbf{S5}_n$.

To explain the interest of this duality, we extend it to a dynamic setting. We found that in this context, a very natural setting is dynamic epistemic logic (DEL) [4, 10] 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 $\mathcal{I}[\mathcal{A}]$.

We apply our framework to study fault-tolerant distributed computability, because its intimate relation to topology is well understood [16]. 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. We define a particular action model of interest, the *immediate snapshot* action model, which is well known in distributed computing because it fully preserves the topology of the initial complex. This model corresponds to wait-free asynchronous processes operating on a shared memory, which means that the

processes run at an arbitrary speed, independent from the others, and are not allowed to wait for events to happen in other processes.

Another goal is to show how DEL can be used to specify a distributed *task*. A task is the equivalent of a function in distributed computability [2]. 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 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 δ that makes the above diagram commute (Definition 2). This intuitively happens when there is sufficient knowledge in $\mathcal{I}[\mathcal{A}]$ to solve the task.



Beyond the applications that we provide in this paper, our main goal is to construct a general framework that connects epistemic logic and distributed computability. 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 [16] 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.

We concentrate on the specific setting of asynchronous wait-free shared read/write memory. However, there are known equivalences between task solvability in our model and other shared memory and message passing models, and this 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 [16]. Nevertheless, this is far from telling the whole story. In the Conclusions section we discuss many interesting avenues that remain to be explored. And additional technical details appear in the companion Technical Reports [13, 14].

Related work. Work on knowledge and distributed systems is of course one of the inspirations of the present work [26], especially where connectivity [7, 8] is used. But the authors know of no previous work using DEL [4, 10] to study such systems, and neither on directly connecting the combinatorial topological setting of [16] with Kripke models. 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 [24], 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 [24] implies that common knowledge is a necessary condition for performing simultaneous actions. Our formulation of carrier maps as products

has been partially observed in [15]. There are other (categorical) connections between Kripke frames and geometry [25].

DEL is often thought of as inherently being capable of modeling only agents that are synchronous, but as discussed in [9], this is not the case. More recently, [21] proposes a variant of public announcement logic for asynchronous systems that introduces two different modal operators for sending and receiving messages. As we show here, 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.

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 with DEL and distributed computing will be developed in the next sections.

Syntax. Let AP be a countable set of propositional variables and A a finite set of agents. The language \mathcal{L}_K is generated by the following BNF grammar:

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

In the following, we work with $n + 1$ agents, and write $A = \{a_0, \dots, a_n\}$.

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 $|X| - 1$. 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 $\chi : \mathcal{V}(C) \rightarrow A$, 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* $f : C \rightarrow D$ 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. 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 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 $F : \mathcal{S}_A \rightarrow \mathcal{K}_A$ and $G : \mathcal{K}_A \rightarrow \mathcal{S}_A$ 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 generated by the relations $X \sim_a Y$ (for X and Y facets of C) if $a \in \chi(X \cap Y)$.

For a morphism $f : C \rightarrow D$ in \mathcal{S}_A , we define $F(f) : F(C) \rightarrow F(D)$ 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 $A = \{a_0, \dots, a_n\}$. Intuitively, what we want to do is take one n -simplex $\{v_0^s, \dots, v_n^s\}$ for each $s \in S$, and glue them together according to the indistinguishability 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 $\{[v_0^s], \dots, [v_n^s]\}$ 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 $\{[v_0^s], \dots, [v_n^s]\}$ for $s \in S$, since the Kripke frame M is proper, we cannot equate two facets together.

Now let $f : M \rightarrow N$ be a morphism in \mathcal{K}_A . We define $G(f) : G(M) \rightarrow G(N)$ 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 $\{[v_0^s], \dots, [v_n^s]\}$ (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 χ 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. $X = \{[v_0^s], \dots, [v_n^s]\}$ and $Y = \{[v_0^t], \dots, [v_n^t]\}$ 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 Kripke frames.

