Defining Logical Obstruction with Fixpoints in Epistemic Logic*

Susumu Nishimura

susumu@math.kyoto-u.ac.jp

Department of Mathematics, Kyoto University

Abstract

The logical method proposed by Goubault, Ledent, and Rajsbaum provides a means of demonstrating the unsolvability of distributed tasks within the epistemic logic framework. To show that a task is unsolvable, we need to find a logical obstruction, which is an epistemic logic formula describing the reason for the unsolvability, or more precisely, the incompatibility between the task, which is a model of what is to be solved, and the protocol, which is a model of what the distributed system can compute.

To date, only a few concrete instances of logical obstructions have been devised. In particular, existing proposals of logical obstruction to the k-set agreement task are unsatisfactory because they work only for the case $k = 1$ or the protocol is restricted to single-round execution. This is because the unsolvability of the k-set agreement task is tied with the higher-dimensional property of the corresponding combinatorial topological model, while the language of epistemic logic has a limited ability to express it. This study proposes the use of an epistemic μ-calculus variant, which extends epistemic logic with distributed knowledge modalities and propositional greatest fixpoints. With these extensions, we can define an epistemic formula whose epistemic content contradicts a property regarding the higher-dimensional connectivity, which is indicated in the proof of Sperner's lemma. This formula thus works as a logical obstruction, showing that the k-set agreement task is unsolvable by the multiple-round immediate snapshot protocol. Further, we show that the same formula works as a logical obstruction for the k-concurrency, which is a protocol of a limited degree of concurrency.

keywords: Sperner's lemma, Epistemic μ-calculus, Task unsolvability, Theory of distributed computing

1 Introduction

In a distributed environment, where processes are subject to failure, certain simple computing problems cannot be solved effectively. For example, the consensus problem is unsolvable in an asynchronous distributed system: it has been shown that there is no wait-free distributed algorithm that allows all the non-faulty processes in the system to agree on a single common value [10, 18], where the wait-freedom indicates that each non-faulty process is guaranteed to finish its execution within a finite number of steps.

A distributed computing problem such as the consensus problem is often specified as a *task* defining what outputs can be decided by the processes for each different combination of initial inputs given to them. For example, *k-set agreement*, which generalizes the consensus problem such that the processes may agree on at most k different values, is defined as a task that satisfies the following specification:

Agreement. At most k different values are decided by the processes in the system.

Validity. A value decided by a process must be one of the initial input values given to the processes.

*This is an author's version of the paper published in Journal of Applied and Computational Topology vol. 8, pp. 941–970 (2024). https://doi.org/10.1007/s41468-023-00151-8

A central topic in the study of distributed computing is *task solvability*, i.e., whether a given task is solvable by a *protocol*, which is an algorithm designed for a particular distributed environment. Particularly, the (un)solvability of the consensus and the general k-set agreement tasks has been intensively studied. The immediate snapshot protocol [5] is most significant in this context, because, as for task solvability, the multiple-round, iterated immediate snapshot protocol is a computational equivalent of the asynchronous read-write shared memory model [21]. A protocol is called a *single-round* protocol if it allows each process to communicate with others only once, while a protocol is called a *multiple-round* protocol if it iterates a single-round protocol an arbitrary number of times, allowing multiple inter-process communications.

The solvability of the consensus and the k-set agreement tasks has been answered negatively using the topological method [18]. In the topological method, each individual process is modeled as a vertex representing the local state of the process. The set of $n+1$ concurrent processes of the distributed system is modeled using an n-dimensional *simplex*, i.e., a set of $n+1$ vertexes. The collection of possible output states produced by nondeterministic execution of the distributed system is modeled by an n-dimensional *complex*, i.e., a set of n-dimensional simplexes (and their lower dimensional faces). In this simplicial complex model, task unsolvability is demonstrated by showing that the protocol complex — which models nondeterministic outputs produced by the protocol — is not topologically compatible with the task complex, which models the possible outputs allowed by the task. The topological method, built on the solid theoretical foundations of combinatorial topology, provides a powerful means of proving the unsolvability, which is difficult to replace by other methods [1].

We may alternatively reason about the properties of a distributed system using epistemic logic [22, 34], by regarding it as a knowledge system whose ability of solving tasks is governed by the information that each process knows about the others. In this study, we show the unsolvability of the k-set agreement tasks, using the recently emerging logical method proposed by [13], where epistemic logic is employed to describe the topological structure of simplicial complexes. The central observation in their development is that simplicial complexes are isomorphic to Kripke models: the facets of a simplicial complex correspond to the epistemic states (or possible worlds) of a Kripke model and the geometric adjacency of facets is interpreted by the indistinguishability relations over the epistemic states. This indicates that geometric information contained in a simplicial complex can be translated into a Kripke model that is suitable for reasoning with epistemic logic. The prominent feature of the logical method is that the unsolvability of a distributed task can be entailed from a *logical obstruction*, which is an epistemic logic formula that is valid in the model of the task but is invalid in the model of the protocol. The logical method uses a logical obstruction to describe a topological incompatibility. For example, [13] provided the logical obstruction to the consensus task and N-approximate agreement task, using the common knowledge [9] for expressing the topological incompatibility on path connectivity.

Devising a logical obstruction formula is not a trivial task, however. The logical obstruction to the k-set agreement task was left as an open problem by [13]. The major difficulty is that the solvability of the k-set agreement tasks ($k > 1$) is topologically tied with higher-dimensional connectivity [20], which is more difficult to express in the language of epistemic logic than the 1-dimensional path connectivity for the consensus task. Using the distributed knowledge modality for expressing higher-dimensional adjacency of simplexes, [27] presented a logical obstruction to the k-set agreement. This obstruction was further refined in the study of [36] to encompass a broader class of distributed systems with superset-closed adversaries. However, these logical obstructions are still unsatisfactory, because they do not work for the multiple-round protocol.

In this study, we present a logical obstruction to the k-set agreement task for the multiple-round immediate snapshot protocol. More precisely, we devise a concrete epistemic logic formula to show that the k-set agreement task is not solvable by any number of iterations of the immediate snapshot protocol, for any k that is less than the number of the processes.

Our strategy is to derive such a logical obstruction from Sperner's lemma [31], which is a classical result of combinatorial topology. Sperner's lemma has already been used to provide a concise topological proof for the unsolvability of the k-set agreement tasks [21, 19]. Sperner's lemma is, at least superficially, a claim about the parity of the number of certain facets contained in the subdivision of a simplex. Epistemic logic cannot express this claim, as it has no effective means of enumerating the facets contained in a complex, or equivalently, the epistemic states contained in a Kripke model. Nevertheless, a careful look at the original proof of Sperner's lemma [8] reveals that it can also be understood as a claim about higher-dimensional connectivity of facets. The notion of higher-dimensional connectivity is obtained by replacing 1-dimensional paths by higher-dimensional ones, which are a sequence of facets where each successive pair in the sequence share not necessarily a vertex, but a face of arbitrary dimension.

To express this higher-dimensional connectivity, we extend the language of epistemic logic with the following features.

- Distributed knowledge modality.

  Similar to Nishida's obstruction and its refinement [27, 36], we use distributed knowledge modalities [9, 15]. A distributed knowledge formula $\mathrm{D}_A\,\varphi$, where A is a subset of the processes, states that φ is true at every epistemic state that is not distinguished by any process in A. To interpret it geometrically, φ is true at every facet sharing a common face, whose vertexes represent the subset A of processes.

- Propositional greatest fixpoint.

  We further introduce the propositional greatest fixpoint $\nu Z.\varphi$ into the language of epistemic logic. This gives rise to the so-called *epistemic μ-calculus* [9, 30], an epistemic variant of modal μ-calculi [24, 6]. A greatest fixpoint formula can be combined with distributed knowledge modalities to express an unbounded nesting of modalities $\mathrm{D}_{A_1}(\cdots \mathrm{D}_{A_2}(\cdots \mathrm{D}_{A_m}(\cdots)\cdots)\cdots)$. We use this fixpoint formula to formally express that a certain invariant property holds at every facet along higher-dimensional paths.

- Atomic formulas for output values.

  We further extend the language of epistemic μ-calculus with an additional set of atomic formulas for mentioning the output values decided by the processes. The formal semantics for the additional atomic formulas is given in an extended Kripke model with factual change of atomic propositions [32].

With these extensions to the language of logic, we provide a logical obstruction to the k-set agreement task, which was left as an open problem in [13]. The logical obstruction is given as a greatest fixpoint formula expressing a contradicting property that no facet witnesses $k+1$ or more distinct output values along any higher-dimensional path in the Kripke model of the multiple-round immediate snapshot protocol. In this way, the logical obstruction precisely describes the reason for the unsolvability in an explicit formula of epistemic μ-calculus. This is in contrast to the topological method, where the reason for the unsolvability is implicitly expressed in the informal text of a mathematical proof.

In this study, we also apply the logical method to the k-concurrency protocol [11], which defines a submodel of the 2-iterated immediate snapshot protocol. The topological structure of the k-concurrency protocol is more intricate to analyze formally than that of the immediate snapshot protocol, but the logical method applies rather straightforwardly. This is because the submodel includes those facets that are relevant to the argument of higher-dimensional connectivity and therefore the very same logical obstruction used for the immediate snapshot protocol can be applied. This yields an example that the logical method provides a proof for task unsolvability, for which the topological method would require a more involved topological analysis.

1.1 Related work

A limited number of logical obstructions have been presented to date. [13] presented logical obstructions to the consensus task and N-approximate agreement task, using common knowledge modalities. [27] and [36] devised logical obstructions to the k-set agreement task, using distributed knowledge modalities. However, these were unsatisfactory in that they allowed only the single-round immediate snapshot protocol. This study presents a logical obstruction to the k-set agreement task that works for the multiple-round immediate snapshot protocol, in the language of epistemic μ-calculus. To the best of the author's knowledge, there has been no proposal of such a logical obstruction containing epistemic content that accounts for the reason for the unsolvability.

[32] presented the generic logical obstruction formula, which is a purely propositional formula describing the input/output specification of the task to be solved. This generic logical obstruction, however, does not contain any epistemic content; therefore, it does not indicate any reason for the

unsolvability. The unsolvability proof with the generic logical obstruction must entirely resort to external topological properties.

Epistemic logic has already been used for reasoning about knowledge of agents (processes), e.g., in interpreted systems [16, 9]. The distinguishing feature of the logical method is that it can express the reason for the unsolvability in the language of epistemic logic, as a logical obstruction. The present study reinforces this advantage of the logical method by providing an instance of logical obstruction to the k-set agreement task, which is a central topic in the study of distributed computing.

Recent studies have shown that not every unsolvable task has a corresponding logical obstruction and moreover the existence of logical obstruction may depend on the language of epistemic logic. [32] showed, by applying a bisimulation technique, that no logical obstruction to the equality negation task is definable in an epistemic language that allows knowledge modality but no atomic formulas on output values. [23] applied a similar technique of (one-way) simulation to show that no logical obstruction to 2-set agreement task that works for the multiple-round immediate snapshot protocol is definable in an epistemic language that allows distributed knowledge modalities, but neither fixpoints nor atomic formulas on output values. As we will see in subsequent sections, using the fixpoints and the additional atomic formulas, we can define a logical obstruction for the multiple-round protocol that expresses a graph path of unbounded length that spans over the higher-dimensional structure of a simplicial complex, where the graph path is navigated by the output values on the vertexes of the simplicial complex. In contrast, the logical obstruction for the single-round protocol given in [27, 36] does not require these constructs, as the graph path has a bounded length and can be simply navigated without the help of output values.

In the earliest development of the logical method [13], a logical obstruction is defined using common knowledge as a primary modality that describes path connectivity. The present study employs the general fixpoint construct, rather than specific modalities, such as common knowledge and common distributed knowledge modalities [9, 3], which can be defined using fixpoints. Herein, we require a more delicate definition of logical obstruction that describes connectivity via higher-dimensional paths, in which facets are connected via faces of varying dimensions. Instead of introducing specific modalities tailored for this purpose, we define the logical obstruction with the flexibility of the general fixpoint construct.

In the study of the topological method, k-connectivity has been identified as the critical higher-dimensional topological property from which the unsolvability of set agreement tasks is deduced [18, 20]. The logical obstruction also describes a similar higher-dimensional property, but it is a far less general one that mentions the connectivity of a graph that is specifically defined for a subdivision of a complex. The formal correspondence of this graph connectivity and k-connectivity is left open for future investigation.

Throughout this paper, we use product update models to incorporate the topological definition of distributed tasks and protocols into Kripke models for epistemic reasoning, following the studies of [13, 32]. Recently it has been shown that such Kripke models can be alternatively defined by communication pattern models [35]. Communication pattern models have the advantage that, in modeling multiple-round protocols, each different round of the protocol can be defined by the same communication pattern, whereas product update models must prepare a different product update 'action' for each different round. Despite of this inefficiency, in this study, we stick to product update models because the theory of product update models has a well-developed set of theorems, such as the knowledge gain (Theorem 3.1) and the logical obstruction theorem with factual change (Theorem 4.1).

1.2 Organization

The rest of this paper is organized as follows. Section 2 reviews Sperner's lemma and describes the higher-dimensional connectivity property indicated in its proof. Section 3 introduces an epistemic μ-calculus and its formal semantics in a simplicial model, namely a Kripke model that is derived from a simplicial complex. Section 4 presents the formal combinatorial definition of product update models, by which we argue the unsolvability of the k-set agreement tasks. In Section 5, we provide the logical obstruction as a concrete formula of the epistemic μ-calculus. We also show that the same obstruction applies to the k-concurrency protocol in Section 6. Finally Section 7 concludes this paper.

2 Sperner's Lemma as Higher-Dimensional Connectivity

Sperner's lemma [31], which is a well-known result in combinatorial topology, is primarily a statement on the parity of the number of certain facets contained in a subdivision. Nevertheless, its proof indicates that it can also be interpreted as a statement about higher-dimensional connectivity of the facets. In this section, we informally describe the geometric intuitions behind Sperner's lemma and observe that the proof exploits a higher-dimensional path that spans over facets in the subdivision. In Section 5, we define an epistemic formula that expresses this higher-dimensional path and formally proves that the formula is a logical obstruction to the k-set agreement task.

For the remainder of this paper, we use $[\ell, m]$ to denote the set of nonnegative integers $\{\ell, \ell + 1, \dots, m\}$, ranging from ℓ to m.

2.1 Sperner's lemma

Suppose $\mathit{\Delta}^n$ is an n-dimensional simplex, i.e., a set of $n+1$ vertexes, and also $\mathcal{D}$ is a complex of a subdivision of $\mathit{\Delta}^n$. Let us write $V(\mathit{\Delta}^n)$ and $V(\mathcal{D})$ for the set of vertexes contained in $\mathit{\Delta}^n$ and $\mathcal{D}$, respectively. A (simplicial) complex is a set of simplexes that is closed under set inclusion. A simplex Δ' is called a *face* of another simplex Δ, if $V(\Delta') \subseteq V(\Delta)$. The subdivision $\mathcal{D}$ is a pure complex of dimension n, that is, every *facet* (i.e., maximal simplex) in $\mathcal{D}$ is of dimension n.

Assume $\mathcal{D}$ is a subdivision of $\mathit{\Delta}^n$ and $\eta_0 : V(\mathit{\Delta}^n) \to [0, n]$ is a labeling function on the $n+1$ vertexes of $\mathit{\Delta}^n$ such that $\eta_0(u) = \eta_0(v)$ implies $u = v$. A labeling function $\eta : V(\mathcal{D}) \to [0, n]$ is referred to as a *Sperner labeling* on $\mathcal{D}$, if η extends η_0 such that $\eta(v) \in \eta_0(\mathit{\Delta}_v)$ for each vertex $v \in V(\mathcal{D})$, where $\mathit{\Delta}_v$ is the minimum face of $\mathit{\Delta}^n$ (geometrically) containing the vertex v. This means that for each $v \in V(\mathcal{D})$ contained in the minimum face $\{v_0, \dots, v_m\} \subseteq \mathit{\Delta}^n$, v is allowed to obtain its label only from $\eta_0(v_0), \dots, \eta_0(v_m)$. A facet $X \in \mathcal{D}$ is called *fully-labeled*, if $\eta(X) = [0, n]$, i.e., X is a set of $n+1$ distinctly labeled vertexes. See Fig. 1 for an instance of subdivision with Sperner labeling.

Lemma 2.1 (Sperner's lemma [31])**.** *Let $\eta : V(\mathcal{D}) \to [0, n]$ be a Sperner labeling on a subdivision $\mathcal{D}$ of an n-dimensional simplex $\mathit{\Delta}^n$, whose vertexes are distinctly labeled by $[0, n]$. Then, $\mathcal{D}$ contains a fully-labeled facet; more precisely, it contains an odd number of fully-labeled facets.*

The unsolvability of the k-set agreement task for the asynchronous, wait-free distributed system with $n+1$ processes is a direct consequence of Sperner's lemma. The outputs by such a system are characterized by a protocol complex that is a subdivision of the input simplicial complex of dimension n [4, 21, 18]. The validity property of the k-set agreement task implies a Sperner labeling on the subdivision, but Sperner's lemma contradicts the agreement property, unless $k \geq n+1$: the existence of a fully-labeled facet implies the inevitable possibility that $n+1$ processes make $n+1$ distinct decisions.

2.2 Sperner's lemma as a statement of higher-dimensional connectivity

Sperner's lemma does not directly entail a logical obstruction. This is because epistemic logic, even with the extensions to be introduced in this paper, has no way to effectively enumerate the facets contained in a complex, and hence cannot express the statement of Sperner's lemma.

In this study, we exploit a different interpretation of Sperner's lemma, which can be expressed using a suitable extension in the language of epistemic logic. The proof of Sperner's lemma [8] suggests that the *higher-dimensional connectivity* of the facets of the subdivision derives a particular graph structure.

The proof proceeds by induction on dimension. Let us consider a subdivision of a 2-dimensional simplex labeled by $[0, 2]$, as illustrated in the top left figure of Fig. 1, with a Sperner labeling on the vertexes. From the Sperner labeling, we derive a graph as follows. The nodes of the graph are the facets of the subdivision and a special graph node, which is designated by $\star$ in the figure. We introduce a graph edge for each $\{0, 1\}$-labeled simplex (i.e., a 1-dimensional simplex comprising one vertex labeled with 0 and the other vertex with 1), as indicated by a red line in the figure. Such a $\{0, 1\}$-labeled simplex is either contained in the interior of the original 2-dimensional simplex $\mathit{\Delta}^2$, or in the subdivision of the 1-dimensional face labeled by 0 and 1. In the former case, a pair of facets sharing a $\{0, 1\}$-labeled simplex is connected by a graph edge. In the latter case, the sole facet containing a $\{0, 1\}$-labeled simplex is connected to the special node by a graph edge.

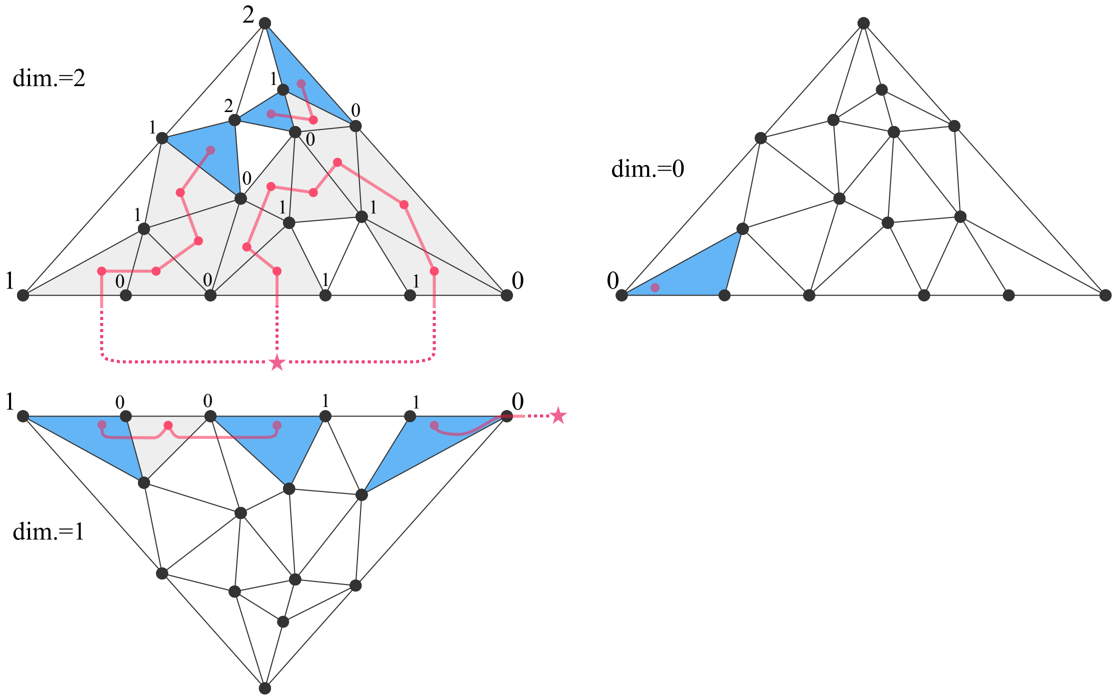

Figure 1: Graphs derived from Sperner labeling, for dimensions 2, 1, and 0

The graph obtained by this particular construction consists of graph nodes of the following three classes of different degrees. (A graph node is of degree m, if it has m graph edges connected to it.) (i) Each fully-labeled facet (shaded in blue in the figure) is a graph node of degree 1; (ii) Each facet whose three vertexes are labeled by either $\{0,0,1\}$ or $\{0,1,1\}$ (shaded in gray) is a graph node of degree 2; (iii) The special graph node has an odd degree, because the subdivision of the 1-dimensional $\{0,1\}$-labeled simplex contains an odd number of $\{0,1\}$-labeled smaller simplexes, by the induction hypothesis for dimension 1. Sperner's lemma is then a consequence of an elementary fact from graph theory that every graph contains an even number of graph nodes of odd degrees.

The general induction case of dimension d ($0 < d \leq n$) is similarly argued. Given a subdivision of a d-dimensional, $[0,d]$-labeled face of the original simplex Δ^n, we derive a graph according to the Sperner labeling as follows: For each $[0,d-1]$-labeled simplex of dimension $d-1$ in the subdivision, we introduce a graph edge to connect a pair of facets sharing the $[0,d-1]$-labeled simplex or to connect the special node to the sole facet containing the $[0,d-1]$-labeled simplex.

The lower left figure of Fig. 1 illustrates the induction case of dimension 1. For the Sperner labeling on the subdivision of the $\{0,1\}$-labeled 1-dimensional face of the original simplex Δ^2, we obtain a graph by introducing a graph edge for each $\{0\}$-labeled (0-dimensional) simplex, i.e., a vertex labeled with 0, such that each graph edge connects a pair of facets sharing the $\{0\}$-labeled simplex; In the case where a $\{0\}$-labeled simplex is contained in a sole single facet of the subdivision (e.g., the rightmost one shaded blue in the figure), we let a graph edge connect the facet with the special node $\star$.

For the induction case of dimension 0, we obtain a trivial graph comprising a single node, as shown in the top right figure of Fig. 1.

Remark 1. *For the induction cases of a lower dimension, a graph node is specified by a facet of the subdivision of the simplex of the lower dimension, but it is represented by a facet of the subdivision of the original simplex Δ^2 that subsumes it as a face. We prefer this representation, since the epistemic states of the Kripke model to be introduced in Section 3.2 are modeled by facets rather than by simplexes of lower dimensions. Special care must be taken here regarding the choice of facet, because a simplex of lower dimension can be shared by many facets. (For instance, in the 0-dimensional case of Fig. 1, there are two different facets that subsume the 0-labeled simplex (the vertex labeled by 0).) In Section 4, we introduce a formal combinatorial notation that uniquely determines such a facet.*

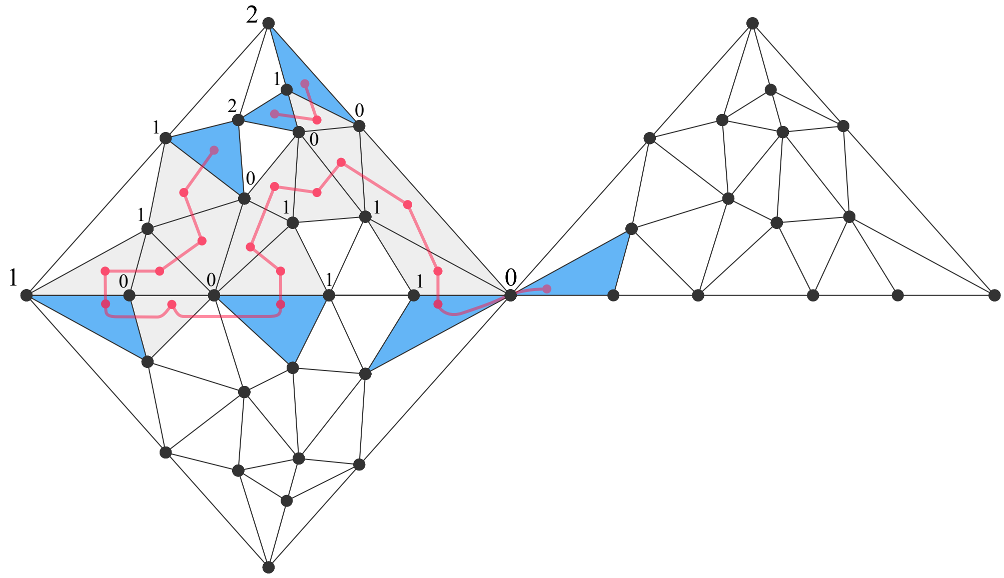

Figure 2: Unified graph of all the induction cases of dimensions 2, 1, and 0

The proof of Sperner's lemma above exploits the higher-dimensional connectivity of facets implied by a Sperner labeling, where a pair of facets can be connected not only by a common vertex, but also by a common face of higher dimension. This higher-dimensional connectivity is better demonstrated by a single unified graph given in Fig. 2, which merges the graphs of different induction cases from Fig. 1 by eliminating the special nodes and filling vacancies with corresponding graph nodes.

The unified graph satisfies the following properties. (i) Each graph node has degree 1 or 2; (ii) A graph node is of degree 1 if and only if it is either a node of a fully-labeled facet or the sole graph node of the trivial graph for the case of dimension 0. From these, we observe that the unified graph contains a cycle-free path that begins from the sole single node (of the graph for dimension 0) and ends at a node of a fully-labeled facet.

In Section 5, we formally define such a graph and demonstrate that the graph has the above-mentioned structure in general. We use this fact to show the unsolvability of the k-set agreement task. Assuming that no fully-labeled facets are contained in a subdivision of a simplicial complex with Sperner labeling, we deduce that the unified graph must have an ever-lasting, cycle-free graph path, which can never happen for a finite graph that is derived from a finite simplicial complex of a subdivision.

3 Logical Method for Distributed Computing

For the remainder of this paper, we assume a distributed system of $n+1$ processes, where each individual process is distinguished by a unique process id taken from the set $\Pi = [0, n]$. We use 'process a' to refer to the process identified by $a \in \Pi$. We use *Value* to denote the set of values that can be assigned to the processes and assume that $\mathit{Value} \supseteq \Pi$.

3.1 Simplicial complex model for distributed computing

In the topological theory of distributed computing [18], distributed systems are modeled by chromatic simplicial complexes, which are abstract simplicial complexes whose vertexes are colored by Π. A (d-dimensional) *chromatic simplex* is a set $\{(a_0, v_0), \ldots, (a_d, v_d)\}$ comprising $d+1$ vertexes, where each $(a_i, v_i) \in \Pi \times \mathit{Value}$ is a vertex distinctly colored by a_i. A *chromatic simplicial complex* $\mathcal{C}$ is a finite collection of chromatic simplexes closed under set inclusion. We notice that a simplicial complex $\mathcal{C}$ is equally determined by the set of facets, which is denoted by $\mathsf{F}(\mathcal{C})$.

Intuitively, an n-dimensional chromatic simplex X models a global state of a distributed system of $n+1$ processes, where each vertex $(a, v) \in X$ models a process a that has its own private value v. The private value v is often referred to as the *view* of process a, and we write $\mathit{view}_a(X)$ for the view of process a in a simplex X. That is, $view_a(X) = v$ iff $(a, v) \in X$. We also define the *coloring function* χ by $\chi((a, v)) = a$. A set of possible global states of a nondeterministic distributed system of $n+1$ processes is modeled by a pure chromatic simplicial complex $\mathcal{C}$ of dimension n, where each facet $\{(0, v_0), (1, v_1), \ldots, (n, v_n)\} \in \mathsf{F}(\mathcal{C})$ denotes a possible global state of the system.

In what follows, we solely consider pure chromatic simplicial complexes. For brevity, we simply write vertexes, simplexes, and complexes to refer to their pure chromatic counterparts.

Suppose we are given complexes $\mathcal{C}$ and $\mathcal{D}$ of dimension n colored by Π. A *simplicial map* $\mu : V(\mathcal{C}) \to V(\mathcal{D})$ is a color-preserving function on vertexes such that $\chi(\mu(v)) = \chi(v)$ for every $v \in V(\mathcal{C})$ and also $\mu(X) \in \mathcal{D}$ for every $X \in \mathcal{C}$. We also define a complex of the cartesian product $\mathcal{C} \times \mathcal{D}$ as follows. For each pair of facets $X \in \mathsf{F}(\mathcal{C})$ and $Y \in \mathsf{F}(\mathcal{D})$, we define a simplex $X \times Y = \{(a, (u, v)) \mid (a, u) \in X, (a, v) \in Y, a \in \Pi\}$. The cartesian product $\mathcal{C} \times \mathcal{D}$ is a complex determined by the set of facets $\mathsf{F}(\mathcal{C} \times \mathcal{D}) = \{X \times Y \mid X \in \mathsf{F}(\mathcal{C}), Y \in \mathsf{F}(\mathcal{D})\}$.