Consider now a pure chromatic simplicial complex $C \in \mathcal{S}_A$. It is easily seen that $GF(C)$ is isomorphic, as a pure chromatic simplicial complex, to C , hence \mathcal{S}_A and \mathcal{K}_A are equivalent categories. \square

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.

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 $\ell : \mathcal{V}(C) \rightarrow \mathcal{P}(AP)$ 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 $X = \{v_0, \dots, v_n\} \in C$, we write $\ell(X) = \bigcup_{i=0}^n \ell(v_i)$. A *morphism* of simplicial models $f : M \rightarrow M'$ is a chromatic simplicial map that preserves the labeling: $\ell'(f(v)) = \ell(v)$ (and χ). 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 $L : S \rightarrow \mathcal{P}(AP)$. 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* $f : M \rightarrow M'$ 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 $F : \mathcal{SM} \rightarrow \mathcal{KM}$ and $G : \mathcal{KM} \rightarrow \mathcal{SM}$. 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 $\{v_0^s, \dots, v_n^s\}$, 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 figure below shows the so-called binary input complex and its associated Kripke model, for 2 and 3 agents. Each agent gets a binary value 0 or 1, but doesn't know which value has been received by the other agents. So, 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 all the edges labeled by only one agent.


In the simplicial model, agents are represented as colors (black, grey, and white). The labeling ℓ 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 [16] that the binary input simplicial complex for $n + 1$ agents is a n -dimensional 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. 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 truth of a formula φ in some epistemic state (M, X) with $M = \langle C, \chi, \ell \rangle$ a simplicial model, $X \in \mathcal{F}(C)$ a facet of C and $\varphi \in \mathcal{L}_K(A, AP)$. The satisfaction relation, determining when a formula is true in an epistemic state, is defined as:

$$\begin{aligned}
 M, X \models p & \quad \text{iff} \quad p \in \ell(X) \\
 M, X \models \neg\varphi & \quad \text{iff} \quad M, X \not\models \varphi \\
 M, X \models \varphi \wedge \psi & \quad \text{iff} \quad M, X \models \varphi \text{ and } M, X \models \psi \\
 M, X \models K_a \varphi & \quad \text{iff} \quad \text{for all } Y \in \mathcal{F}(C), a \in \chi(X \cap Y) \text{ implies } M, Y \models \varphi
 \end{aligned}$$

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 $\{v_0^s, \dots, v_n^s\}$ of $G(N)$.

Proof. This is straightforward by induction on the formula φ . \square

It is well-known that the axiom system $\mathbf{S5}_n$ is sound and complete with respect to the class of Kripke models [10]. 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:

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

Corollary 1. *The axiom system $\mathbf{S5}_n + \mathbf{Loc}$ is sound and complete w.r.t. the class of simplicial models.*

Proof. Adapting the proof of [10] for $\mathbf{S5}_n$, it can be shown that $\mathbf{S5}_n + \mathbf{Loc}$ is sound and complete w.r.t. the class of local proper Kripke models. Then, we transpose it to simplicial models using Proposition 1.

Indeed, suppose a formula φ 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 φ is true in every simplicial model. Similarly, the converse also holds. \square

The following Theorem shows 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 $f : M \rightarrow M'$. Let $X \in \mathcal{F}(C)$ be a facet of M , a an agent, and φ a formula 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 φ . First, for p an atomic proposition, since morphisms preserves the valuation ℓ , 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 φ 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. In fact, Theorem 3 is still true if the formula φ contains other knowledge operators such as group and common knowledge. For B a subgroup of agents, group knowledge is defined as $E_B \varphi = \bigwedge_{b \in B} K_b \varphi$ and common knowledge for group B is, semantically, the least solution of the equation $C_B \varphi = \varphi \wedge E_B(C_B \varphi)$.

3 DEL via simplicial complexes

We describe here our adaptation of Dynamic Epistemic Logic (DEL) to simplicial models, and an action model that is fundamental in distributed computing.

3.1 DEL basic notions

DEL is the study of modal logics of model change [4, 10]. 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. We first recall the usual notion of action model; then describe a dual version, appropriate to represent epistemic change in simplicial models.

Dynamic Epistemic Logic. An *action model* is a structure $\langle T, \sim, \text{pre} \rangle$, where T is a domain of *action points*, such that for each $a \in A$, \sim_a is an equivalence relation on T , and $\text{pre} : T \rightarrow \mathcal{L}_{\mathcal{K}}$ is a function that assigns a *precondition* $\text{pre}(t)$ 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 $\mathcal{A} = \langle T, \sim, \text{pre} \rangle$ 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 $\text{pre}(t)$ 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 2. *Let M be a local proper Kripke model and $\mathcal{A} = \langle T, \sim, \text{pre} \rangle$ 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 DEL. To work in the category of simplicial models, we consider a simplicial version of action models. First, let us define cartesian products. 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 $X \times Y = \{(u_0, v_0), \dots, (u_k, v_k)\}$ where $X = \{u_0, \dots, u_k\} \in C$, $Y = \{v_0, \dots, v_k\} \in T$ and $\chi(u_i) = \chi(v_i)$.

A *simplicial action model*, $\langle T, \chi, \text{pre} \rangle$ consists of a pure chromatic simplicial complex $\langle T, \chi \rangle$, where the facets $\mathcal{F}(T)$ represent communicative *actions*, and pre assigns to each facet $X \in \mathcal{F}(T)$ a precondition formula $\text{pre}(X)$ in $\mathcal{L}_{\mathcal{K}}$. Let $M = \langle C, \chi, \ell \rangle$ be a simplicial model, and $\mathcal{A} = \langle T, \chi, \text{pre} \rangle$ 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 sub-complex of the cartesian product $C \times T$, induced by all the facets of the form $X \times Y$ such that $\text{pre}(Y)$ holds in X , i.e., $M, X \models \text{pre}(Y)$. The valuation $\ell : \mathcal{V}(C[\mathcal{A}]) \rightarrow \mathcal{P}(AP)$ 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. The simplicial version of the product update model agrees with the usual one on Kripke models:

Proposition 3. *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. \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 snapshots 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. [1, 16] but for the point we want to make about using DEL, the main issues can be studied with this very simple action model.

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 value. 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. Based on this information, each process decides an output value. 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 define instead an action model \mathcal{IS} corresponding to a subset of all the executions in the previous situation. And we do so without loss of generality, because from the task computability perspective, they are known to be equivalent [1].

The interleavings we consider can be represented by a sequence of concurrency classes, c_1, c_2, \dots, c_m . 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.


Let us define formally the simplicial action model corresponding to such immediate snapshot schedules. A *sequential partition* of agents A is a sequence $c = c_1, c_2, \dots, c_m$, 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 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. Define $\text{Aview}_a(c)$ ('A' stands for "agent" view) to be the set of agents whose inputs are seen by a in c : if $a \in c_j$, $\text{Aview}_a(c) = \bigcup_{i \leq j} c_i$. Notice that two executions of the immediate snapshot are indistinguishable by a when the corresponding sequential partitions yield the same Aview for a , and additionally, the agents in Aview have the same inputs.

Consider for instance 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 $b_0b_1b_2$, corresponding to the three values of b, g, w , in that order.

In the *immediate snapshot* simplicial action model $\mathcal{IS} = \langle T, \chi, \text{pre} \rangle$ for three agents $A = \{b, g, w\}$, each action in T is associated with a sequential partition of A . Furthermore, there is one copy of each sequential partition c for each facet $X \in \mathcal{F}(C)$ of model M . Thus, an action of T is given by the data c, b_0, b_1, b_2 , which we write $c^{b_0b_1b_2}$. The precondition of the action $c^{b_0b_1b_2}$ is true precisely in the facet $b_0b_1b_2$ of M (it is a conjunction of atomic propositions).