The definition of task solvability in the topological model is summarized as follows. (For full details, see [18, Chapter 3].) Suppose $\mathcal{I}$ is the input complex, which models the possible set of input values to the system. A task is given by a *carrier map* $\Theta : \mathcal{I} \to 2^{\mathcal{O}}$, which is a monotonic function that maps each simplex $X \in \mathcal{I}$ to $\Theta(X)$, a subcomplex of $\mathcal{O}$ that is intended to model a collection of possible output states of the system. Similarly, a protocol is given by a carrier map $\Psi : \mathcal{I} \to 2^{\mathcal{P}}$. A task $\Theta : \mathcal{I} \to 2^{\mathcal{O}}$ is said to be solved by a protocol $\Psi : \mathcal{I} \to 2^{\mathcal{P}}$, if there exists a simplicial map μ such that $\mu(\Psi(X)) \subseteq \Theta(X)$ for every $X \in \mathcal{I}$. This means that, for any $Y \in \Psi(X)$, which is a set of output values produced by a nondeterministic execution of the protocol, the simplicial map μ yields a set of final outputs required by the task, namely $\mu(Y) \in \Theta(X)$. Here, μ defines how each output value $(a, v_a) \in Y$ produced by the protocol must be processed locally in each process a to yield a final output.

3.2 Epistemic logic for distributed computing

In this study, we exploit the *epistemic μ-calculus* [9], which extends the language of epistemic logic considered in [13, 33] with propositional greatest fixpoint construct.

3.2.1 The epistemic mu-calculus

The syntax of the language of epistemic μ-calculus is given by:

$$\varphi ::= p \mid \neg p \mid Z \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \mathrm{D}_A\, \varphi \mid \nu Z.\varphi,$$

where p ranges over the set AP of atomic propositions, Z ranges over the set PV of propositional variables, and A ranges over 2^{Π}, i.e., the powerset of processes[1].

Notice that this defines a class of *positive* formulas, in which atomic formulas are allowed to be negated, whereas non-propositional constructs $\mathrm{D}_A\, \varphi$ and $\nu Z.\varphi$ are not. Due to de Morgan's law, we may also allow a positive formula to contain negated forms of propositional formulas, which combine atomic formulas with propositional connectives $\neg$, $\vee$, and $\wedge$. Moreover, as a usual convention, we regard $\varphi \Rightarrow \psi$ as a synonym for the positive formula $\neg\varphi \vee \psi$.

This positiveness restriction is essential for the proof with logical obstruction. We show in Theorem 3.1 that any positive formula φ satisfies the knowledge gain property, which implies that φ is less likely to hold along a morphism between two epistemic models. (In the topological model, this amounts to the fact that a simplicial map transfers a complex to a more connected complex.) A logical obstruction thus implies the non-existence of a morphism, which indicates the task unsolvability.

A *Kripke frame* is a pair $\langle S, \sim \rangle$, where S is the set of (epistemic) *states*, and $\sim$ is a family $\{\sim_a \subseteq S \times S \mid a \in \Pi\}$ of *indistinguishability relations*, where each $\sim_a$ is an equivalence relation over S. The Kripke model[2] $M = \langle S, \sim, L \rangle$ augments the Kripke frame with a function $L : S \to 2^{\mathrm{AP}}$, which assigns to each state $X \in S$ a set $L(X)$ of atomic propositions that are satisfied at X. For a subset A of Π, we write $\sim_A$ for a derived equivalence relation defined by $X \sim_A Y$ iff $X \sim_a Y$ for all $a \in A$.

The formal semantics of epistemic formulas is given in Fig. 3, by induction on the structure of formulas. The semantics of a formula φ is given by the set $\|\varphi\|^M_\rho$ of states of M at which φ is satisfied,

[1] In this paper, we use "process" as a synonym for "agent" in multi-agent epistemic logic.

[2] More precisely, this is a Kripke model for multi-agent epistemic logic with axiom system $\mathbf{S5}_n$, where each indistinguishability relation $\sim_a$ is an equivalence relation [9].

where $\rho : \mathrm{PV} \to 2^S$ gives an interpretation of free propositional variables occurring in φ. The notation $\rho[S/Z]$ defines a modified interpretation of ρ, such that $\rho[S/Z](Z) = S$ and $\rho[S/Z](Z') = \rho(Z')$ for any propositional variable Z' other than Z.

$$
\begin{aligned}
\|p\|^M_\rho &= \{X \in S \mid p \in L(X)\} \qquad \|Z\|^M_\rho = \rho(Z) \qquad \|\neg p\|^M_\rho = S \setminus \|p\|^M_\rho \\
\|\varphi_1 \vee \varphi_2\|^M_\rho &= \|\varphi_1\|^M_\rho \cup \|\varphi_2\|^M_\rho \qquad\qquad \|\varphi_1 \wedge \varphi_2\|^M_\rho = \|\varphi_1\|^M_\rho \cap \|\varphi_2\|^M_\rho \\
\|\,\mathrm{D}_A\,\varphi\|^M_\rho &= \{X \in S \mid Y \in \|\varphi\|^M_\rho \text{ for every } Y \text{ such that } Y \sim_A X\} \\
\|\nu Z.\varphi\|^M_\rho &= \bigcup\{S' \in 2^S \mid S' \subseteq \|\varphi\|^M_{\rho[S'/Z]}\}
\end{aligned}
$$

Figure 3: Kripke semantics of epistemic μ-calculus

In addition to propositional formulas, the epistemic μ-calculus provides *distributed knowledge* $\mathrm{D}_A\,\varphi$, which is an epistemic modality that is intended to assert "every process in the group A knows φ." It further provides *greatest fixpoint* $\nu Z.\varphi$, which denotes the greatest solution satisfying the equation $Z = \varphi$ on the propositional variable Z [9, 30].[3] A greatest fixpoint formula $\nu Z.\varphi$ is logically equivalent to its unfolding, i.e., a formula obtained by replacing every free occurrence of Z in φ by $\nu Z.\varphi$.

In what follows, we write $M, X \models_\rho \varphi$ iff $X \in \|\varphi\|^M_\rho$, to mean that φ is *satisfied* at a particular state X of Kripke model M, under an interpretation ρ. In particular, we write $M, X \models \varphi$, when φ is a *closed* formula, i.e., φ contains no free occurrences of propositional variables. We also write $M \models \varphi$ to mean that a closed formula φ is *valid*, that is, $M, X \models \varphi$ holds for every state $X \in S$.

3.2.2 Simplicial model: a Kripke model induced from simplicial complex

The topological structure of a simplicial complex can be turned into a Kripke model, where a pair of facets of the complex are interpreted as epistemic states that are related by $\sim_a$ if and only if they share a vertex of color a [13]. More precisely, a simplicial complex $\mathcal{C}$ induces a Kripke model $M = \langle \mathsf{F}(\mathcal{C}), \sim, L\rangle$, where

- the set of (epistemic) states is $\mathsf{F}(\mathcal{C})$, the set of facets of $\mathcal{C}$;
- the indistinguishability relation $\sim$ is a family of equivalence relations $\{\sim_a \mid a \in \Pi\}$ over $\mathsf{F}(\mathcal{C})$, where each relation $\sim_a$ is defined by $X \sim_a Y$ iff $a \in \chi(X \cap Y)$;
- $L : \mathsf{F}(\mathcal{C}) \to 2^{\mathrm{AP}}$, where $\mathrm{AP} = \{\mathsf{input}^v_a \mid a \in \Pi, v \in \mathit{Value}\}$, is a function defined by $L(X) = \{\mathsf{input}^v_a \mid (a, v) \in X\}$.

There is a rigid correspondence between (pure chromatic) simplicial complexes and local proper Kripke models: they are isomorphic up to categorical equivalence under an appropriate categorical setting [13]. A Kripke model $M = \langle W, \sim, L\rangle$ is called local, if a process a always knows its own values, that is, $p \in L(w)$ iff $p \in L(w')$ whenever $w \sim_a w'$, for any atomic proposition $p \in \mathrm{AP}$ concerning process a; M is called proper, if any pair of distinct states $w, w' \in M$ are distinguishable, that is, $w \not\sim_a w'$ for some $a \in \Pi$.

This isomorphism indicates that a simplicial complex $\mathcal{C}$ and a Kripke model $\langle \mathsf{F}(\mathcal{C}), \sim, L\rangle$ induced from it are different representations of the same model. In this respect, we use these interchangeably in the rest of the paper, and refer to both as the *simplicial model*. Moreover, in abuse of notation, we may occasionally write $\mathcal{C}$ to denote a Kripke model $\langle \mathsf{F}(\mathcal{C}), \sim, L\rangle$ induced from the complex $\mathcal{C}$.

Fig. 4 illustrates a simplicial model in two different representations, by a complex and a Kripke model. Fig. 4(a) is a 2-dimensional complex comprising three facets, which models a system of three processes $\Pi = \{0, 1, 2\}$. Each facet comprises three vertexes, which are designated by $\circ$, $\bullet$, and $\bullet$, corresponding to processes 0, 1, and 2, respectively. (We follow this coloring convention throughout the paper.) Each vertex has its input value as its view. For example, the sole blue vertex in the figure designates the vertex $(1, 1)$, which models a process 1 of input 1. Fig. 4(b) shows the corresponding

[3] The logic in this study does not include the least fixpoint $\mu Z.\varphi$. Its logical equivalent $\neg\nu Z.\neg\varphi$ is not included either, as it is not a positive formula.

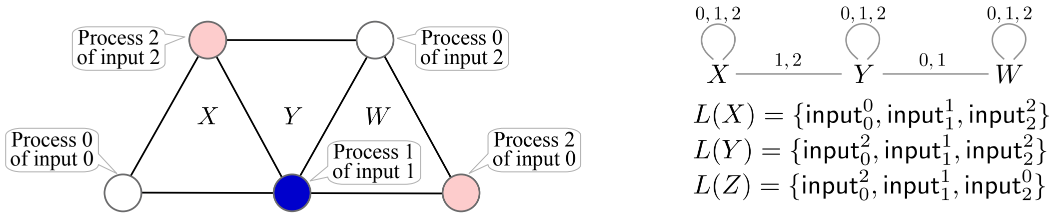

(a) A simplicial complex $\mathcal{C}$

(b) The corresponding Kripke model $\langle \mathsf{F}(\mathcal{C}), \sim, L\rangle$

Figure 4: A topological model and its corresponding Kripke model for a 3 process system

Kripke model. The Kripke frame is depicted by an undirected graph, where each graph node stands for a state, corresponding to a facet in $\mathcal{C}$, and each graph edge labeled by $a \in \Pi$ stands for a pair of nodes related by $\sim_a$. The function L gives the set of atomic formulas, where $\mathsf{input}_i^{v_i} \in L(w)$ indicates that process i is given the input value v_i in the facet that corresponds to the state w.

Here, it is instructive to observe the geometric interpretation of epistemic constructs. In the Kripke model $\langle \mathsf{F}(\mathcal{C}), \sim, L\rangle$ induced from a complex $\mathcal{C}$, $\sim_a$ relates a pair of facets if and only if they share a common vertex of color a in $\mathcal{C}$. Similarly, the derived relation $\sim_A$ relates a pair of facets if and only if they share a common face U such that $\chi(U) = A$. For example, in the complex of Fig. 4(a), $X \sim_{\{1,2\}} Y$ holds, because X and Y share a 1-dimensional simplex comprising vertexes $\bullet$ and $\bullet$. In contrast, X and W are not related by $\sim_{\{1,2\}}$ but by $\sim_{\{1\}}$ via the common 0-dimensional simplex (i.e., the vertex $\bullet$). Therefore $\mathcal{C}, X \models \mathrm{D}_{\{1,2\}}\, \mathsf{input}_2^2$ holds, because $\sim_{\{1,2\}}$ relates X with itself and Y, while $\mathcal{C}, X \not\models \mathrm{D}_{\{1\}}\, \mathsf{input}_2^2$ because $X \sim_{\{1\}} W$ but $\mathcal{C}, W \not\models \mathsf{input}_2^2$.[4]

The greatest fixpoint provides extra power for expressing epistemic modalities. For example, the common distributed knowledge $\mathrm{Cd}_{\mathcal{B}}\, \psi$ [3] can be defined as the greatest fixpoint $\nu Z.(\psi \wedge \bigwedge_{B\in\mathcal{B}} \mathrm{D}_B\, Z)$, where $\mathcal{B}$ is a class of subsets of Π. Informally, this can be understood as the conjunction of the formulas of infinitely nested modalities of the form $\psi \wedge \mathrm{D}_{B_1}(\psi \wedge \mathrm{D}_{B_2}(\psi \wedge \mathrm{D}_{B_3}(\cdots)))$, where $B_1, B_2, B_3, \ldots \in \mathcal{B}$. This means that the formula $\mathrm{Cd}_{\mathcal{B}}\, \psi$ is satisfied at some X_0 if and only if ψ is satisfied at every X_i along any infinite path $X_0 \sim_{B_1} X_1 \sim_{B_2} X_2 \sim_{B_3} \cdots$. In other words, every X_i is in the reach of the transitive closure of the relation $\bigcup_{B\in\mathcal{B}} \sim_B$. The common knowledge $\mathrm{C}_A\, \psi$ [9], where $A \subseteq \Pi$, is a special case of the common distributed knowledge $\mathrm{Cd}_{\mathcal{B}}\, \psi$ with $\mathcal{B} = \{\{a\} \mid a \in A\}$. For example, in the complex $\mathcal{C}$ of Fig. 4(a), $\mathcal{C}, X \models \mathrm{C}_{\{2\}}\, \mathsf{input}_2^2$ holds but $\mathcal{C}, X \models \mathrm{C}_{\{0,2\}}\, \mathsf{input}_2^2$ does not. This is because W, at which input_2^2 does not hold, is included in the transitive closure of $\sim_0 \cup \sim_2$ but not included in the transitive closure of $\sim_2$.

With this extra power of the greatest fixpoint, in Section 5 we define the logical obstruction in the form $\nu Z.\big(\psi \wedge \bigwedge_{\emptyset\subsetneq A\subseteq\Pi}(\varphi_A \Rightarrow \mathrm{D}_A(\psi' \wedge Z))\big)$. We remark that this greatest fixpoint formula cannot be substituted by a common distributed knowledge. This greatest fixpoint formula is, informally, understood as an infinitely nested sequence of modalities $\psi \wedge (\varphi_{A_1} \Rightarrow \mathrm{D}_{A_1}(\psi' \wedge \psi \wedge (\varphi_{A_2} \Rightarrow \mathrm{D}_{A_2}(\psi' \wedge \psi \wedge (\varphi_{A_3} \Rightarrow \cdots))))))$, where each nesting has different form $\mathrm{D}_{A_i}(\psi' \wedge \psi \wedge (\varphi_{A_i} \Rightarrow \cdots))$, depending on the choice of sequence $A_1, A_2, A_3, \ldots$ of nonempty subsets of A. In contrast, the common distributed knowledge $\mathrm{Cd}_{\mathcal{B}}\, \psi$ is interpreted by the nested modalities with the same formula ψ, for any choice of $B_1, B_2, \ldots \in \mathcal{B}$.

3.2.3 Morphisms between simplicial models

Let $\mathcal{C}$ and $\mathcal{D}$ be complexes and $\langle \mathsf{F}(\mathcal{C}), \sim, L\rangle$ and $\langle \mathsf{F}(\mathcal{D}), \sim', L'\rangle$ be Kripke models induced from them, respectively. We say a function $\delta : V(\mathcal{C}) \to V(\mathcal{D})$ is a *morphism*, if δ is a color-preserving simplicial map and furthermore $L(X) = L'(\delta(X))$ holds for every $X \in \mathsf{F}(\mathcal{C})$.

The knowledge gain theorem, which is essential for demonstrating task unsolvability in the logical method, is conservatively extended to allow the additional logic constructs, i.e., greatest fixpoints and distributed knowledge modalities.