Consider an action $c^{b_0b_1b_2}$, where $c = c_1, \dots, c_m$. Then $c^{b_0b_1b_2}$ is interpreted as follows. If an agent a is in c_j then a learns the values (of facet $b_0b_1b_2$) of all agents in c_i for $i \leq j$, and only those values. We write $\text{view}_a(c^{b_0b_1b_2})$ the vector of values that a learned, namely, the vector obtained from $b_0b_1b_2$ by replacing a value b_i by \emptyset for agents not in $A\text{view}_a(c)$. Formally, the chromatic simplicial complex $\langle T, \chi \rangle$, consists of all facets of the form:

$$\{ \langle b, \text{view}_b(c^{b_0b_1b_2}) \rangle, \langle g, \text{view}_g(c^{b_0b_1b_2}) \rangle, \langle w, \text{view}_w(c^{b_0b_1b_2}) \rangle \}$$



The figure above illustrates (part of) the action model \mathcal{IS} . It consists of the subdivisions of two triangles; the green copy above has one triangle for each sequential partition (and has four sequential partitions depicted). Similarly, the yellow subdivided triangle below repeats again all sequential partitions, but for a different precondition. The precondition for all facets in the subdivided triangle above is 000, while for the facets of the subdivided triangle below it is 100. The subdivision on top has four facets identified, X, Y, Z, W , each one corresponding to one of the four types of sequential partitions of $A = \{b, w, g\}$, along with the corresponding $A\text{views}$ shown in bubbles. The colors black, grey, white of the vertices correspond respectively to agents b, g, w . Notice that, for example, neither b nor w distinguish between actions Y and Z , and indeed, their views are equal in Y and Z : the view of b consists of itself and the view of w consists of the three inputs. The numbers on the subdivided triangles indicate the views. In the corners, an agent does not learn the input of any other agent. In the boundary, two agents learn each other's inputs, and in the center, all three learn each other's inputs. 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.

An 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*, that has been thoroughly studied. It is clear from the figure that C is a subdivision, but for an arbitrary number of agents, the proof is not simple [16, 23]. It has been shown to have several other topological properties, such as being collapsible [5]. But in fact, for many applications such as consensus and set agreement, it is sufficient to observe the following (see ch.9 of [16]).

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

For a detailed proof see [1]. To complete the example, notice that the effect of applying the action model \mathcal{IS} to the model M of Example 2, which consists of a triangulated sphere, is to subdivide each of the triangles in the sphere. Remarkably, the topology of the initial simplicial complex is preserved.

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 [1, 27], where each agent executes r consecutive immediate snapshots. Each component is a chromatic subdivision, where every triangle is subdivided r times.

4 A DEL semantics for distributed task computability

4.1 Tasks

Consider the situation where a set of agents A starts in an initial global state, defined by 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 [6]. Here we provide, for the first time, a DEL semantics for tasks.

Consider a simplicial model $\mathcal{I} = \langle I, \chi, \ell \rangle$ called the *initial simplicial model*. Each facet of I , with its labeling ℓ , 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 $\mathcal{T} = \langle T, \chi, \text{pre} \rangle$ 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 $\text{pre}(X)$ holds, and every agent $a \in A$ decides the value d_a , then this is a valid execution”.

The most important task in distributed computing is *binary consensus*, where the agents must agree on a value 0 or 1, such that at least one agent had the agreed value as input. Thus, T has only two facets, X_0 where all decisions are 0 and X_1 , where all decisions are 1. $\text{pre}(X_0)$ is true in all facets of \mathcal{I} , except for the one where all agents start with input 1. Similarly, $\text{pre}(X_1)$ is true in all facets of \mathcal{I} , except for the one where all agents start with input 0. The following generalization of consensus has been well studied in distributed computability [16]. In the *k-set agreement* task, agents start with inputs from a set of at least $k + 1$ values, and have to decide on at most k different inputs.

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 δ of simplicial models is a chromatic simplicial map that preserves the labeling: $\ell'(f(v)) = \ell(v)$. 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 $\delta : \mathcal{I}[\mathcal{A}] \rightarrow \mathcal{I}[\mathcal{T}]$ such that $\pi_{\mathcal{I}} \circ \delta = \pi_{\mathcal{I}}$, i.e., the diagram of simplicial complexes below commutes.

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 δ 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, $\text{pre}(dec)$ 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 $\chi(v) = a \in \mathcal{A}$. 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 the same value in both situations.



The diagram above has two illuminating interpretations. First, by Theorem 3, we know that the knowledge about the world of each agent can only decrease (or stay constant) along the δ arrow. So agents should improve knowledge through communication, by going from \mathcal{I} to $\mathcal{I}[\mathcal{A}]$. The task is solvable if and only if there is enough knowledge in $\mathcal{I}[\mathcal{A}]$ 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. This leads us to the connection with distributed computability described in this extended abstract; further details are in [13, 14].

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, set agreement, and approximate agreement. Their solvability is already well-understood; our aim here is to understand the epistemic logic content of the known topological arguments that are used to prove unsolvability.

Consensus Let $\mathcal{I} = \langle I, \chi, \ell \rangle$ be the initial simplicial model for binary input values, and $\mathcal{T} = \langle T, \chi, \text{pre} \rangle$ be the action model for binary consensus. 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 $\delta : \mathcal{I}[\mathcal{A}] \rightarrow \mathcal{I}[\mathcal{T}]$ of Definition 2 does not exist. The usual proof of impossibility uses a topological obstruction to the existence of δ . Here, instead, we exhibit a logical obstruction.

Theorem 4. *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 φ_i be a formula denoting that at least one agent has input i . We claim that, for $i \in \{0, 1\}$, at any facet Y of $I_i \times X_i$, there is common knowledge that at least one input is i , $\mathcal{I}[\mathcal{T}], Y \models C_A \varphi_i$.

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, and hence there is a path from the facet with valuation indicating that all inputs are 0 to the facet where all inputs are 1. Namely, we claim that at any facet X of $\mathcal{I}[\mathcal{IS}]$, it is not the case that $\mathcal{I}[\mathcal{IS}], X \models C_A \varphi_i$, 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 hence, there cannot be a morphism δ from $\mathcal{I}[\mathcal{IS}]$ to $\mathcal{I}[\mathcal{T}]$, by the two previous claims. \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 [19], and even if only one may crash in a message passing system [11], which are the classic consensus impossibility results. Secondly, the usual topological argument for impossibility is the following: because simplicial maps preserve connectivity, δ 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.

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 $\mathcal{T} = \langle T, \chi, \text{pre} \rangle$ be the action model for 2-set agreement, requiring that each agent decides on one of the input values, and at most 2 different values are decided. Thus, T has facets X_{d_0, d_1, d_2} , for each vector d_0, d_1, d_2 , such that $d_i \in \{0, 1, 2\}$, $|\{d_0, d_1, d_2\}| \leq 2$, and $\text{pre}(X_{d_0, d_1, d_2}) = \varphi_{d_0} \wedge \varphi_{d_1} \wedge \varphi_{d_2}$, where φ_i is as above.

Theorem 5. *The 2-set agreement task is not solvable by \mathcal{IS} .*

Proof (Sketch). The usual topological argument [16] roughly goes as follows. We can visualize the complex $\mathcal{I}[\mathcal{T}]$ as having the structure of a triangle with a hole in the middle. The three “corners” of the triangle, indexed by $i \in \{0, 1, 2\}$, consist of just one facet of the form $I_i \times X_i$, where I_i is the input facet with only one input, i , and in X_i all decisions are i . The three “edges” of the triangle are of the form $I_{ij} \times X_{ij}$, where I_{ij} consists of all facets with inputs in $\{i, j\}$, and X_{ij} the facets whose decisions are in $\{i, j\}$. Notice that $I_{ij} \times X_{ij}$ contains $I_i \times X_i$ and $I_j \times X_j$. This triangle must have a hole in the middle: otherwise, by Sperner’s lemma (see e.g. [22, 16]), there would be a facet with three distinct decision values. Thus, $\mathcal{I}[\mathcal{T}]$ is not 2-connected. But since $\mathcal{I}[\mathcal{IS}]$ is 2-connected, and simplicial maps preserve 2-connectivity, there cannot exist a suitable δ . \square