[4] The distributed knowledge operator with a singleton set of processes $\mathrm{D}_{\{a\}}\, \varphi$ is known as the knowledge operator $\mathrm{K}_a\, \varphi$, which is omitted from the language of the logic in this study.

Theorem 3.1 (knowledge gain)**.** *Suppose $\mathcal{C} = \langle \mathsf{F}(\mathcal{C}), \sim, L\rangle$ and $\mathcal{D} = \langle \mathsf{F}(\mathcal{D}), \sim', L'\rangle$ are simplicial models and δ is a morphism from $\mathcal{C}$ to $\mathcal{D}$. Then, for any state $X \in \mathsf{F}(\mathcal{C})$ and closed positive formula φ, $\mathcal{D}, \delta(X) \models \varphi$ implies $\mathcal{C}, X \models \varphi$.*

Proof. It suffices to show that $\delta^{-1}(\|\varphi\|^{\mathcal{D}}_{\rho}) \subseteq \|\varphi\|^{\mathcal{C}}_{\delta^{-1}\circ\rho}$ holds for any interpretation ρ and any positive, but not necessarily closed, formula φ. Assuming $X \in \delta^{-1}(\|\varphi\|^{\mathcal{D}}_{\rho})$, we show $X \in \|\varphi\|^{\mathcal{C}}_{\delta^{-1}\circ\rho}$ by induction on φ.

Suppose φ is an atomic proposition p. Then, $\delta(X) \in \|p\|^{\mathcal{D}}_{\rho} = \{Y \in \mathsf{F}(\mathcal{D}) \mid p \in L'(Y)\}$ implies $p \in L'(\delta(X)) = L(X)$ and thus $X \in \|p\|^{\mathcal{C}}_{\delta^{-1}\circ\rho}$. The case of negated atomic formula $\neg p$ is proved in a similar manner. Suppose φ is a propositional variable Z. Then, $X \in \delta^{-1}(\|Z\|^{\mathcal{D}}_{\rho}) = \delta^{-1}(\rho(Z)) = \|Z\|^{\mathcal{C}}_{\delta^{-1}\circ\rho}$. Suppose φ is a disjunction $\varphi_1 \vee \varphi_2$. Then, $\delta(X) \in \|\varphi_1\|^{\mathcal{D}}_{\rho} \cup \|\varphi_2\|^{\mathcal{D}}_{\rho}$. Without loss of generality, we may assume $\delta(X) \in \|\varphi_1\|^{\mathcal{D}}_{\rho}$. By the induction hypothesis, we have $X \in \delta^{-1}(\|\varphi_1\|^{\mathcal{D}}_{\rho}) \subseteq \|\varphi_1\|^{\mathcal{C}}_{\delta^{-1}\circ\rho} \cup \|\varphi_2\|^{\mathcal{C}}_{\delta^{-1}\circ\rho}$. The case of conjunction is similarly proved.

Suppose φ is $\mathrm{D}_A\,\psi$. Let $X' \in \mathsf{F}(\mathcal{C})$ be any state such that $X \sim^{\mathcal{C}}_{A} X'$, which implies $\delta(X) \sim^{\mathcal{D}}_{A} \delta(X')$. Since $\delta(X) \in \|\,\mathrm{D}_A\,\psi\|^{\mathcal{D}}_{\rho}$, we have $\delta(X') \in \|\psi\|^{\mathcal{D}}_{\rho}$. Then, by the induction hypothesis, $X' \in \delta^{-1}(\|\psi\|^{\mathcal{D}}_{\rho}) \subseteq \|\psi\|^{\mathcal{C}}_{\delta^{-1}\circ\rho}$. As X' was arbitrarily chosen, this proves $X \in \|\,\mathrm{D}_A\,\psi\|^{\mathcal{C}}_{\delta^{-1}\circ\rho}$.

Suppose φ is $\nu Z.\psi$. We have $\delta(X) \in \|\nu Z.\psi\|^{\mathcal{D}}_{\rho} = \bigcup\{S' \mid S' \subseteq \|\psi\|^{\mathcal{D}}_{\rho[S'/Z]}\}$. Thus there exists S' such that $\delta(X) \in S' \subseteq \|\psi\|^{\mathcal{D}}_{\rho[S'/Z]}$, which implies $X \in \delta^{-1}(S') \subseteq \delta^{-1}(\|\psi\|^{\mathcal{D}}_{\rho[S'/Z]})$. Since $\delta^{-1}(\|\psi\|^{\mathcal{D}}_{\rho[S'/Z]}) \subseteq \|\psi\|^{\mathcal{C}}_{\delta^{-1}\circ(\rho[S'/Z])}$ by the induction hypothesis, we obtain $X \in \delta^{-1}(S') \subseteq \|\psi\|^{\mathcal{C}}_{\delta^{-1}\circ(\rho[S'/Z])} = \|\psi\|^{\mathcal{C}}_{(\delta^{-1}\circ\rho)[\delta^{-1}(S')/Z]}$. Therefore $X \in \bigcup\{S'' \mid S'' \subseteq \|\psi\|^{\mathcal{C}}_{(\delta^{-1}\circ\rho)[S''/Z]}\} = \|\nu Z.\psi\|^{\mathcal{C}}_{\delta^{-1}\circ\rho}$ holds, by taking $S'' = \delta^{-1}(S')$. $\square$

4 Product Update Models for k-Set Agreement

This section defines the product update models for the task and the protocol. The notion of product update [2, 34] was introduced for modeling dynamic updates of Kripke models. The logical method employs product updates to interpret carrier maps in the topological model as Kripke models, which are suitable for epistemic reasoning. For the general definition of the product update, see [13].

In the remainder of this paper, we assume that the initial inputs given to the processes are taken from Π, up to appropriate renaming of the input values.

4.1 Product updates and task solvability

The logical method specifies distributed computation by using product update models. A product update model is a subcomplex of the cartesian product $\mathcal{C} \times \mathcal{D}$ of two complexes (of the same dimension n).

Let $\langle \mathsf{F}(\mathcal{C}), \sim^{\mathcal{C}}, L\rangle$ be the simplicial model derived from $\mathcal{C}$. A *product update model* $\mathcal{C}[\mathcal{D}]$ is a simplicial model $\langle \mathsf{F}(\mathcal{C}[\mathcal{D}]), \sim^{\mathcal{C}[\mathcal{D}]}, L'\rangle$, where $\mathcal{C}[\mathcal{D}]$ is a subcomplex of $\mathcal{C} \times \mathcal{D}$, $\langle \mathsf{F}(\mathcal{C}[\mathcal{D}]), \sim^{\mathcal{C}[\mathcal{D}]}\rangle$ is the Kripke frame induced from $\mathcal{C}[\mathcal{D}]$, and L' is defined by $L'(X \times Y) = L(X)$ for every $X \times Y \in \mathsf{F}(\mathcal{C}[\mathcal{D}])$.

In a product update model $\mathcal{C}[\mathcal{D}]$, $X \times Y \in \mathsf{F}(\mathcal{C}[\mathcal{D}])$ indicates that $Y \in \mathsf{F}(\mathcal{D})$ is a possible output for the input $X \in \mathsf{F}(\mathcal{C})$. In other words, $\mathsf{F}(\mathcal{C}[\mathcal{D}])$ can be understood as a binary relation over $\mathsf{F}(\mathcal{C}) \times \mathsf{F}(\mathcal{D})$ that encodes a carrier map $\Psi : \mathcal{C} \to 2^{\mathcal{D}}$ such that $X \times Y \in \mathsf{F}(\mathcal{C}[\mathcal{D}])$ iff $Y \in \Psi(X)$. (Formally, such a binary relation is specified by a *precondition* $\mathsf{pre}(Y)$ for each $Y \in \mathsf{F}(\mathcal{D})$, where $\mathsf{pre}(Y)$ is an epistemic formula such that $\mathcal{C}, X \models \mathsf{pre}(Y)$ holds iff $Y \in \Psi(X)$, for every $X \in \mathsf{F}(\mathcal{C})$.)

In what follows, we write $\mathcal{I}$ to denote the input complex, i.e., a complex determined by the set of facets $\mathsf{F}(\mathcal{I}) = \{\{(0, v_0), \ldots, (n, v_n)\} \mid v_0, \ldots, v_n \in \Pi\}$.

Suppose $\mathcal{I}[\mathcal{T}]$ and $\mathcal{I}[\mathcal{P}]$ are product update models of a task and a protocol, respectively. Then, the notion of task solvability is defined as follows, by means of morphisms over these product update models.

Definition 1 (task solvability[13])**.** *A task $\mathcal{I}[\mathcal{T}]$ is solvable by a protocol $\mathcal{I}[\mathcal{P}]$ if there exists a morphism $\delta : \mathcal{I}[\mathcal{P}] \to \mathcal{I}[\mathcal{T}]$ such that $\pi_{\mathcal{I}} \circ \delta = \pi_{\mathcal{I}}$, where $\pi_{\mathcal{I}}$ is the first projection defined by $\pi_{\mathcal{I}}((X \times Y)) = X$.*

However, the product update model presented above is not sufficient to show the unsolvability of the k-set agreement, as mentioned in Section 1. We must extend it with factual change [32] such that not only the input values, but also the output decision values are mentioned by atomic propositions.

Let $\mathcal{I}[\mathcal{T}]$ and $\mathcal{I}[\mathcal{P}]$ be product update models $\langle \mathsf{F}(\mathcal{I}[\mathcal{T}]), \sim^{\mathcal{I}[\mathcal{T}]}, L\rangle$ and $\langle \mathsf{F}(\mathcal{I}[\mathcal{P}]), \sim^{\mathcal{I}[\mathcal{P}]}, L'\rangle$, respectively. Let us assume an augmented set of atomic propositions $\widehat{\mathrm{AP}} = \mathrm{AP} \cup \{\mathsf{decide}_a^d \mid a \in \Pi, d \in \mathit{Value}\}$, where decide_a^d is an atomic proposition asserting that process a decides d as its output value. Suppose $\delta : \mathcal{I}[\mathcal{P}] \to \mathcal{I}[\mathcal{T}]$ is a morphism satisfying $\pi_{\mathcal{I}} \circ \delta = \pi_{\mathcal{I}}$. Then, the product update model with factual change for the task, written $\widehat{\mathcal{I}[\mathcal{T}]}$, is a modified product update model $\langle \mathsf{F}(\mathcal{I}[\mathcal{T}]), \sim^{\mathcal{I}[\mathcal{T}]}, \widehat{L}\rangle$, where $\widehat{L}$ augments the assignment of true atomic propositions by $\widehat{L}(X \times Y) = \{\mathsf{input}_a^v \mid (a, v) \in X\} \cup \{\mathsf{decide}_a^d \mid (a, d) \in Y\}$. Furthermore, the product update model with factual change for the protocol is given by $\mathcal{I}[\mathcal{P}]_\delta = \langle \mathsf{F}(\mathcal{I}[\mathcal{P}]), \sim^{\mathcal{I}[\mathcal{P}]}, L_\delta\rangle$, where $L_\delta(X \times Y) = \widehat{L}(\delta(X \times Y))$ for every $X \times Y \in \mathsf{F}(\mathcal{I}[\mathcal{P}])$. Since $\pi_{\mathcal{I}} \circ \delta = \pi_{\mathcal{I}}$, L_δ is an extension of L', i.e., $L'(X \times Y) \subseteq L_\delta(X \times Y)$.

In the proof of task unsolvability, we resort to the following property of product updates with factual change. (This claim follows from Theorem 3.1. For the details, see the discussion in the proof of Theorem 19 of [32].)

Theorem 4.1 (knowledge gain with factual change)**.** *Let $\mathcal{I}[\mathcal{T}]$ and $\mathcal{I}[\mathcal{P}]$ be the product update models of a task and a protocol, respectively. Suppose there exists a morphism $\delta : \mathcal{I}[\mathcal{P}] \to \mathcal{I}[\mathcal{T}]$ such that $\pi_{\mathcal{I}} \circ \delta = \pi_{\mathcal{I}}$. Then, $\widehat{\mathcal{I}[\mathcal{T}]}, \delta(X) \models \varphi$ implies $\mathcal{I}[\mathcal{P}]_\delta, X \models \varphi$, for any $X \in \mathsf{F}(\mathcal{I}[\mathcal{P}]_\delta)$ and positive formula φ.*

This lemma provides a logical means of proving task unsolvability. Suppose we have a positive formula φ, such that φ is valid in the product update model of the task $\mathcal{I}[\mathcal{T}]$ but not valid in that of the protocol $\mathcal{I}[\mathcal{P}]$, for any morphism δ satisfying $\pi_{\mathcal{I}} \circ \delta = \pi_{\mathcal{I}}$. This implies that the task is unsolvable, in the sense of Definition 1. We call such a formula φ a *logical obstruction*.

4.2 The product update model for k-set agreement task

The output complex of the k-set agreement task is given by the set of facets $\mathsf{F}(\mathcal{SA}_k) = \big\{\{(0, d_0), \ldots, (n, d_n)\} \mid |\{d_0, \ldots, d_n\}| \leq k, d_0, \ldots, d_n \in \Pi\big\}$, which respects the agreement property. The product update model $\mathcal{I}[\mathcal{SA}_k]$ of the k-set agreement task is a simplicial model $\langle \mathsf{F}(\mathcal{I}[\mathcal{SA}_k]), \sim^{\mathcal{I}[\mathcal{SA}_k]}, L\rangle$, where $\mathcal{I}[\mathcal{SA}_k]$ is determined by $\mathsf{F}(\mathcal{I}[\mathcal{SA}_k]) = \{I \times O \mid I \in \mathsf{F}(\mathcal{I}), O \in \mathsf{F}(\mathcal{SA}_k), \{v \mid (a, v) \in I\} \supseteq \{d \mid (a, d) \in O\}\}$, which is the set of facets respecting the validity condition.

This product update model can be extended to the one with factual change $\widehat{\mathcal{I}[\mathcal{SA}_k]} = \langle \mathsf{F}(\mathcal{I}[\mathcal{SA}_k]), \sim^{\mathcal{I}[\mathcal{SA}_k]}, \widehat{L}\rangle$, where $\widehat{L}(I \times O) = \{\mathsf{input}_a^v \mid (a, v) \in I\} \cup \{\mathsf{decide}_a^d \mid (a, d) \in O\}$ for each $I \times O \in \mathsf{F}(\mathcal{I}[\mathcal{SA}_k])$.

The following formulas are valid in the product update model $\widehat{\mathcal{I}[\mathcal{SA}_k]}$.

$$\mathrm{OFUN} = \textstyle\bigwedge_{a\in\Pi}\big((\bigwedge_{d,e\in\Pi,d\neq e} \neg(\mathsf{decide}_a^d \wedge \mathsf{decide}_a^e)) \wedge \bigvee_{d\in\Pi} \mathsf{decide}_a^d\big) \tag{1}$$

$$\mathrm{VALID} = \textstyle\bigwedge_{a\in\Pi} \bigwedge_{d\in\Pi}\big(\mathsf{decide}_a^d \Rightarrow \bigvee_{b\in\Pi} \mathsf{input}_b^d\big) \tag{2}$$

$$\mathrm{AGREE}_k = \textstyle\bigvee_{A\subseteq\Pi,0<|A|\leq k} \bigwedge_{a\in\Pi} \bigvee_{d\in A} \mathsf{decide}_a^d \tag{3}$$

$$\mathrm{KNOW} = \textstyle\bigwedge_{a\in\Pi} \bigwedge_{d\in\Pi}\big(\mathsf{decide}_a^d \Rightarrow \mathrm{D}_{\{a\}}\, \mathsf{decide}_a^d\big) \tag{4}$$

The formula OFUN means that each process a decides a unique output value. The formula VALID expresses the validity condition: any output value d must be an input to some of the processes. The formula AGREE_k specifies that processes may decide at most k different values. Finally, the formula KNOW indicates that, given two facets $\sigma, \sigma' \in \mathsf{F}(\mathcal{I}[\mathcal{SA}_k])$ such that $\sigma \sim_a \sigma'$, process a decides the same output value at both σ and σ'.