Finding a logical obstruction to the existence of δ is an open question. This would amount to finding a formula φ which is true $\mathcal{I}[\mathcal{T}]$, but false $\mathcal{I}[\mathcal{IS}]$, and applying Theorem 3. Doing so, we would understand better what knowledge is necessary to solve set agreement, which is not available in $\mathcal{I}[\mathcal{IS}]$.

Approximate agreement We discuss now the approximate agreement task, where agents have to decide values which are $1/N$ apart from each other. Its solvability depends on the number of immediate snapshot communication rounds r that the agents perform. We did not describe in detail the action model \mathcal{IS}^r , so we only briefly mention that the task is solvable if and only if the number of rounds is large

enough, with respect to N . Very roughly, there is a center facet X^c in the product update of the task, where $\mathcal{I}[\mathcal{T}], X^c \models E^k \phi_c$, where k is roughly $N/2$, for a formula ϕ_c representing the input values in X^c . On the other hand, there is no facet X in $\mathcal{I}[\mathcal{LS}']$ where this knowledge exists, unless r is large enough. A detailed proof is in the technical report [13].

5 Conclusions

We have 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. Also, we have defined problem specifications based on DEL using simplicial complexes, instead of based on formula specifications. 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{LS} , that corresponds to a well-studied situation in distributed computing, but many other models can be treated similarly.

Many interesting questions are left for future work. 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 failures within DEL itself. In two preliminary reports we give additional details, and explore further some of these issues [13, 14]. It is known to be undecidable whether a task is solvable in the immediate snapshot model, even for three processes [12, 17], and hence the connection we establish with DEL implies that it is undecidable if certain knowledge has been gained in multi-round immediate snapshot action models, but further work is needed to understand this issue. Future work is needed to study bisimulations and their relation to simulations studied in task computability [18]. 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.

References

- [1] H. Attiya & S. Rajsbaum (2002): *The Combinatorial Structure of Wait-Free Solvable Tasks*. *SIAM J. Comput.* 31(4), pp. 1286–1313, doi:10.1137/S0097539797330689.
- [2] H. Attiya & J. Welch (2004): *Distributed Computing: Fundamentals, Simulations, and Advanced Topics*, 2 edition. Wiley, doi:10.1002/0471478210.
- [3] A. Baltag, L.S. Moss & S. Solecki (1998): *The logic of common knowledge, public announcements, and private suspicions*. In: *TARK VII*, pp. 43–56, doi:10.1007/978-3-319-20451-2_38.
- [4] A. Baltag & B. Renne (2016): *Dynamic Epistemic Logic*. In: *The Stanford Encyclopedia of Philosophy*, see <https://plato.stanford.edu/archives/win2016/entries/dynamic-epistemic/>, Metaphysics Research Lab, Stanford University.
- [5] F. Benavides & S. Rajsbaum (2018): *Collapsibility of read/write models using discrete morse theory*. *Journal of Applied and Computational Topology*, pp. 1–32, doi:10.1007/s41468-018-0011-7.
- [6] O. Biran, S. Moran & S. Zaks (1990): *A Combinatorial Characterization of the Distributed 1-Solvable Tasks*. *J. Algorithms* 11(3), pp. 420–440, doi:10.1016/0196-6774(90)90020-F.
- [7] A. Castañeda, Y. A. Gonczarowski & Y. Moses (2016): *Unbeatable Set Consensus via Topological and Combinatorial Reasoning*. In: *PODC*, ACM, pp. 107–116, doi:10.1145/2933057.2933120.
- [8] A. Castañeda, Y. A. Gonczarowski & Y. Moses (2014): *Unbeatable Consensus*. In: *DISC, LNCS 8784*, Springer, pp. 91–106, doi:10.1007/978-3-662-45174-8_7.