4.3 Product update model of single-round immediate snapshot protocol

The immediate snapshot protocol is a fundamental piece in the topological method that bridges distributed computing and combinatorial topology [5, 14, 28]. In its execution, each process obtains a *snapshot view* through the concurrent execution in the order $A_1, \ldots, A_r$ ($r \geq 1$), where A_i's are disjoint subsets of processes, referred to as concurrency classes. Up to a nondeterministic choice of a sequence of concurrency classes, the processes in the same concurrency class A_m simultaneously obtain the same snapshot view collected from the processes in $\bigcup_{i=1}^m A_i$, i.e, the processes in A_m and its preceding concurrency classes.

The immediate snapshot protocol has a tight connection with combinatorial topology. Herein, we review the combinatorial aspects of the single-round immediate snapshot protocol. (We discuss

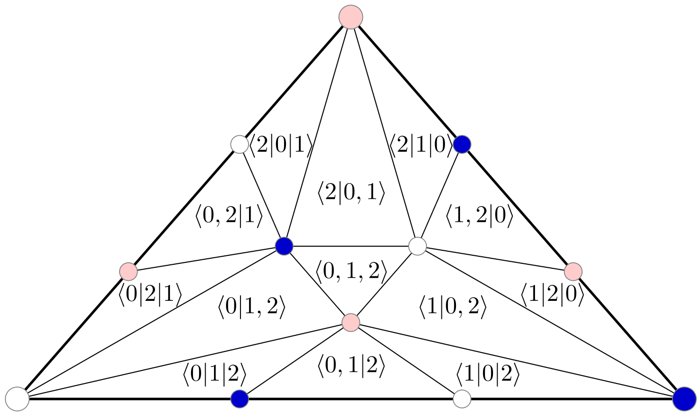

Figure 5: Standard chromatic subdivision of a 2-dimensional simplex, where the facets are labeled with ordered set partitions.

the multiple-round immediate snapshot protocol in Section 4.4.) A single-round immediate snapshot protocol corresponds to a geometric operation of subdivision, called the *standard chromatic subdivision* [21, 18]. Given an n-dimensional simplex that models a distributed system of $n+1$ processes, a single-round immediate snapshot protocol subdivides the simplex into a complex of the standard chromatic subdivision, where each facet of the subdivision corresponds to a sequence of concurrency classes or an *ordered set partition* [26]. An ordered set partition of Π, written $\langle A_1|A_2|\cdots|A_r\rangle$ $(r \geq 1)$, is a sequence of pairwisely disjoint, nonempty subsets of Π that satisfy $\Pi = \bigcup_{i=1}^{r} A_i$.

Figure 5 illustrates the standard chromatic subdivision of a 2-dimensional simplex, where facets are indexed by ordered set partitions. We occasionally omit curly braces in ordered set partitions, e.g., $\langle 0,1|2|3\rangle$ instead of $\langle\{0,1\}|\{2\}|\{3\}\rangle$.

In what follows, we write γ, γ', etc. to denote ordered set partitions. Given a facet X and an ordered set partition $\gamma = \langle A_1|A_2|\cdots|A_r\rangle$ of the set $\chi(X)$, we write $X \rtimes \gamma$ for a facet of the standard chromatic subdivision of X that is determined by γ. More concretely, $X \rtimes \gamma$ denotes a facet comprising the set of vertexes $\{(a, \mathit{view}_a(X \rtimes \gamma)) \mid a \in \chi(X)\}$, where $\mathit{view}_a(X \rtimes \gamma)$ is the view of process a defined by $\mathit{view}_a(X \rtimes \gamma) = \{(b, v) \in X \mid b \in \bigcup_{i=1}^{q} A_i,\ \text{where } a \in A_q\}$.

Below we present some properties concerning the adjacency of facets, which are useful in the proof of Section 5.

The following holds, because $(a, v) \in X$ iff $(a, v) \in \mathit{view}_a(X \rtimes \gamma)$.

Lemma 4.2. *$X \rtimes \gamma \sim_a Y \rtimes \gamma'$ implies $X \sim_a Y$.*

The product update model $\mathcal{I}[\mathcal{IS}]$ of the single-round immediate snapshot protocol is a simplicial model $\langle \mathsf{F}(\mathcal{I}[\mathcal{IS}]), \sim^{\mathcal{I}[\mathcal{IS}]}, L'\rangle$ where

- $\mathsf{F}(\mathcal{I}[\mathcal{IS}]) = \{X \rtimes \gamma \mid X \in \mathsf{F}(\mathcal{I}) \text{ and } \gamma \text{ is an ordered set partition of } \Pi\}$[5] is the set of facets of the standard chromatic subdivision of the complex $\mathcal{I}$;
- $L'(X \rtimes \gamma) = \{\mathsf{input}_a^v \mid (a, v) \in X\}$, for each $X \rtimes \gamma \in \mathsf{F}(\mathcal{I}[\mathcal{IS}])$.

Following [26], below we provide a combinatorial condition for a pair of facets of the standard chromatic subdivision of an n-dimensional simplex to share a common face of dimension $n-1$.

Lemma 4.3. *Suppose $X \rtimes \gamma, X \rtimes \gamma' \in \mathsf{F}(\mathcal{I}[\mathcal{IS}])$, such that $\gamma \neq \gamma'$, i.e., $X \rtimes \gamma$ and $X \rtimes \gamma'$ are distinct facets of the standard chromatic subdivision of $X \in \mathsf{F}(\mathcal{I})$. Let $b \in [0, n]$. Then, $X \rtimes \gamma \sim_{[0,n]\setminus\{b\}} X \rtimes \gamma'$ and $X \rtimes \gamma \not\sim_b X \rtimes \gamma'$ holds if and only if γ and γ' are a pair of ordered set partitions of the forms $\langle A_1|\cdots|A_{s-1}|\, b\,|A_s|\cdots|A_r\rangle$ and $\langle A_1|\cdots|A_{s-1}|A_s \cup \{b\}|\cdots|A_r\rangle$ $(1 \leq s \leq r)$.*

[5] We abbreviate a facet of $\mathcal{I}[\mathcal{IS}]$ as $X \rtimes \gamma$, where we omit the duplicate of X in the formal notation $X \times (X \rtimes \gamma)$. The redundancy in the formal notation is due to the definition of product update, where the action model $\mathcal{IS}$ must be defined relative to the input complex $\mathcal{I}$, in order for a precondition to relate each action in $\mathcal{IS}$ (i.e., a facet of the standard chromatic subdivision) with an appropriate set of facets in $\mathcal{I}$.

When the ordered set partition γ does not have either of the forms in the above lemma, i.e., when $\gamma = \langle A_1|\cdots|A_{r-1}|\,b\rangle$, the facet $X \rtimes \gamma \in \mathsf{F}(\mathcal{I}[\mathcal{IS}])$ touches on a boundary face W of X, where W is an $(n-1)$-dimensional simplex satisfying $\chi(W) = [0,n] \setminus \{b\}$. In this case, the above lemma implies that $X \rtimes \gamma$ is the sole facet of $\mathcal{I}[\mathcal{IS}]$ that touches on the boundary face W at $W \rtimes \langle A_1|\cdots|A_{r-1}\rangle$, a facet of the subdivision of W. (For example, in Fig. 5, the three facets of the standard chromatic subdivision that touch on the lower edge of the triangle are identified by ordered set partitions $\langle 0|1|2\rangle$, $\langle 0,1|2\rangle$, and $\langle 1|0|2\rangle$.)

In Section 2, we discussed Sperner's lemma as a property derived from a graph that spans over the facets of a subdivision of a simplex, where each graph node is a facet of the subdivision or a facet that touches on a particular lower dimensional face of the simplex. To uniquely index such facets using ordered set partitions, for each d ($0 \le d \le n$), we use a subset OSP_d of ordered set partitions of the form $\langle A_1|\cdots|A_r|d+1|\cdots|n\rangle$, where $A_1,\ldots,A_r$ are a partition of the subset $[0,d]$. (By definition, $OSP_d \subseteq OSP_{d+1}$ holds for every d ($0 \le d < n$). In particular, OSP_n is the set of ordered set partitions of $[0,n]$.) For $\gamma = \langle A_1|\cdots|A_r|d+1|\cdots|n\rangle \in OSP_d$, $X \rtimes \gamma$ designates a facet of $\mathcal{I}[\mathcal{IS}]$ such that the first half $\langle A_1|\cdots|A_r\rangle$ identifies a facet $W \rtimes \langle A_1|\cdots|A_r\rangle$ of the subdivision of a d-dimensional face W of X such that $\chi(W) = [0,d]$ and furthermore the remaining half $\langle d+1|\cdots|n\rangle$ uniquely determines a facet of the subdivision of X, out of many possible facets touching the d-dimensional face at $W \rtimes \langle A_1|\cdots|A_r\rangle$.

Generalizing Lemma 4.3, we obtain the following proposition that gives a combinatorial condition for a pair of facets that touches on the d-dimensional boundary to have a common face of dimension $d-1$.

Proposition 4.4. *Let $d \in [1,n]$, $b \in [0,d]$, and $A = [0,d] \setminus \{b\}$. Suppose $X \rtimes \gamma, Y \rtimes \gamma' \in \mathsf{F}(\mathcal{I}[\mathcal{IS}])$ and $\gamma, \gamma' \in OSP_d$, where $\gamma = \langle A_1|\cdots|A_r|d+1|\cdots|n\rangle$ and $\gamma' = \langle A'_1|\cdots|A'_{r'}|d+1|\cdots|n\rangle$. Then, $X \rtimes \gamma \sim_A Y \rtimes \gamma'$ holds if and only if either of the following conditions are satisfied:*

(i) $\gamma' = \gamma$ and either $X \sim_{[0,d]} Y$ or $A_r = \{b\}$.

(ii) $\gamma' = \underline{\gamma}_A$ and $A_r \ne \{b\}$, where

$$\underline{\gamma}_A = \begin{cases} \langle A_1|\cdots|A_{s-1}|\,b\,|A_s\backslash\{b\}|\cdots|A_r|d+1|\cdots|n\rangle \\ \qquad\qquad\qquad\qquad \text{if } b \in A_s \text{ and } |A_s| > 1 \ (s \le r), \\ \langle A_1|\cdots|A_{s-1}|A_{s+1}\cup\{b\}|\cdots|A_r|d+1|\cdots|n\rangle \\ \qquad\qquad\qquad\qquad \text{if } A_s = \{b\} \ (s < r). \end{cases}$$

Proof. The if direction is easy to show. In both cases, $view_a(X \rtimes \gamma) = view_a(Y \rtimes \gamma')$ holds for every $a \in A$.

For the converse, suppose $X \rtimes \gamma \sim_A Y \rtimes \gamma'$. We write β (resp., β') for $\langle A_1|\cdots|A_r\rangle$ (resp., $\langle A'_1|\cdots|A'_{r'}\rangle$). Let $U \subseteq X$ and $W \subseteq Y$ be the faces satisfying $\chi(U) = \chi(W) = [0,d]$. Then, $U \rtimes \beta$ and $W \rtimes \beta'$ are faces of $X \rtimes \gamma$ and $Y \rtimes \gamma'$, respectively, and they satisfy $U \rtimes \beta \sim_A W \rtimes \beta'$. Consider the case $\gamma' = \gamma$. We show that the condition (i) holds. Assume $A_r \ne \{b\}$ and let $c \in A_r \setminus \{b\}$. (There exists such c, because $A_r \ne \emptyset$.) Since $c \in A$ and $U \rtimes \beta \sim_A W \rtimes \beta'$, we must have $\{(a,v) \in U \mid a \in [0,d]\} = view_c(U \rtimes \beta) = view_c(W \rtimes \beta') = \{(a,v) \in W \mid a \in [0,d]\}$. This implies $U \sim_{[0,d]} W$, and therefore $X \sim_{[0,d]} Y$. Let us consider the other case $\gamma' \ne \gamma$. We show that the condition (ii) holds. Suppose $A_r = \{b\}$. If $A'_{r'} \ne \{b\}$, there exists $c \in A'_{r'} \setminus \{b\} \subseteq A$, which implies $view_c(U \rtimes \beta) \ne view_c(W \rtimes \beta')$. If $A'_{r'} = \{b\}$, there exists $c \in A$ such that $view_c(U \rtimes \beta) \ne view_c(W \rtimes \beta')$. In both cases, we have $U \rtimes \beta \not\sim_A W \rtimes \beta'$, a contradiction. Therefore, $A_r \ne \{b\}$, and we have $\gamma' = \underline{\gamma}_A$ by Lemma 4.3. □

4.4 Product update model of multiple-round m-iterated immediate snapshot protocol

Let us consider the multiple-round immediate snapshot protocol, i.e., the m-iterated execution of the single immediate snapshot protocol ($m \ge 1$). This is modeled by m-iterated standard chromatic subdivision, which subdivides the initial input complex to become finer at each iteration.

A facet of the m-iterated subdivision on an input simplex X is designated by $X \rtimes \gamma_1 \rtimes \cdots \rtimes \gamma_m$, where each ordered set partition γ_i identifies a facet of the i-th subdivision, which further subdivides a facet of the preceding $(i-1)$-th subdivision. (Assuming that $\rtimes$ associates to left, we omit the parentheses in $(\cdots(X \rtimes \gamma_1) \rtimes \cdots) \rtimes \gamma_m$.) The view of a vertex in this facet is defined by induction on m, that is, $\mathit{view}_a(X \rtimes \gamma_1 \rtimes \cdots \rtimes \gamma_m \rtimes \gamma_{m+1}) = \mathit{view}_a(Y \rtimes \gamma_{m+1})$, where $Y = \{(a, \mathit{view}_a(X \rtimes \gamma_1 \rtimes \cdots \rtimes \gamma_m)) \mid a \in \chi(X)\}$.

The product update model of $\mathcal{I}[\mathcal{IS}^m]$ is a simplicial model $\langle \mathsf{F}(\mathcal{I}[\mathcal{IS}^m]), \sim^{\mathcal{I}[\mathcal{IS}^m]}, L' \rangle$ where

- $\mathsf{F}(\mathcal{I}[\mathcal{IS}^m]) = \{X \rtimes \gamma_1 \rtimes \cdots \rtimes \gamma_m \mid X \in \mathsf{F}(\mathcal{I}), \gamma_1, \ldots, \gamma_m \in \mathit{OSP}_n\}$ is the set of facets of the m-th iterated standard chromatic subdivision of the complex $\mathcal{I}$;
- $L'(X \times \gamma_1 \rtimes \cdots \rtimes \gamma_m) = \{\mathsf{input}_a^v \mid (a, v) \in X\}$.

The Proposition 4.4 is further generalized to iterated subdivision.