- [9] C. Dégremon, B. Löwe & A. Witzel (2011): *The Synchronicity of Dynamic Epistemic Logic*. In: *TARK XIII*, ACM, pp. 145–152, doi:10.1145/2000378.2000395.
- [10] H. van Ditmarsch, W. van der Hoek & B. Kooi (2007): *Dynamic Epistemic Logic*. Springer, doi:10.1007/978-1-4020-5839-4.
- [11] M. Fischer, N. A. Lynch & M. S. Paterson (1985): *Impossibility Of Distributed Commit With One Faulty Process*. *Journal of the ACM* 32(2), pp. 374–382, doi:10.1145/3149.214121.
- [12] E. Gafni & E. Koutsoupias (1999): *Three-Processor Tasks Are Undecidable*. *SIAM J. Comput.* 28(3), pp. 970–983, doi:10.1137/S0097539796305766.
- [13] E. Goubault & S. Rajsbaum (2017): *A simplicial complex model of dynamic epistemic logic for fault-tolerant distributed computing*. Technical Report, arXiv:1703.11005.
- [14] E. Goubault & S. Rajsbaum (2017): *Models of fault-tolerant distributed computation via dynamic epistemic logic*. Technical Report, arXiv:1704.07883.
- [15] J. Havlicek (2000): *Computable obstructions to wait-free computability*. *Distributed Computing* 13(2), pp. 59–83, doi:10.1007/s004460050068.
- [16] M. Herlihy, D. Kozlov & S. Rajsbaum (2013): *Distributed Computing Through Combinatorial Topology*. Elsevier-Morgan Kaufmann, doi:10.1016/C2011-0-07032-1.
- [17] M. Herlihy & S. Rajsbaum (1997): *The Decidability of Distributed Decision Tasks*. In: *STOC*, ACM, pp. 589–598, doi:10.1145/258533.258652.
- [18] M. Herlihy & S. Rajsbaum (2012): *Simulations and reductions for colorless tasks*. In: *PODC*, ACM, pp. 253–260, doi:10.1145/2332432.2332483.
- [19] M. P. Herlihy (1988): *Impossibility and Universality Results for Wait-free Synchronization*. In: *PODC*, ACM, pp. 276–290, doi:10.1145/62546.62593.
- [20] Y. Hirai (2010): *An Intuitionistic Epistemic Logic for Sequential Consistency on Shared Memory*. In: *LPAR*, Springer Berlin Heidelberg, pp. 272–289.
- [21] S. Knight, B. Maubert & F. Schwarzentruber (2017): *Reasoning about knowledge and messages in asynchronous multi-agent systems*. *Mathematical Structures in Computer Science*, pp. 1–42, doi:10.1017/S0960129517000214.
- [22] D. Kozlov (2007): *Combinatorial Algebraic Topology*. Springer, doi:10.1007/978-3-540-71962-5.
- [23] D. N. Kozlov (2012): *Chromatic subdivision of a simplicial complex*. *Homology Homotopy Appl.* 14(2), pp. 197–209, doi:10.4310/HHA.2012.v14.n2.a12.
- [24] Yoram Moses (2015): *Relating Knowledge and Coordinated Action: The Knowledge of Preconditions Principle*. In: *TARK, EPTCS*, pp. 231–245, doi:10.4204/EPTCS.215.17.
- [25] T. Porter (2004): *Interpreted systems and Kripke models for multiagent systems from a categorical perspective*. *Theoretical Computer Science* 323(1), pp. 235 – 266, doi:10.1016/j.tcs.2004.04.005.
- [26] Y. Moses R. Fagin, J. Halpern & M. Vardi (1995): *Reasoning About Knowledge*. MIT Press.
- [27] S. Rajsbaum (2010): *Iterated Shared Memory Models*. In: *LATIN, LNCS 6034*, Springer, pp. 407–416, doi:10.1007/978-3-642-12200-2_36.