Proposition 4.5. *Let $d \in [1, n]$, $b \in [0, d]$, and $A = [0, d] \backslash \{b\}$. Suppose $X \rtimes \gamma_1 \rtimes \cdots \rtimes \gamma_m, Y \rtimes \gamma_1' \rtimes \cdots \rtimes \gamma_m' \in \mathsf{F}(\mathcal{I}[\mathcal{IS}^m])$, where $X, Y \in \mathsf{F}(\mathcal{I})$ are facets, and $\gamma_1, \gamma_1', \ldots, \gamma_m, \gamma_m' \in \mathit{OSP}_d$ are ordered set partitions such that $X \sim_A Y$ and $\gamma_i = \langle A_{i,1} | \cdots | A_{i,r_i} | d+1 | \cdots | n \rangle$ $(1 \leq i \leq m)$. Then, $X \rtimes \gamma_1 \rtimes \cdots \rtimes \gamma_m \sim_A Y \rtimes \gamma_1' \rtimes \cdots \rtimes \gamma_m'$ holds if and only if either of the following conditions are satisfied:*

(i) $\gamma_i' = \gamma_i$ for every i $(1 \leq i \leq m)$ and either $X \sim_{[0,d]} Y$ or $A_{1,r_1} = \cdots = A_{m,r_m} = \{b\}$.

(ii) There exists j $(1 \leq j \leq m)$, such that $\gamma_j' = \underline{\gamma_j}_A$, $A_{j,r_j} \neq \{b\}$, $\gamma_i' = \gamma_i$ for every i $(1 \leq i < j)$, and $\gamma_i' = \gamma_i$ and $A_{i,r_i} = \{b\}$ for every i $(j < i \leq m)$.

Proof. The if direction is easy to show. In both cases, $\mathit{view}_a(X \rtimes \gamma_1 \rtimes \cdots \rtimes \gamma_m) = \mathit{view}_a(Y \rtimes \gamma_1' \rtimes \cdots \rtimes \gamma_m')$ holds for every $a \in A$.

For the converse, the proof proceeds by induction on m. The case $m = 1$ is already demonstrated in Proposition 4.4. For the case $m > 1$, assume that $X \rtimes \gamma_1 \rtimes \cdots \rtimes \gamma_m \sim_A Y \rtimes \gamma_1' \rtimes \cdots \rtimes \gamma_m'$. By Lemma 4.2, we have $X \rtimes \gamma_1 \rtimes \cdots \rtimes \gamma_{m-1} \sim_A Y \rtimes \gamma_1' \rtimes \cdots \rtimes \gamma_{m-1}'$, to which the induction hypothesis applies. Consider the case $\gamma_m = \gamma_m'$. Then, the condition (i) or (ii) follows from Proposition 4.4 and the induction hypothesis. Consider the other case $\gamma_m \neq \gamma_m'$. Suppose that the condition (ii) holds for the induction hypothesis, that is, there exists p $(1 \leq p \leq m-1)$ such that $\gamma_p' = \underline{\gamma_p}_A$, $A_{p,r_p} \neq \{b\}$, $\gamma_i' = \gamma_i$ for every i $(1 \leq i < p)$, and $\gamma_i' = \gamma_i$ and $A_{i,r_i} = \{b\}$ for every i $(p < i \leq m-1)$. This contradicts $X \rtimes \gamma_1 \rtimes \cdots \rtimes \gamma_{m-1} \sim_A Y \rtimes \gamma_1' \rtimes \cdots \rtimes \gamma_{m-1}'$, because $\mathit{view}_c(X \rtimes \gamma_1 \rtimes \cdots \rtimes \gamma_{m-1}) \neq \mathit{view}_c(Y \rtimes \gamma_1' \rtimes \cdots \rtimes \gamma_{m-1}')$ for some $c \in A_{j,r_j} \setminus \{b\} \subseteq A$. Thus the induction hypothesis must satisfy the condition (i). Furthermore, according to Proposition 4.4(ii), we have $\gamma_m' = \underline{\gamma_m}_A$ and $A_{m,r_m} \neq \{b\}$. Therefore the condition (ii) holds with $j = m$. $\square$

5 Logical Obstruction to k-Set Agreement

We apply the logical method to show that the k-set agreement task is not solvable by the m-iterated immediate snapshot protocol, if k is less than $n+1$, i.e., the number of processes.

Let us consider the following positive epistemic formula:

$$\Phi_k = \nu Z. \left[\mathrm{OFUN} \wedge \mathrm{VALID} \wedge \mathrm{KNOW} \wedge \bigwedge_{\emptyset \subsetneq A \subseteq \Pi} \left(\mathrm{DEC}_A \Rightarrow \mathrm{D}_A(\mathrm{AGREE}_k \wedge Z) \right) \right],$$

where $\mathrm{DEC}_A = \bigwedge_{d=0}^{|A|-1} \bigvee_{a \in A} \mathsf{decide}_a^d$.

We show that Φ_k is a logical obstruction to the k-set agreement task. Let $\mathcal{I}[\mathcal{SA}_k]$ be the product update model of the k-set agreement task (Section 4.2) and $\mathcal{I}[\mathcal{IS}^m]$ be the product update model of the m-iterated immediate snapshot protocol (Section 4.4). Given a morphism $\delta : \mathcal{I}[\mathcal{IS}^m] \to \mathcal{I}[\mathcal{SA}_k]$ such that $\pi_{\mathcal{I}} \circ \delta = \pi_{\mathcal{I}}$, we obtain product update models with factual change $\widehat{\mathcal{I}[\mathcal{SA}_k]} = \langle \mathsf{F}(\mathcal{I}[\mathcal{SA}_k]), \sim^{\mathcal{I}[\mathcal{SA}_k]}, \widehat{L} \rangle$ for the task and $\mathcal{I}[\mathcal{IS}^m]_\delta = \langle \mathsf{F}(\mathcal{I}[\mathcal{IS}^m]), \sim^{\mathcal{I}[\mathcal{IS}^m]}, L_\delta \rangle$ for the protocol, as defined in Section 4.1. By definition, $\mathcal{I}[\mathcal{IS}^m]_\delta, X \times Y \models \mathsf{decide}_a^d$ iff $\widehat{\mathcal{I}[\mathcal{SA}_k]}, \delta(X \times Y) \models \mathsf{decide}_a^d$. This means that whenever a process a decides an output value d by δ, d is a valid output of the k-set agreement task.

In the sequel, let σ, τ, ... range over the facets of these product update models.

5.1 Unsolvability proof with logical obstruction

We show the unsolvability result by appealing to Theorem 4.1. Let us assume, by contradiction, that there exists a morphism $\delta : \mathcal{I}[\mathcal{IS}^m] \to \mathcal{I}[\mathcal{SA}_k]$ such that $\pi_{\mathcal{I}} \circ \delta = \pi_{\mathcal{I}}$. We show that there exists $\sigma \in \mathsf{F}(\mathcal{I}[\mathcal{IS}^m]_\delta)$ such that $\widehat{\mathcal{I}[\mathcal{SA}_k]}, \delta(\sigma) \models \Phi_k$ but $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma \not\models \Phi_k$, which contradicts Theorem 4.1.

5.1.1 Proof of $\widehat{\mathcal{I}[\mathcal{SA}_k]}, \delta(\sigma) \models \Phi_k$

As we have seen in Section 4.2, OFUN, VALID, KNOW, AGREE$_k$ are all valid formulas in the product update model $\widehat{\mathcal{I}[\mathcal{SA}_k]}$. Therefore Φ_k is trivially valid, meaning that $\widehat{\mathcal{I}[\mathcal{SA}_k]}, \delta(\sigma) \models \Phi_k$ holds for any $\sigma \in \mathsf{F}(\mathcal{I}[\mathcal{IS}^m]_\delta)$.

5.1.2 Proof of $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma \not\models \Phi_k$

Assume, by contradiction, Φ_k is valid in $\mathcal{I}[\mathcal{IS}^m]_\delta$. In what follows, for each d $(0 \le d \le n)$, let F_d be a subclass of the facets of $\mathcal{I}[\mathcal{IS}^m]_\delta$ defined by:

$$F_d = \{I_d \rtimes \gamma_1 \rtimes \cdots \rtimes \gamma_m \in \mathsf{F}(\mathcal{I}[\mathcal{IS}^m]_\delta) \mid \gamma_i \in OSP_d (1 \le i \le m)\},$$

where $I_d = \{(i,i) \mid i \in [0,d]\} \cup \{(i,d) \mid i \in [d+1,n]\}$ is a facet of the input complex $\mathcal{I}$.

Fig. 6 illustrates the sets of facets F_2, F_1, and F_0 in $\mathcal{I}[\mathcal{IS}^2]_\delta$. Each F_d is a subset of the facets of the standard chromatic subdivision of I_d, over which the subgraph for the induction case of dimension d in Section 2 spans. In general, every facet $\sigma \in F_d$ $(0 \le d \le n)$ touches on the boundary $I_d \cap I_{d+1} = \{(i,i) \mid i \in [0,d]\}$, i.e., the d-dimensional common face shared by I_d and I_{d+1}, such that each facet of the subdivision of the common face is subsumed by exactly two facets, one in F_d and the other in F_{d+1}. In particular, the subdivision of the 0-dimensional boundary $I_0 \cap I_1 = \{(0,0)\}$ is a single vertex and is subsumed by the sole facet in F_0 and another facet in F_1.

In the remainder of this section, we follow the steps below to construct a graph path that spans over $\bigcup_{d=0}^{k} F_d$, using the output decision values determined by the morphism δ as a Sperner labeling on the vertexes of the complex $\mathcal{I}[\mathcal{IS}^m]$.

1. We define a family of irreflexive symmetric relations $\frown_A$ over facets, where $\sigma \frown_A \sigma'$ indicates that σ and σ' are adjacent facets connected by a graph edge.
2. Starting from a particular facet $\sigma_0 = I_0 \rtimes \gamma_1 \rtimes \cdots \rtimes \gamma_m \in F_0$, where $\gamma_1 = \cdots = \gamma_m = \langle 0|1|\cdots|n\rangle$, we construct a graph path $\sigma_0 \frown_{A_1} \sigma_1 \frown_{A_2} \sigma_2 \frown_{A_3} \cdots$, as we did in Section 2.
3. Assuming that OFUN $\wedge$ VALID $\wedge$ AGREE$_k$ $\wedge$ KNOW is an invariant that holds at every facet σ_i $(i \ge 1)$, we show that every σ_i is connected with exactly two other facets (Proposition 5.3). This indicates that the graph path is cycle-free.[6]
4. We obtain $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma_0 \models \Phi_k$ using Theorem 4.1. From this, we deduce that OFUN $\wedge$ VALID $\wedge$ AGREE$_k$ $\wedge$ KNOW is an invariant that holds at every σ_i, by unfolding the fixpoint construct $\nu Z. \cdots$ in Φ_k for each time we traverse a graph edge (Theorem 5.4).

Overall, we obtain a cycle-free graph path of unbounded length that spans over $\mathcal{I}[\mathcal{IS}^m]_\delta$, albeit this contradicts the fact that $\mathcal{I}[\mathcal{IS}^m]_\delta$ is a finite model.

Let us first show some geometric properties satisfied by the facets in $F_0, \ldots, F_n$.

Lemma 5.1. *The set $F_0 = \{\sigma_0\}$ is a singleton, where $\sigma_0 = I_0 \rtimes \langle 0|1|\cdots|n\rangle$. Moreover, $\sigma_1 = I_1 \rtimes \langle 1|\cdots|n\rangle$ is the sole facet in F_1 satisfying $\sigma_0 \sim_{\{0\}} \sigma_1$.*

Proof. By the definition of F_0 and Proposition 4.5(i). □

Lemma 5.2. *Let $\sigma = I_d \rtimes \gamma_1 \rtimes \cdots \rtimes \gamma_m \in F_d$ and $\sigma' = I_{d'} \rtimes \gamma'_1 \rtimes \cdots \rtimes \gamma'_m \in F_{d'}$ $(1 \le d \le d' \le n)$ be distinct facets of $\mathcal{I}[\mathcal{IS}^m]_\delta$, where $\gamma_i = \langle A_{i,1}|\cdots|A_{i,r_i}|d{+}1|\cdots|n\rangle \in OSP_d$, and $\gamma'_i = \langle A'_{i,1}|\cdots|A'_{i,r'_i}|d'{+}1|\cdots|n\rangle \in OSP_{d'}$ for each i $(1 \le i \le m)$. Suppose $b \in [0,d']$ and $A = [0,d'] \setminus \{b\}$. Then, $\sigma \sim_A \sigma'$ holds if and only if either of the following conditions hold.*

[6] In topological arguments, some of the invariant properties such as OFUN, VALID, and KNOW are often left implicit, because they are readily obtained from the simplicial complex model. In contrast, we must make them explicit to promote logical reasoning, at every step of unfolding of the fixpoint.

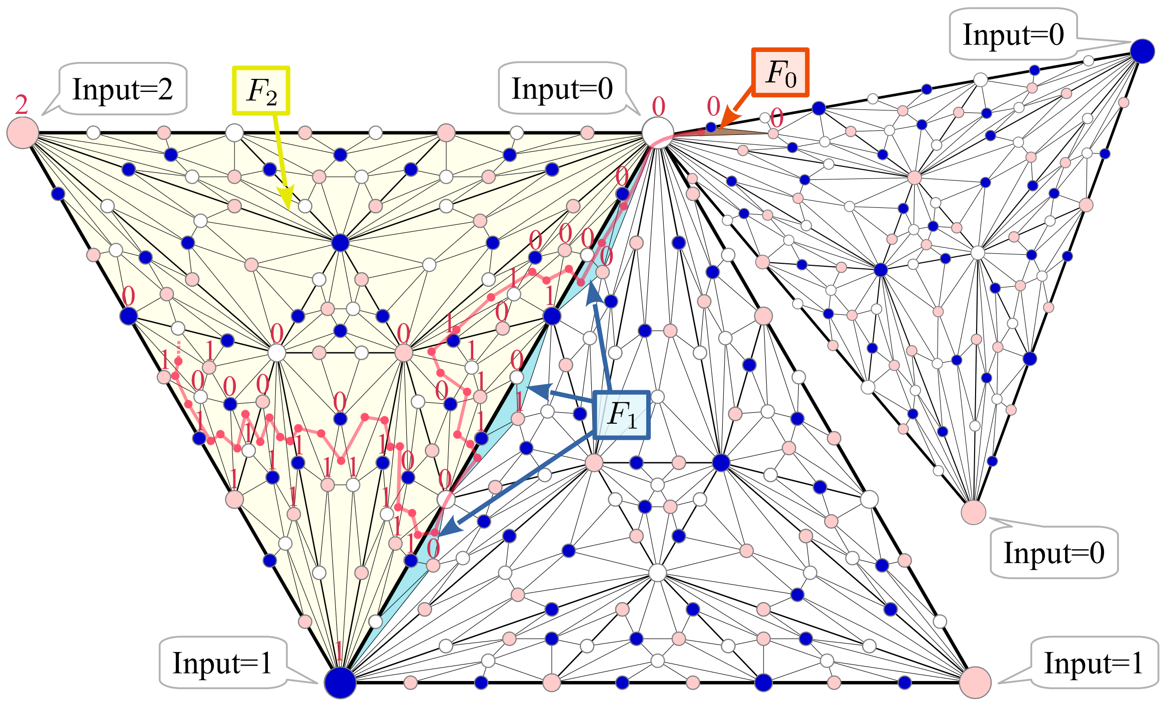

Figure 6: Collections of facets F_0, F_1, and F_2 that are contained in the product update model $\mathcal{I}[\mathcal{IS}^2]_\delta$. The numbers on the vertexes indicate the output values decided by δ. The solid red lines indicate an (incomplete) graph path over the facets that are connected by either of the relations $\frown_{\{0\}}$ and $\frown_{\{0,1\}}$.

(i) $d' = d+1$, $b = d'$, *and for every* i $(1 \le i \le m)$, $\gamma'_i = \gamma_i$ *and* $A'_{i,r'_i} = \{d'\}$.

(ii) $d' = d$ *and there exists* j $(1 \le j \le m)$ *and* $b \in [0,d]$ *such that* $\gamma'_j = \underline{\gamma_j}_A$ *and* $A_{j,r_j} \ne \{b\}$, $\gamma'_i = \gamma_i$ *for every* i $(1 \le i < j)$, *and* $A_{i,r_i} = \{b\}$ *for every* i $(j < i \le m)$.

Proof. The if direction holds by Proposition 4.5. To show the converse, assume $\sigma \sim_A \sigma'$. Since $\chi(I_d \cap I_{d'}) = [0,d]$, either $d' = d+1$ or $d' = d$. For the case $d' = d+1$, we obtain the condition (i) by applying Proposition 4.5(i) with $A = [0,d']$ and $b = d'$. For the case $d' = d$, the condition (ii) follows from Proposition 4.5(ii). □

We define a Sperner labeling η on the vertexes of $\mathcal{I}[\mathcal{IS}^m]_\delta$ as follows. For each facet $\sigma \in \mathsf{F}(\mathcal{I}[\mathcal{IS}^m]_\delta)$, let $\hat{\eta}_\sigma$ be a function defined by $\hat{\eta}_\sigma(a) = d$ iff $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma \models \mathsf{decide}^d_a$, which means that a vertex of color a contained in σ has an output value d. Then, we define a labeling function η on the vertexes of $\mathcal{I}[\mathcal{IS}^m]_\delta$ by $\eta(v) = d$ iff $\eta(v) = \hat{\eta}_\sigma(\chi(v))$ for some facet σ containing v. This labeling function η is well-defined: For any pair of facets σ and σ' with a common vertex $v \in \sigma \cap \sigma'$, $\hat{\eta}_\sigma(\chi(v)) = \hat{\eta}_{\sigma'}(\chi(v))$ follows from $\sigma \sim_{\chi(v)} \sigma'$.

The following presents a generalized definition of the graph that we constructed in Section 2.2 for a specific subdivision of a simplex.

Definition 2. *Let us define a family of irreflexive symmetric relations* $\frown_A$ *over* $\bigcup_{d=0}^k F_d$, *where* A *ranges over nonempty subsets of* Π, *as follows.*

- $\sigma \frown_A \sigma'$ *holds if and only if the following conditions are satisfied for some* $\sigma \in F_d$ $(0 \le d \le k)$ *and* $\sigma' \in F_e$ $(0 \le e \le k)$*:*

  (i) $\sigma \sim_A \sigma'$, *and* $\sigma \ne \sigma'$;
  (ii) $|A| = \max(d,e)$, $A \subseteq [0, \min(d,e)]$, *and* $|d-e| \le 1$;
  (iii) $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma \models \mathrm{DEC}_A$ *and* $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma' \models \mathrm{DEC}_A$.

The relation $\frown_A$ defines the conditions for a pair of facet σ and σ' being connected by a graph edge. The facets σ and σ' share a $(|A|-1)$-dimensional face whose vertexes are colored by A (condition (i)); σ and σ' are facets of $F_{|A|-1}$ or $F_{|A|}$ but not both of them are a facet of $F_{|A|-1}$ (condition (ii)); For every $d \in [0, |A|-1]$, both facets σ and σ' contain a vertex v such that $\chi(v) \in A$ and $\hat{\eta}_\sigma(\chi(v)) = d$, that is, $\hat{\eta}_\sigma(A) \supseteq [0, |A|-1]$ and $\hat{\eta}_{\sigma'}(A) \supseteq [0, |A|-1]$ (condition (iii)). If both σ and σ' are facets of $F_{|A|}$, they share a common face comprising vertexes of colors $[0, |A|] \setminus \{b\}$, for some $b \in [0, |A|]$; If one of σ and σ' is a facet of $F_{|A|}$ and the other is a facet of $F_{|A|-1}$, they share a common face, which is a facet of the subdivision of the boundary face $I_{|A|-1} \cap I_{|A|} = \{(i,i) \mid 0 \le i \le |A|-1\}$.

Fig. 6 shows a graph that spans over the facets in $\mathcal{I}[\mathcal{IS}^2]_\delta$, which is defined by a particular Sperner labeling on the vertexes.

Proposition 5.3. *Suppose $k \le n$ and $\sigma \in F_d$ for some d $(1 \le d \le k)$. If $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma \models \mathrm{OFUN} \wedge \mathrm{VALID} \wedge \mathrm{AGREE}_k \wedge \mathrm{KNOW}$, the number of facets $\sigma' \in \bigcup_{d=0}^{k} F_d$ satisfying $\sigma \frown_A \sigma'$ for some A is either 0 or 2.*

Proof. Suppose $\sigma \in F_d$ $(0 < d \le k)$, such that $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma \models \mathrm{OFUN} \wedge \mathrm{VALID} \wedge \mathrm{AGREE}_k \wedge \mathrm{KNOW}$. Then, the labeling function $\hat{\eta}_\sigma$ is a well-defined total function satisfying $\hat{\eta}_\sigma([0,n]) \subseteq [0,d]$ and $|\hat{\eta}_\sigma([0,n])| \le k$.

Let us assume that $\sigma \frown_A \tau$ holds for some τ and A. We show that σ has exactly two such distinct facets τ_1 and τ_2. That is, $\sigma \frown_{A_p} \tau_p$ holds for some A_p, for each $p = 1, 2$. In what follows, let $\sigma = I_d \rtimes \gamma_1 \rtimes \cdots \rtimes \gamma_m$, where $\gamma_i = \langle A_{i,1} | \cdots | A_{i,r_i} | d+1 | \cdots | n \rangle \in \mathit{OSP}_d$ for every i $(1 \le i \le m)$. Since we have assumed $\sigma \frown_A \tau$, $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma \models \mathrm{DEC}_A$ (i.e., $\hat{\eta}_\sigma(A) \supseteq [0, |A|-1]$), $A \subseteq [0,d]$, and $d \le |A| \le d+1$ holds. Then, either (a) $A = [0,d]$ and $\hat{\eta}_\sigma$ is a bijection on $[0,d]$ or (b) $A = [0,d] \setminus \{b\}$ for some $b \in [0,d]$, $\hat{\eta}_\sigma(A) \supseteq [0, d-1]$, and $\hat{\eta}_\sigma$ is not a bijection on $[0,d]$. In either case, it suffices to show that there are two distinct facets τ_1 and τ_2 satisfying the conditions (i) and (ii) of Definition 2. (Condition (iii) follows from $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma \models \mathrm{KNOW}$ and $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma \models \mathrm{DEC}_A$.)

Let us first consider the case (a). In this case, we have $d+1 = |A| = |\hat{\eta}_\sigma(A)| \le |\hat{\eta}_\sigma([0,n])| \le k$, i.e., $d \le k-1$. Define $\tau_1 = I_{d+1} \rtimes \gamma_1 \rtimes \cdots \rtimes \gamma_m$. By Lemma 5.2(i), τ_1 is the sole facet in F_{d+1} satisfying $\sigma \sim_{[0,d]} \tau_1$. For the sake of another facet τ_2, there are two cases to consider. Suppose $A_{i,r_i} = \{d\}$ for every i $(1 \le i \le m)$. By Lemma 5.2(i), $\tau_2 = I_{d-1} \rtimes \gamma_1 \rtimes \cdots \rtimes \gamma_m$ is the sole facet in F_{d-1} satisfying $\sigma \sim_{[0,d-1]} \tau_2$. Suppose the other case, in which there is the largest number j $(1 \le j \le m)$, such that $A_{j,r_j} \ne \{d\}$. Therein, by Lemma 5.2(ii), $\tau_2 = I_d \rtimes \gamma'_1 \rtimes \cdots \rtimes \gamma'_m$ is the sole facet in F_d satisfying $\sigma \sim_{[0,d-1]} \tau_2$, where $\gamma'_j = \underline{\gamma_j}_{[0,d-1]}$ and $\gamma'_i = \gamma_i$ for every i other than j. In either case, τ_2 is the sole facet satisfying $\sigma \sim_{[0,d-1]} \tau_2$ and τ_2 is different from τ_1. Therefore, exactly two facets τ_1 and τ_2 are related to σ.

Consider the other case (b). In this case, $\hat{\eta}_\sigma([0,d]) = [0, d-1]$. Hence, there must be exactly two distinct sets $B_1, B_2 \subseteq [0,d]$, such that $B_1 = [0,d] \setminus \{b_1\}$, $B_2 = [0,d] \setminus \{b_2\}$, and $\hat{\eta}_\sigma(B_1) = \hat{\eta}_\sigma(B_2) = [0, d-1]$, where $b_1, b_2 \in [0,d]$. For each $p = 1, 2$, we show that there exists a unique facet τ_p such that $\sigma \sim_{B_p} \tau_p$.

For each $p = 1, 2$, there are three cases to consider. Suppose $b_p = d$ and $A_{i,r_i} = \{d\}$ for every i $(1 \le i \le m)$. In this case, $\gamma_i \in \mathit{OSP}_{d-1}$ for every i $(1 \le i \le m)$ and hence $\tau_p = I_{d-1} \rtimes \gamma_1 \rtimes \cdots \rtimes \gamma_m \in F_{d-1}$ is a facet satisfying $\sigma \sim_{B_p} \tau_p$, by Lemma 5.2(i). Suppose $b_p \ne d$ and $A_{i,r_i} = \{d\}$ for every i $(1 \le i \le m)$. In this case, according to Lemma 5.2(ii), there exists a facet $\tau_p \in F_d$ satisfying $\sigma \sim_{B_p} \tau_p$. Suppose that there exists i $(1 \le i \le m)$, such that $A_{i,r_i} \ne \{d\}$. Let j be the largest number satisfying $A_{j,r_j} \ne \{d\}$ and define $\tau_p = I_d \rtimes \gamma'_1 \rtimes \cdots \rtimes \gamma'_m \in F_d$, where $\gamma'_j = \underline{\gamma_j}_{B_p}$ and $\gamma'_i = \gamma_i$ for every i other than j. Then, according to Lemma 5.2(ii), $\sigma \sim_{B_p} \tau_p$. In either case, $\tau_1 \ne \tau_2$ and $\sigma \sim_{B_p} \tau_p$ holds for each $p = 1, 2$. □

Theorem 5.4. *Let $\mathcal{I}[\mathcal{SA}_k]$ and $\mathcal{I}[\mathcal{IS}^m]$ be the product update models of the k-set agreement task and the m-iterated immediate snapshot protocol, respectively. If $k < n+1$, where $n+1$ the number of processes in the system, there is no morphism δ from $\mathcal{I}[\mathcal{IS}^m]$ to $\mathcal{I}[\mathcal{SA}_k]$ satisfying $\pi_\mathcal{I} \circ \delta = \pi_\mathcal{I}$, meaning that the k-set agreement task is not solvable by the multiple-round immediate snapshot protocol.*

Proof. Suppose, by contradiction, that there exists such a morphism δ. Let $\sigma_0 = I_0 \rtimes \langle 0|1|\cdots|n \rangle$. By Lemma 5.1, σ_0 is the sole facet of F_0 and $\sigma_1 = I_1 \rtimes \langle 0|1|\cdots|n \rangle$ is the sole facet in F_1 satisfying $\sigma_0 \sim_{\{0\}} \sigma_1$.

Since Φ_k is valid in $\widehat{\mathcal{I}[\mathcal{SA}_k]}$, it is also valid in $\mathcal{I}[\mathcal{IS}^m]_\delta$ by Theorem 3.1. We obtain $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma_0 \models \mathrm{OFUN} \wedge \mathrm{VALID} \wedge \mathrm{KNOW} \wedge (\mathrm{DEC}_{\{0\}} \Rightarrow \mathrm{D}_{\{0\}}(\mathrm{AGREE}_k \wedge \Phi_k))$, by unfolding the fixpoint $\nu Z.\cdots$. Furthermore, since $\pi_{\mathcal{I}}(\delta(\sigma_0)) = \pi_{\mathcal{I}}(\sigma_0) = I_0$ and $\widehat{\mathcal{I}[\mathcal{SA}_k]}, \delta(\sigma_0) \models \bigwedge_{a\in\Pi} \mathsf{input}_a^0$, we have $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma_0 \models \bigwedge_{a\in\Pi} \mathsf{input}_a^0$ by Theorem 3.1. By this and $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma_0 \models \mathrm{OFUN} \wedge \mathrm{VALID}$, we obtain $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma_0 \models \mathrm{DEC}_{\{0\}}$. Hence, $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma_0 \models \mathrm{DEC}_{\{0\}} \wedge \mathrm{KNOW} \wedge \mathrm{D}_{\{0\}}(\mathrm{AGREE}_k \wedge \Phi_k)$. Since $\sigma_0 \sim_{\{0\}} \sigma_1$, we have $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma_1 \models \mathrm{DEC}_{\{0\}} \wedge \mathrm{AGREE}_k \wedge \Phi_k$. Therefore, $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma_1 \models \mathrm{AGREE}_k \wedge \mathrm{OFUN} \wedge \mathrm{VALID} \wedge \mathrm{KNOW} \wedge \Phi_k$ holds. Furthermore, $\sigma_0 \frown_{\{0\}} \sigma_1$ holds, because $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma_i \models \mathrm{DEC}_{\{0\}}$ for $i = 0, 1$.

Let us show that, for every $q \geq 1$, there exists a path of facets $\sigma_0 \frown_{A_0} \sigma_1 \frown_{A_1} \cdots \frown_{A_{q-1}} \sigma_q$ such that $\sigma_0, \sigma_1, \ldots, \sigma_q$ are pairwisely distinct facets in $\bigcup_{i=0}^{k} F_k$ and also $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma_i \models \mathrm{OFUN} \wedge \mathrm{VALID} \wedge \mathrm{KNOW} \wedge \mathrm{AGREE}_k \wedge \Phi_k$ holds for every σ_i. We show this by induction on q. We have already examined the case $q = 1$. Suppose the claim holds up to q ($q > 0$), that is, $\sigma_{q-1} \frown_{A_{q-1}} \sigma_q$ and $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma_q \models \mathrm{OFUN} \wedge \mathrm{VALID} \wedge \mathrm{KNOW} \wedge \mathrm{AGREE}_k \wedge \Phi_k$. By Proposition 5.3, σ_q must have a facet σ_{q+1}, other than σ_{q-1}, such that $\sigma_q \frown_{A_q} \sigma_{q+1}$ for an appropriate A_q. The relation $\sigma_q \frown_{A_q} \sigma_{q+1}$ implies $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma_q \models \mathrm{DEC}_{A_q}$ and hence we obtain $\mathcal{I}[\mathcal{IS}^m]_\delta, \sigma_{q+1} \models \mathrm{OFUN} \wedge \mathrm{VALID} \wedge \mathrm{KNOW} \wedge \mathrm{AGREE}_k \wedge \Phi_k$ by unfolding the fixpoint.

Finally, we show that every σ_q ($q > 0$) is distinct from any other preceding facets, by induction on q. The first three facets σ_0, σ_1, and σ_2 are pairwisely distinct, because $\sigma_0 \frown_{A_0} \sigma_1 \frown_{A_1} \sigma_2$ and $\sigma_0 \neq \sigma_2$. We show that $\sigma_{q+1} \neq \sigma_j$ for any $q \geq 2$ and j ($0 \leq j < q$). Suppose $\sigma_{q+1} = \sigma_0$. Since σ_1 is the sole facet such that $\sigma_0 \frown_A \sigma_1$ for some A, $\sigma_0 \frown_{A_{q-1}} \sigma_q$ and $\sigma_0 \frown_{A_0} \sigma_1$ implies $\sigma_q = \sigma_1$, which contradicts the induction hypothesis. Suppose $\sigma_{q+1} = \sigma_j$ for some $0 < j < q$. This implies $\sigma_{j-1} \frown_{A_{j-1}} \sigma_j \frown_{A_j} \sigma_{j+1}$ and $\sigma_q \frown_{A_q} \sigma_j$. By Proposition 5.3, σ_q must be either σ_{j-1} or σ_{j+1}, which contradicts the induction hypothesis.

Therefore, there exists an ever-lasting path of distinct facets $\sigma_0 \frown_{A_0} \sigma_1 \frown_{A_1} \sigma_2 \frown_{A_2} \cdots$. This contradicts the fact that $\bigcup_{i=0}^{k} F_k$ is a finite set. Hence, there exists no morphism δ. $\square$

6 Set Agreement in k-Concurrency Model

The unsolvability argument in Section 5 can also be applied to a submodel of $\mathcal{I}[\mathcal{IS}^m]$, if the submodel subsumes the set of facets $\bigcup_{d=1}^{k} F_d$. In this section, we examine k-concurrency model [11], as an instance of such a submodel.

In the *k-concurrency model*, at most k processes are allowed to be concurrently active in their execution. As shown in [11], the k-concurrent execution of the 2-iterated immediate snapshot protocol precisely captures the ability for solving the k-set agreement task, and the protocol complex of the k-concurrency model is given by $\mathcal{R}_k$, a subcomplex of $\mathcal{IS}^2$ that consists of only those facets produced by k-concurrent executions. The exact correspondence between the k-set agreement task and the protocol complex $\mathcal{R}_k$ implies that the k-concurrency model cannot solve the ℓ-set agreement task, if $\ell < k$. In what follows, we demonstrate this unsolvability result, using the product update model for k-concurrency.

Following [11], we define the protocol complex $\mathcal{R}_k$ of the k-concurrency model with a formal combinatorial description given below. For an ordered set partition $\gamma = \langle A_1 | A_2 | \cdots | A_r \rangle$ of Π and $a \in \Pi$, we define $\mathit{view}_a(\gamma) = \bigcup_{i=1}^{t} A_i$, where $a \in A_t$. Further, for a facet $X \rtimes \gamma_1 \rtimes \gamma_2 \in \mathsf{F}(\mathcal{IS}^2)$ and a vertex of color $a \in \Pi$, we define the *carrier set* by $\mathit{carrier}_a(X \rtimes \gamma_1 \rtimes \gamma_2) = \bigcup_{b \in \mathit{view}_a(\gamma_2)} \mathit{view}_b(\gamma_1)$. For each facet $\sigma \in \mathsf{F}(\mathcal{IS}^2)$, the *contention set* is defined by $\mathit{Cont}(\sigma) = \{A \subseteq \Pi \mid \mathit{carrier}_a(\sigma) = \mathit{carrier}_b(\sigma)$ for every $a, b \in A\}$.[7] Then the k-concurrency model $\mathcal{R}_k$ is a subcomplex of $\mathcal{IS}^2$ with a restricted set of facets $\mathsf{F}(\mathcal{R}_k) = \{\sigma \in \mathsf{F}(\mathcal{IS}^2) \mid |A| \leq k \text{ for every } A \in \mathit{Cont}(\sigma)\}$.

Figure 7 illustrates the 2-concurrency model $\mathcal{R}_2$ for a 3-process system. The product update model $I[\mathcal{R}_2]$ is similarly defined as in Section 4 and $I[\mathcal{R}_2]$ is geometrically isomorphic to $\mathcal{R}_2$. Observe that this model contains the set of facets $F_0 \cup F_1$ of Fig. 6. This containment property holds in general: $I[\mathcal{R}_k]$ contains $\bigcup_{d=0}^{k-1} F_d$, if $k \leq n$. Suppose $\sigma = I_d \rtimes \gamma_1 \rtimes \gamma_2 \in F_d$ ($0 \leq d \leq k-1$), where $\gamma_1, \gamma_2 \in \mathit{OSP}_d$. Then, $\mathit{carrier}_a(\sigma) = [0, a]$ if $a \in [d+1, n]$, while $\mathit{carrier}_a(\sigma) \subseteq [0, d]$ if $a \in [0, d]$. This implies $|A| \leq d + 1 \leq k$ for every $A \in \mathit{Cont}(\sigma)$ and therefore $\sigma \in \mathsf{F}(I[\mathcal{R}_k])$. This means that $I[\mathcal{R}_k]$ contains all facets that are relevant for the proof of unsolvability of the $k-1$ (or fewer) set agreement task. We obtain the following theorem by the same argument presented in Section 5.

[7] Here we define the contention set differently from the original one [11] for clarity, albeit the two definitions are equivalent.

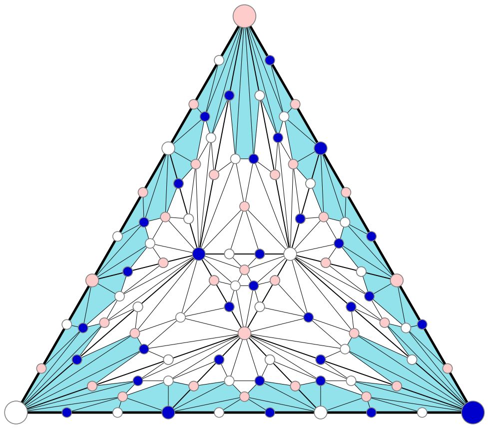

Figure 7: 2-concurrency model R2 for a 3-process system, a subcomplex of $\mathcal{IS}^2$ that consists of the shaded facets

Theorem 6.1. *Let $\mathcal{I}[\mathcal{SA}_\ell]$ be the product update model for the ℓ-set agreement task, $\mathcal{R}_k$ be the k-concurrency model. Then, there is no morphism δ from $I[\mathcal{R}_k]$ to $\mathcal{I}[\mathcal{SA}_\ell]$ satisfying $\pi_{\mathcal{I}} \circ \delta = \pi_{\mathcal{I}}$, if $\ell < k$.*

7 Conclusion

We presented a logical obstruction to the k-set agreement task in the language of epistemic μ-calculus. Extended with distributed knowledge modalities and propositional greatest fixpoints, epistemic μ-calculus can describe higher dimensional properties of the topological model of distributed computing in the formal language of logic. This allowed us to define the logical obstruction in epistemic μ-calculus, which contradicts Sperner's lemma and thereby proves that the k-set agreement task is not solvable by the multiple-round immediate snapshot protocol. This result extends the existing proposals of logical obstructions [13, 27, 36], which are limited to the consensus task or the single-round immediate snapshot protocol. Furthermore, we showed that the unsolvability of set agreement task is entailed for the k-concurrency model, using the same logical obstruction formula.

We showed that an equivalent of Sperner's lemma can be expressed in the language of epistemic μ-calculus. This suggests that the logical method could provide an alternative platform for demonstrating task unsolvability. However, further investigation is needed to devise concrete logical obstructions to other distributed tasks. As pointed out in [32, 23], it heavily depends on the language of the logic whether a logical obstruction to a particular distributed task is definable. Furthermore, the unsolvability results found in the literature [12, 29, 7, 18] often resort to sophisticated results from combinatorial topology, such as the Index lemma [17] and Nerve lemma [25]. To devise logical obstructions corresponding these results, we would need to express these lemmas or their equivalents in an appropriate language of logic. This remains as a problem for future investigation, and it would contribute to a deeper understanding of the interplay between combinatorial topology and epistemic logic.

Acknowledgments

The author is supported by JSPS KAKENHI Grant Number 20K11678.

References

[1] Dan Alistarh, James Aspnes, Faith Ellen, Rati Gelashvili, and Leqi Zhu. Why extension-based proofs fail. *SIAM Journal on Computing*, 52(4):913–944, 2023.

[2] Alexandru Baltag, Lawrence S. Moss, and Sławomir Solecki. The logic of public announcements, common knowledge, and private suspicions. In Horacio Arló-Costa, Vincent F. Hendricks, and

Johan van Benthem, editors, *Readings in Formal Epistemology: Sourcebook*, chapter 39, pages 773–812. Springer, 2016.

[3] Alexandru Baltag and Sonja Smets. Learning what others know. In *LPAR-23: 23rd International Conference on Logic for Programming, Artificial Intelligence and Reasoning*, volume 73 of *EPiC Series in Computing*, pages 90–119, 2020.

[4] Elizabeth Borowsky and Eli Gafni. Generalized FLP impossibility result for t-resilient asynchronous computations. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing, STOC*, pages 91–100. ACM, 1993.

[5] Elizabeth Borowsky and Eli Gafni. Immediate atomic snapshots and fast renaming (extended abstract). In *Proceedings of the 12th Annual ACM Symposium on Principles of Distributed Computing*, pages 41–51. ACM, 1993.

[6] Julian C. Bradfield and Colin Stirling. Modal mu-calculi. In Patrick Blackburn, J. F. A. K. van Benthem, and Frank Wolter, editors, *Handbook of Modal Logic*, volume 3 of *Studies in logic and practical reasoning*, pages 721–756. North-Holland, 2007.

[7] Armando Castañeda, Pierre Fraigniaud, Ami Paz, Sergio Rajsbaum, Matthieu Roy, and Corentin Travers. A topological perspective on distributed network algorithms. *Theor. Comput. Sci.*, 849:121–137, 2021.

[8] Mark de Longueville. *A Course in Topological Combinatorics*. Universitext. Springer, 2013.

[9] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.

[10] Michael J. Fischer, Nancy A. Lynch, and Mike Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, 1985.

[11] Eli Gafni, Yuan He, Petr Kuznetsov, and Thibault Rieutord. Read-write memory and k-set consensus as an affine task. In *20th International Conference on Principles of Distributed Systems (OPODIS 2016)*, pages 6:1–6:17, 2016.

[12] Éric Goubault, Marijana Lazic, Jérémy Ledent, and Sergio Rajsbaum. Wait-free solvability of equality negation tasks. In *33rd International Symposium on Distributed Computing, DISC 2019*, volume 146 of *LIPIcs*, pages 21:1–21:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[13] Éric Goubault, Jérémy Ledent, and Sergio Rajsbaum. A simplicial complex model for dynamic epistemic logic to study distributed task computability. *Information and Computation*, 278:104597, 2021. An earlier version appeared in Proc. of 9th International Symposium on Games, Automata, Logics, and Formal Verification, GandALF 2018.

[14] Rachid Guerraoui and Petr Kuznetsov. *Algorithms for Concurrent Systems*. EPFL press, 2018.

[15] Joseph Y. Halpern and Yoram Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, 1990.

[16] Joseph Y. Halpern and Moshe Y. Vardi. The complexity of reasoning about knowledge and time. I. lower bounds. *Journal of Computer and System Sciences*, 38(1):195–237, 1989.

[17] Michael Henle. *A Combinatorial Introduction to Topology*. Dover, 1983.

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

[19] Maurice Herlihy and Sergio Rajsbaum. Set consensus using arbitrary objects (preliminary version). In *Proceedings of the 13th Annual ACM Symposium on Principles of Distributed Computing*, pages 324–333. ACM, 1994.

[20] Maurice Herlihy and Sergio Rajsbaum. Algebraic spans. *Mathematical Structures in Computer Science*, 10(4):549–573, 2000.

[21] Maurice Herlihy and Nir Shavit. The topological structure of asynchronous computability. *Journal of the ACM*, 46(6):858–923, 1999.

[22] Jaakko Hintikka. *Knowledge and belief: An introduction to the logic of the two notions*. Cornell University Press, 1962.

[23] Sou Hoshino. Determining existence of logical obstructions to the distributed task solvability, 2022. arXiv:2203.05153. URL: `https://arxiv.org/abs/2203.05153/`.

[24] Dexter Kozen. Results on the propositional μ-calculus. *Theoretical Computer Science*, 27(3):333–354, 1983.

[25] Dmitry Kozlov. *Combinatorial Algebraic Topology*. Springer, 2008.

[26] Dmitry N. Kozlov. Chromatic subdivision of a simplicial complex. *Homology, Homotopy and Applications*, 14(2):197–209, 2012.

[27] Yutaro Nishida. Impossibility of k-set agreement via dynamic epistemic logic (in Japanese). In *Algebraic system, Logic, Language and Related Areas in Computer Sciences II*, volume 2188 of *RIMS Kôkyûroku*, pages 96–105, Feb. 2020. URL: `http://hdl.handle.net/2433/265621`.

[28] Susumu Nishimura. Schlegel diagram and optimizable immediate snapshot protocol. In *21st International Conference on Principles of Distributed Systems, OPODIS 2017*, pages 22:1–22:16, 2017.

[29] Vikram Saraph, Maurice Herlihy, and Eli Gafni. Asynchronous computability theorems for t-resilient systems. In *Distributed Computing, 30th International Symposium, DISC 2016*, pages 428–441, 2016.

[30] Nikolay V. Shilov and Natalya Olegovna Garanina. Model checking knowledge and fixpoints. In *Fixed Points in Computer Science, FICS 2002*, volume NS-02-2 of *BRICS Notes Series*, pages 25–39, 2002.

[31] Emanuel Sperner. Neuer beweis für die invarianz der dimensionszahl und des gebietes. *Abhandlungen Hamburg*, 6:265–272, 1928.

[32] Hans van Ditmarsch, Éric Goubault, Marijana Lazic, Jérémy Ledent, and Sergio Rajsbaum. A dynamic epistemic logic analysis of equality negation and other epistemic covering tasks. *Journal of Logical and Algebraic Methods in Programming*, 121:100662, 2021.

[33] Hans van Ditmarsch, Eric Goubault, Jérémy Ledent, and Sergio Rajsbaum. Knowledge and simplicial complexes. *CoRR*, abs/2002.08863, 2020. URL: `https://arxiv.org/abs/2002.08863`.

[34] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer, 2008.

[35] Diego A. Velázquez, Armando Castañeda, and David A. Rosenblueth. Communication pattern models: An extension of action models for dynamic-network distributed systems. In *Proceedings Eighteenth Conference on Theoretical Aspects of Rationality and Knowledge, TARK 2021*, volume 335 of *EPTCS*, pages 307–321, 2021.

[36] Koki Yagi and Susumu Nishimura. Logical obstruction to set agreement tasks for superset-closed adversaries, 2020. arXiv:2011.13630. URL: `https://arxiv.org/abs/2011.13630/`.